

Characterization of Single-Electron Tunneling Transistors for Designing Low-Power Embedded Systems

Changyun Zhu, *Student Member, IEEE*, Zhengyu (Peter) Gu, *Student Member, IEEE*, Robert P. Dick, *Member, IEEE*, Li Shang, *Member, IEEE*, Robert Knobel, *Member, IEEE*

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 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. 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.

Evaluating the use of SETs in large-scale digital systems requires novel architectural and circuit design. SET-based design imposes numerous challenges resulting from low driving strength, relatively large static power consumption, and the presence of reliability problems resulting from random background charge effects. We propose a fault-tolerant, hybrid SET/CMOS, reconfigurable architecture, named IceFlex, that can be tailored to specific requirements and allows trade-offs among power consumption, performance requirements, operation temperature, fabrication cost, and reliability. Using IceFlex as a testbed, we characterize the benefits and limitations of SETs in embedded system designs. In particular, we focus on the use of SETs in room-temperature ultra-low-power embedded systems such as wireless sensor network nodes. We also consider higher-performance applications such as multimedia consumer electronics. We see this work as a first step in determining the potential of ultra-low-power embedded system design using SETs.

Index Terms—Single-electron tunneling transistor, Low-power design, Reconfigurable architecture, Embedded system, High-performance computing, Nanotechnology, Reliability.

I. INTRODUCTION

Energy consumption and thermal issues are now central issues in electronic system design. In high-performance applications, temperature affects integration density, performance,

reliability, power consumption, and cost. For battery-powered embedded systems, power consumption determines 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 primary motivations for replacing vacuum tubes with semiconductor devices in the 1960s and replacing bipolar junction transistors with CMOS in the 1990s. Although CMOS is the mainstream fabrication technology used today, as IC and system integration further increase, it will reach fabrication, power consumption, and thermal limits; it may soon be time for another transition to a dramatically different technology.

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]. The International Technology Roadmap for Semiconductors projects that SETs have the potential to achieve the lowest projected energy per switching event of any known device (1×10^{-18} J) [4]. However, their use poses unique architectural, circuit design, and fabrication challenges. For example, SETs are susceptible to reliability problems caused by random background offset charges. 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 using CMOS and threshold logic. Although the fabrication of SETs capable of operating at low temperatures is now common, feature sizes of only a few nanometers are required for room-temperature operation, making fabrication challenging.

A. Past Work

After their discovery in the 1980s [5, 6], there has been extensive research on fabrication, design, and modeling of SETs [3]. SET fabrication and use in high-sensitivity amplifiers at cryogenic temperatures has been the main research focus [7]. SETs and simple circuits with a variety of structures were proposed and fabricated using different methods and materials [8–10]. Recently, researchers have fabricated SETs that operate at room-temperature [11–13]. Various SET-based circuit applications, such as logic [14–17] and memory [18–20] have been developed. This work provides a promising start for SET circuit design. However, these articles did not provide an architectural evaluation. We do not claim to have improved

C. Zhu is with the Department of Electrical and Computer Engineering, Queen's University, Kingston, ON K7L 3N6, Canada.

Z. Gu is with Synplicity Inc, Sunnyvale, CA 94086, U.S.A.

R. Dick is with the Department of Electrical Engineering and Computer Science, Northwestern University, Evanston, IL 60208, U.S.A.

L. Shang is with the Department of Electrical and Computer Engineering, University of Colorado at Boulder, Boulder, CO 80305, U.S.A.

R. Knobel is with the Department of Physics, Queen's University, Kingston, ON K7L 3N6, Canada.

This work is supported in part by NSERC Discovery Grants #388694-01 and #298185-04, in part by the NSF under awards CNS-0347941, CNS-0721978, and CCF-0702761, and in part by the SRC under award 2007-HJ-1593.

the performance of SET-based logic gates. Instead, we are the first to develop the modules necessary to support architectural design and synthesis and evaluate the architectural performance and power consumption implications of using SETs. They demonstrate orders of magnitude improvement in power consumption and energy efficiency compared to CMOS.

Research on SET modeling and simulation has been an active area. Monte Carlo simulation has been widely used to model SETs. SIMON [21] and MOSES [22] are the two most popular SET simulators. However, they are too slow for analysis of large circuits. Uchida et al. proposed an analytical SET model and incorporated it into SPICE [23]. Recently, Inokawa et al. extended this model to a more general form to include asymmetric SETs [24]. Mahapatra et al. propose a simulation framework for hybrid SET/CMOS circuit design and analysis [25]. Their model for SET behavior is similar to that of Uchida et al. These compact modeling techniques are efficient enough for use in SET circuit design and analysis and closely match Monte Carlo simulation results.

Significant challenges still remain for large-scale integration of SETs and for room-temperature operation. SETs that operate reliably at room temperature have critical dimensions of $\sim 1\text{--}10\text{ nm}$. They are challenging to fabricate using current top-down lithographic techniques. However, several exciting advances make the evaluation of architectures for high-density logic based on SETs worthwhile. Scanning-probe microscopes can be used to create devices smaller than those using conventional lithography [11]. Continual progress has been made on bottom-up nano-fabrication techniques, where chemical techniques are used to make individual molecules with useful electronic properties. Molecular quantum dots [26] can display SET behavior. Larger structures, such as carbon nanotubes and nanowires, can act as SETs [10]. These bottom-up techniques can create structures supporting room-temperature SET operation. However, more research is needed in order to integrate individual devices into large-scale circuits. Very recent advances in graphene [27] devices show promise for SETs. Reliable methods for cooling to very low temperatures without supplies of liquid helium or nitrogen are also becoming more common [28]. For high-performance computing, the added complexity of operating at cryogenic temperatures may not be a limiting factor. Similarly, cryogenic temperatures are readily attained using passive methods in outer space.

B. Contributions

In this paper, we explore the potential use of SETs in low-power embedded systems. In order to take advantage of the power efficiency of SETs, it is critical to bring SET-based design to the system level, characterize the impacts of SETs on system design metrics, and evaluate the benefits and limitations of SETs. 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




Fig. 1. SET structure and schematic.

of SETs while working within their limitations, we propose a fault-tolerant, reconfigurable, hybrid SET/CMOS based architecture called IceFlex. IceFlex is regular and cell-based. It is reconfigurable, permitting compensation for fabrication defects. It incorporates flexible, modular circuits to enable tolerance of run-time faults. In addition to compensating for the weaknesses of SETs, IceFlex exploits their strengths, e.g., we develop a two-SET design to implement Boolean functions that are not linearly separable.

We tailor IceFlex to both high-performance and battery-powered embedded systems and characterize its energy efficiency, performance, and power consumption by using it for a number of instruction processors and application-specific cores. Compared to CMOS-based designs, IceFlex improves energy efficiency by two orders of magnitude for both battery-powered and high-performance applications, while maintaining good performance. However, our results also indicate great challenges to the use of SET-based designs in portable embedded systems. Their use will either require advances in the compact cooling technologies or the fabrication of features with sizes approaching physical limits.

II. SET MODELING

In this section, we introduce the physical properties of SETs, and discuss SET analytical device modeling.

A. SET Basics

The operation of a single-electron tunneling device is governed by the Coulomb charging effect. As shown in Figure 1, a single-electron tunneling device consists of a nanometer-scale conductive island embedded in an insulating material. Electrons travel between the island, source (S), and drain (D) through thin insulating 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 large devices, this change in potential is negligible due to the high island capacitance C_Σ . However, for nanometer-scale islands, C_Σ is much smaller. As a result, the electrostatic energy change due to the addition or removal of a single electron can be larger than the thermal energy, particularly at low temperatures.

Changes to SET island potential results in an energy gap at the 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




Fig. 2. SET Coulomb oscillation ($C_g = 3.2$ aF, $C_s = C_d = 1.0$ aF, and $R_s = R_d = 10$ M Ω).

TABLE I
ISLAND SIZE ESTIMATION

Temperature (K)	$C_\Sigma = e^2/(10k_B T)$		$C_\Sigma = e^2/(40k_B T)$	
	Island capacitance (aF)	Island diameter (nm)	Island capacitance (aF)	Island diameter (nm)
40	4.65	52.48	1.16	13.12
77	2.41	27.26	0.60	6.82
103	1.80	20.38	0.45	5.10
120	1.55	17.49	0.39	4.37
200	0.93	10.50	0.23	2.62
250	0.74	8.40	0.19	2.10
300	0.62	7.00	0.15	1.75

Assuming disc capacitor model ($C_\Sigma = 8\epsilon r$). One side of island embedded in silicon dioxide. Other side exposed to Nitrogen.

by changing the voltage of a conductor capacitively coupled to the island, thereby turning tunneling on and off. Although their transfer functions differ significantly from those of CMOS transistors, with careful circuit design, SETs can be used to realize logic functions using circuits analogous to CMOS, or using radically different design techniques [3].

As shown in Figure 1, a SET typically has four terminals. The source and drain terminals (S , D) serve as electron reservoirs. When the SET is turned on, electrons tunnel from one terminal, through the junction, to the conductive island. They then tunnel through the other junction to the other terminal. Each tunneling junction is modeled as a resistor (R_S or R_D) and a capacitor (C_S or C_D) in parallel. A gate terminal (G), with coupling capacitance C_G , controls the transport of electrons. A SET may also contain an optional second gate terminal (G_2), which is generally used to tune SET V_{GS} bias. The Coulomb blockade effect is maximized when $V_{GS} = me/C_G$, where $m = 0, \pm 1, \pm 2, \dots$ [29] because, at these voltages, the system is in a minimal-energy state when an integer number of electrons are present on the island. Any single tunneling event between island and either source or drain would move the system from this state. The Coulomb blockade effect vanishes when $m = \pm 1/2, \pm 3/2, \dots$, i.e., when m is a half-integer value because, at these voltages, the system is in a minimal-energy state when a half-integer number of electrons are present on the island. In this case, a single tunneling event does not move the system from a minimum energy state. Electrons can therefore tunnel, in single-file, through the island as determined by V_{DS} . The I-V curve of a SET is shown in Figure 2; drain current changes as a function of the gate voltage, with a period of e/C_g . The periodic changes are called Coulomb Oscillations.

In order to observe the Coulomb blockade effect, the following constraints must be satisfied.

1) Since thermal fluctuations can suppress the Coulomb Blockade effect, the electrostatic charging energy, e^2/C_Σ , must be much greater than $k_B T$, where k_B is Boltzmann's constant and T is the temperature. In order to ensure reliability, $e^2/C_\Sigma \geq 10k_B T$ or the more conservative $e^2/C_\Sigma \geq 40k_B T$ constraint is enforced. These equations imply that the maximum allowed island capacitance is inversely proportional to temperature. At room temperature, an island capacitance below 1 aF is required. Island capacitance is a function of island size. As shown in Table I, room-temperature operation requires an island size in the nanometer range, making fabrication challenging. At present, the smallest island capacitance of a fabricated device is around 0.15 aF [12].

2) To observe single-electron charging effects, electrons must be confined to the island, which requires that the junction resistance be higher than the quantum resistance, i.e., $R_S, R_D > h/e^2$, $h/e^2 = 25.8$ k Ω , where h is Planck's constant. Therefore, SETs have high resistances and low driving currents.

In order to operate voltage-state logic, SETs must exhibit voltage gain. The low-temperature voltage gain is equal to the gate capacitance divided by the sum of the junction capacitances: $G = C_G/(C_S + C_D)$. Achieving this gain requires low tunneling junction capacitances. It also requires close coupling of gate and island without a large increase in the total island capacitance. High gain has only been demonstrated for a few devices and has required operation at low temperatures [30, 31]. However, further advances in nanofabrication may overcome this limitation.

B. Random Background Charge Effects

Constant background charge effects have been a persistent problem for SETs. Charges near the SET island influence its equilibrium state [32]. Although the resulting voltage offsets can be compensated for with a biased second gate terminal, the required bias is unknown until fabrication. Worse yet, some devices are affected by random background charge effects, which result in run-time voltage fluctuations.

It is the tentative consensus of the research community that random background charge effects are caused by multiple, closely-spaced charge traps near the island, among which charge carriers tunnel. This produces run-time variation in gate bias, and may cause logic errors. Much work has been done to understand the nature and density of these defects [33–35]. Most SETs have been fabricated with aluminum islands. Some researchers have attempted to eliminate random background charge effects by fabricating SETs with alternative island materials such as silicon, based on the thesis that the use of non-crystalline, non-stoichiometric aluminum oxide junctions in conventional SETs leads to numerous charge-trapping defects. Silicon island based devices have high immunity to random background charge noise, with operation unchanged over several weeks [36]. However, random background charge effects remain the main source of run-time reliability problems for most SET designs. In this work, we describe a reconfigurable

architecture that provides architectural resistance to the effects of random background charges.

C. SET Modeling

Circuit design involves extensive simulation. Despite their accuracy, Monte Carlo methods are too slow for large-scale circuit analysis. We build upon the SET analytical model developed by Inokawa et al. [24], which has been incorporated into SPICE. Combined with MOS transistor models, it provides an efficient and accurate simulation solution for hybrid SET/CMOS circuits. Inokawa's model ignores random background charge effects and multi-gate effects. We incorporate these effects into the model.

The I–V characteristics of a SET with island charge equal to n or $n + 1$ electrons follow:

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

where

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

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

$$r = \frac{R_D - R_S}{R_D + R_S}, \quad R_T = \frac{2}{\frac{1}{R_S} + \frac{1}{R_D}} \quad (4)$$

$$C_\Sigma = C_S + C_D + \sum C_{G_i} \quad (5)$$

In this model, $\frac{2 \sum C_{G_i} V_{GS_i}}{e}$ models the Coulomb charging effects of the multiple gate terminals. ζ is a real number that characterizes the random background charge effect.

This compact model is derived based on the steady-state master equation, which is not directly applicable to transient circuit analysis. However, when used in circuits, SETs are connected by metal wires. Based on existing fabrication processes, the capacitance of local interconnect is at least two orders of magnitude higher than SET island capacitance, thereby eliminating inter-SET Coulomb interaction. The independence of SETs enables the use of quasi-steady-state analysis [24, 37].

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

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

Our evaluation of the architectural constraints imposed by SETs led to four main conclusions.

- 1) 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 improve reliability.
- 2) 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.
- 3) SETs have poor driving strength; this must be remedied, especially when driving global interconnect.
- 4) SETs have the ability to efficiently implement some functions that are inefficient using BJTs, CMOS logic, or threshold logic, e.g., non-linearly-separable functions. SET-based architectures should exploit such special properties.

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 two design configurations, each of which is shown in Table II. One targets high-performance embedded applications such as multimedia consumer electronics and one targets ultra-low-power embedded applications such as wireless sensor networks.

1) *Temperature*: IceFlex was evaluated at seven temperature settings (see Table II). IceFlex is a hybrid SET/CMOS design; the temperature range starts at 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 as 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.

2) *Capacitance*: To observe well-defined Coulomb blockade effects, electron charging energy must be higher than the thermal energy, i.e., $\frac{e^2}{C_\Sigma} \geq 10k_B T$ or $\frac{e^2}{C_\Sigma} \geq 40k_B T$, where k_B is Boltzmann's constant and T is temperature. At room temperature, this constraint requires an island capacitance below 1 aF, making fabrication challenging but possible [12]. 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 $C_\Sigma \leq e^2 / (10k_B T)$, $C_\Sigma \leq e^2 / (40k_B T)$ and $G = 1.5$, the maximum allowed gate and junction capacitances are derived and shown in the ‘‘Low power, Capacitance’’ columns of Table II.

The maximal allowed capacitance decreases with increasing temperature. However, fabricating SETs with low gate capacitance is challenging. We assume the capacitances at 300 K are the minimum allowed. Given $\frac{e^2}{C_\Sigma} \geq 10k_B T$, for high-performance applications, these minimal gate and junction capacitances are used at all the temperature settings and shown in the corresponding ‘‘High Performance, Capacitance’’ columns of Table II. Given $\frac{e^2}{C_\Sigma} \geq 40k_B T$, which requires very low SET

capacitance at room temperature, $C_G = 0.09$ aF. This makes fabrication very challenging. Due to fabrication concerns, for high-performance design, the capacitance and voltage are determined at the appropriate operation temperature, instead of room temperature.

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} ; we use this range. Once the range of V_{GS} is known, a V_{SS} of $-e/(4C_G)$ and a V_{DD} of $e/(4C_G)$ naturally follow, shown in the “Voltage” columns of Table II. 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.

4) *Junction Resistance*: To observe single-electron charging effects, electrons must be confined to the island. This requires junction resistances that are much higher than the quantum resistance, i.e., $R_S, R_D \gg h/e^2$, $h/e^2 = 25.8$ k Ω , where h is Planck’s constant. Therefore, SETs have high resistances and low driving currents. In this work, we pick two resistance settings: 100 K Ω for high-performance applications and 10 M Ω for battery-powered systems, shown in the “Resist.” columns of Table II.

5) *Reliability Implications*: Researchers have pointed out the dangers posed by thermal noise as charging (state change) energy approaches thermal energy. We explicitly consider the effects of temperature on steady-state current during circuit analysis and its effects are reflected in our design decisions. We implicitly consider, and guard against, the effects of temperature-dependent shot noise by requiring charging energy to be a large multiple of the thermal energy. Designs with charging energies of both 10 and 40 times the thermal energy are evaluated in this paper ($10k_B T$ or $40k_B T$). Researchers have reported device operation at each level but the $40k_B T$ requirement is more reliable. At charging energies over $10k_B T$, the model we use is accurate to within 4% of the time-dependent master equation [23, 38].

Random background charge effects [34, 35] 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. Several recent devices have shown improved immunity to this noise, as described in Section II-B. Currently, the distribution of random background offset charges can only be determined after fabrication [3]. Susceptible SETs may suffer transient errors infrequently, e.g., only once per day. 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.

B. IceFlex Design

In this section, we present the architecture and circuit design of IceFlex. The microarchitecture of IceFlex is shown in

Figure 3. IceFlex is a cell-based design. 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 arithmetic unit that allows efficient implementations of non-linearly separable arithmetic operations; (3) a SET-based reconfiguration memory array that caches multiple configuration contexts to support efficient run-time reconfiguration; (4) a multi-gate SET-based input switch fabric; and (5) SET registers. In addition, IceFlex includes SET threshold logic-based majority voting logic units, allowing a flexible solution to run-time reliability problems. In IceFlex, a multi-level on-chip interconnect fabric forms inter-SELB connections. Local connections rely on a custom-designed, SET-driven, variable-length, constant-latency interconnect. Using a constant-latency interconnect structure reduces power consumption and simplifies physical-level design automation, e.g., placement and routing. SETs have limited driving strength. Therefore, IceFlex uses hybrid SET/CMOS circuits to drive global interconnects.

We now explain each IceFlex component and discuss both circuit and architecture design tradeoffs.

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, which will be described in the next section.

The proposed multi-gate SET multiplexer tree differs from existing CMOS-based 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 required driving strength. SETs may have multiple gate terminals. As described in Equation 5, the gate charging effect is a function of $\sum C_{G_i} V_{GS_i}$. Therefore, multiple control signals, e.g., the select signals for a multiplexer, can be supplied to a single SET, enabling a more compact circuit structure with better performance and power efficiency.

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, in which each of the q paths consists of a single multi-gate SET controlled by $\lceil \log_2 q \rceil$ select signals. Using this design, the logic depth of a n -to-1 multiplexer tree reduces to $\lceil \log_q m \rceil$ instead of $\lceil \log_2 m \rceil$. Figure 4 also shows a design case for $q = 4$. The output SET buffer is used to break long resistive path and improve the driving strength.

As described in Section II, thermal energy has significant impact on electron tunneling and the ratio of on to off currents, i.e., the ratio of the off to on resistance. This ratio decreases as the ratio of Coulomb charging energy (e^2/C) to thermal energy ($k_B T$) decreases. On the other hand, as the number of gate control signals per SET (hence the number of off paths connected in parallel) increases, the impact of the off paths on the circuit output increases. Consider, for the sake of example, the dual-gate 4-to-1 multiplexer design shown in Figure 4. The four logic inputs are 0001 and both select signals are logic one, i.e., $V_a = V_b = V$. Assume $C_a = C_b = C$. As shown in

TABLE II
DESIGN SPACE CHARACTERIZATION

Temp. (K)	$C_{\Sigma} = e^2/10k_B T$								$C_{\Sigma} = e^2/40k_B T$							
	Low power				High performance				Low power				High performance			
	Capacitance (aF)		Voltage (mV)	Resist. (M Ω)	Capacitance (aF)		Voltage (mV)	Resist. (k Ω)	Capacitance (aF)		Voltage (mV)	Resist. (M Ω)	Capacitance (aF)		Voltage (mV)	Resist. (k Ω)
	C_G	C_S C_D	V_{dd}, V_{in} $e/4C_G$	R_S R_D	C_G	C_S C_D	V_{dd}, V_{in} $e/4C_G$	R_S R_D	C_G	C_S C_D	V_{dd}, V_{in} $e/4C_G$	R_S R_D	C_G	C_S 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	0.70	0.23	57.46	10	0.70	0.23	57.46	100
77	1.45	0.48	27.65	10	0.37	0.12	107.70	100	0.36	0.12	110.60	10	0.36	0.12	110.60	100
103	1.08	0.36	36.99	10	0.37	0.12	107.70	100	0.27	0.09	147.95	10	0.27	0.09	147.95	100
120	0.93	0.31	43.09	10	0.37	0.12	107.70	100	0.23	0.08	172.37	10	0.23	0.08	172.37	100
200	0.56	0.19	71.82	10	0.37	0.12	107.70	100	0.14	0.05	287.28	10	0.14	0.05	287.28	100
250	0.45	0.15	89.77	10	0.37	0.12	107.70	100	0.11	0.04	359.10	10	0.11	0.04	359.10	100
300	0.37	0.12	107.70	10	0.37	0.12	107.70	100	0.09	0.03	430.91	10	0.09	0.03	430.91	100




Fig. 3. IceFlex microarchitecture.

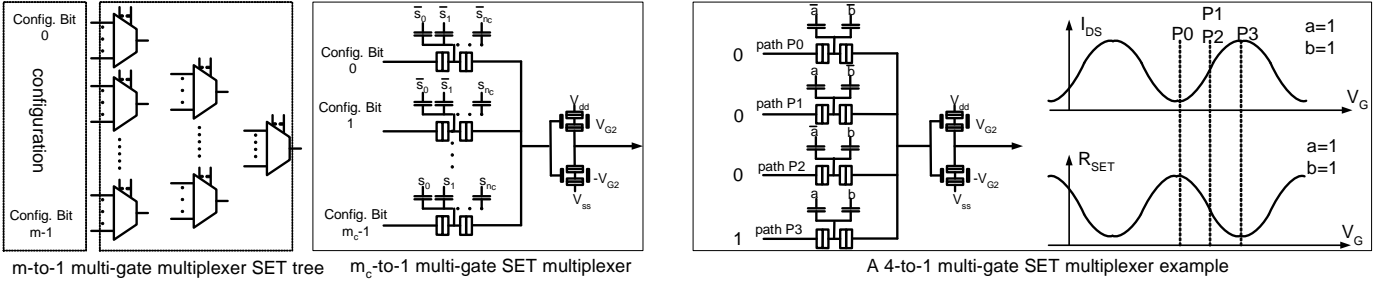


Fig. 4. Multi-gate SET multiplexer tree.

the I-V curve on the right side of Figure 4, for the SET on path P_3 , the overall gate charge equals $2CV$. Therefore, the SET becomes fully conductive. For paths P_1 and P_2 , the gate charges both equal $CV - CV = 0$, hence both switches are partially conductive. For path P_0 , even though the overall gate charge equals $-2CV$, at high temperature its resistance may still be within the same order of magnitude as that of path P_3 . Since the inputs of paths P_0 , P_1 and P_3 are all connected to logic zero (the worst-case scenario), these three parallel paths may reduce the output voltage, producing incorrect results.

In the high-performance setting, the same capacitance settings are used across the whole temperature range. Therefore, the ratio of Coulomb charging energy to thermal energy increases as the temperature decreases. Therefore, lower temperatures permit fewer multiplexer levels in the multiplexer tree, with more inputs to each individual multiplexer.

Detailed circuit analysis shows that, using the high-performance setting and $e^2/C_{\Sigma} \geq 10k_B T$, the dual-gate design may be used at temperatures up to 200 K. At 250 K and 300 K, only the single-gate design is feasible. For the low-power setting, capacitance scaling maintains the same

$e^2/C_{\Sigma}k_B T$ ratio. Therefore, the same design should be used for the whole temperature range. In addition, since both the low-power setting and the high-performance setting at room temperature use the same $e^2/C_{\Sigma}k_B T$ ratio, only the single-gate design is feasible for low-power, room-temperature operation. For the $e^2/C_{\Sigma} \geq 40k_B T$ configurations of IceFlex, the dual-gate design may be used at all temperatures due to the increased charging energy.

2) *SET Configuration Memory*: In IceFlex, run-time reconfiguration is enabled by SET configuration memory, which consists of SET configuration cache and current configuration memory. In each SELB, the configuration cache stores multiple configurations. During run-time reconfiguration, one set of configuration bits stored in the configuration cache are placed into the current configuration memory to program SELB logic and interconnect. If k copies of configuration sets are stored in the configuration cache, then the circuit can be reconfigured k times during run-time execution without the need to access off-chip memory.

The left portion of Figure 5 shows the circuit structure of the configuration memory in IceFlex. The SET configuration

cache is the main on-chip configuration memory. Each storage cell consists of a dual-island SET [3]. A dual-island SET contains two capacitively-coupled SETs: a primary SET and a secondary SET. 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 shifts the phase of the Coulomb oscillations of the primary gate, i.e., its conductivity condition shifts as a function of gate control voltage, V_{GS} . Therefore, under a certain V_{GS} , the primary SET is either conductive or open due to different island charges, representing either a logic one or logic zero.

In the configuration cache, selecting a configuration forms a short-circuit path between the pull-up resistor and SETs with a stored zero within the selected configuration set. The power consumption will be high if the configuration cache constantly controls the logic and interconnect. To minimize power consumption, separate on-chip memories are used to store the currently-used configuration.

We designed two types of SET-based on-chip storage to hold the current configuration. The first design is a dual-island based SET buffer. As shown in the *SET SRAM* portion of Figure 5, this buffer uses two opposite 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 cache updates the island charge of its secondary SET, hence the conductivity of its primary SET, thereby controlling the buffer output. The second design is a SET SRAM design, which is similar to CMOS SRAM.

3) *Efficient SET Implementations of Non-Unate Functions and Implications for 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 highly-useful 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 six transistors [39]. Moreover, it relies on strong input signals because it is not capable of signal restoration. A restoring version would require at least 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. In this design, each SET has two gates, each of which is connected to one of the exclusive-or inputs. The circuit structure for a SET-based n -input parity gate is shown to the right of Figure 5. This design is capable of signal restoration. Thanks to the periodic SET I-V curve, it is possible to directly determine whether the number of high inputs is odd or even. 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). Note that even or odd parity functions with additional inputs may

be implemented using only two SETs. The number of inputs is bounded primarily by geometrical constraints on fabrication of additional gates.

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 approximately 75% more energy-efficient and 25% faster than a design based on a conventional CMOS-style exclusive-or sum implementation, when both are implemented using SETs. This design style is impossible for threshold logic, BJTs, and CMOS technologies. Note that carry-out logic is equivalent to 2-out-of-3 majority vote logic.

4) *Reconfigurable Interconnect Network*: IceFlex consists of a variety of reconfigurable interconnect resources, including SET local interconnects, hybrid SET/CMOS global interconnects, and SET switch fabric.

Interconnect consumes a substantial proportion of total power consumption in IceFlex: its power efficiency is important. For SET-based interconnect, the static power consumption dominates due to the impact of thermal energy on device conductance, especially at high temperatures. In addition, static power consumption increases with wireload because maintaining unchanged communication latency with higher wireload requires lower junction resistance. In contrast, the dynamic power consumption of SETs is low due to the low SET gate capacitance and low voltage swing. For hybrid SET/CMOS-based interconnect, SETs are only used to drive CMOS buffers, which in turn drive wires. In this case, SETs with low driving strength, hence high junction resistance, are allowed. 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. Circuit analysis shows that, given the same performance constraint, SET-based design is more energy-efficient for local interconnect and the hybrid SET/CMOS design is more energy-efficient for global interconnect.

In IceFlex, local interconnects driven directly by SET buffers support communication between nearby SELBs. Three types of local interconnects are supported: single length, double length, and hex length. The proposed SET local interconnect design guarantees a constant latency across different routing lengths. Consider, for the sake of example, a local communication architecture in which the maximum interconnect delay is constrained and the longest interconnect is appropriately buffered to meet this constraint. In this case, it would be possible to similarly drive shorter interconnects, thereby decreasing their delays, relative to that of the longest interconnect. It would also be possible to reduce the driving strength on shorter interconnects to reduce power consumption and produce a local interconnect architecture in which all interconnects have uniform delay. We propose the second design because it improves interconnect power efficiency and also simplifies placement and routing during physical design.

The proposed SET local interconnect is designed as follows. A SET buffer with minimal driving strength (hence high junction resistance) is first determined. Next, for local interconnects with different routing lengths, minimal driving

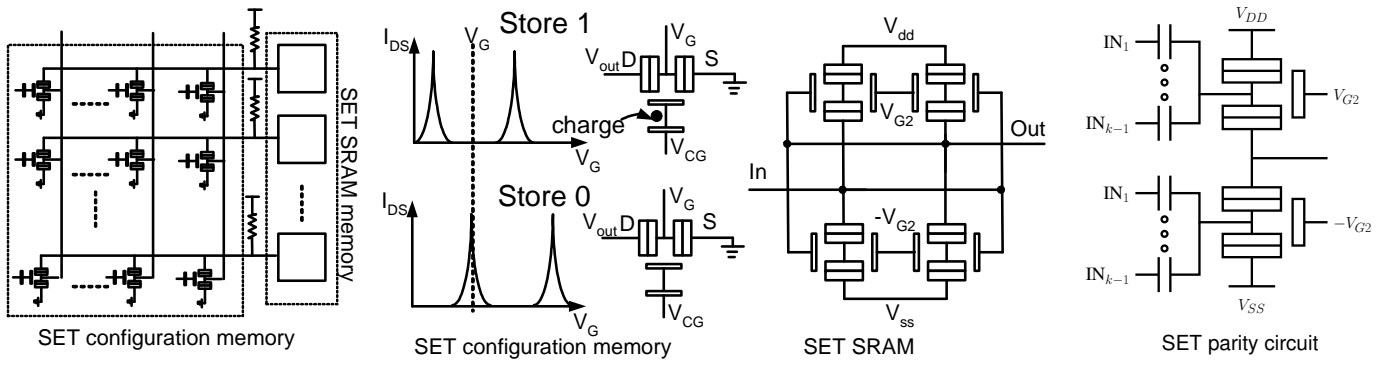


Fig. 5. SET configuration memory and parity circuit.




Fig. 6. Hybrid SET/CMOS interface circuitry.

strength SET buffers are connected in parallel to meet driving strength requirements imposed by performance constraints. The main motivation for using parallel SET buffers is that SET junction resistance cannot be reduced arbitrarily ($R_D, R_S \gg h/e^2$). Using homogeneous SET buffers in parallel instead of heterogeneous SET buffers may also simplify fabrication.

Remote connections introduce the high capacitive loads of long metal wires. To address the driving strength problem of SET-only circuits, we have designed hybrid SET/CMOS interface circuitry to drive global interconnect. Figure 6 shows the circuit structure, which contains two complementary SET inverters and two CMOS inverters. A SELB output is first fed to the input of SET inverter SINV1. SINV1 drives the CMOS inverter, CINV1. Unlike the SET logic used inside SELBs, SINV1 uses a low-resistance design to improve driving strength. Fortunately, it is possible to achieve sufficient driving strength with a single SET. Since the voltage range of SET logic is much smaller than that of CMOS logic, the output signal of SINV1 is within the switching range of the CMOS inverter. Since both MOS transistors are conductive within the switching region, short-circuit power is high. To solve the short-circuit power consumption problem, CINV1 is designed to satisfy the following two constraints. First, $V_{tn} + |V_{tp}| > V_{dd} - V_{ss}$ ensures that at least one MOS transistor is off at all times, reducing static power consumption. Second, the output signal range of SINV1 must be greater than $V_{tn} + |V_{tp}| - (V_{dd} - V_{ss})$. Therefore, the NMOS (PMOS) transistor of CINV1 is conductive when SINV1 has a high (low) output signal. Therefore, CINV1 serves as a signal converter, and CINV2 provides driving strength.

CINV2 cannot be used to drive the input SET logic of a SELB directly. SET current is a periodic function of the gate control voltage and has a period of e/C_G , which is much smaller than the output voltage range of CINV2. Therefore, this output voltage range cannot be used directly. To solve this problem, we design a special SET inverter, SINV2, that

is used for SELB inputs. SINV2 is fabricated with a large distance between gate and island in order to reduce the gate capacitance, C_G . Thus, e/C_G can match the output signal range of CMOS inverter CINT2. Although source–island and drain–island junctions must be short to permit tunneling, there is no such bound on gate–island separation.

In IceFlex, each SELB is equipped with a reconfigurable input switch fabric that selects the connections among local and global interconnects. The input switch fabric is implemented using multi-gate SET multiplexor tree, similar to that in the reconfigurable look-up table described in Section III-B1.

5) *Design and Modeling of IceFlex Majority Voting Logic:* Although researchers are making progress on reducing the severity of noise resulting from random background offset charge effects, it may continue to pose run-time noise problems in the future. Even if this problem can be entirely solved, resistance to run-time faults may be useful in SETs, e.g., to allow resistance to Alpha particle induced faults or other single event upsets. IceFlex incorporates support for hierarchical spatial redundancy to improve fault tolerance. Although much of the literature predicts the need for fault-tolerant architectures in nanoelectronics, the level of fault tolerance is currently unknown. Therefore, we consider the results for a number of possible SET failure rates and in the presence of three fault-tolerance configurations.

Other researchers have proposed a number of architectural techniques to support reliable computation using nanoscale electronics that are susceptible to fabrication-time and run-time faults. Dehon described the use of structural redundancy and programming-time defect-aware configuration in a carbon nanotube and silicon nanowire based programmable logic array architecture [40]. Goldstein et al. describe the use of a defect map that is generated during post-fabrication testing to avoid the use of faulty devices [41]. Bahar et al. present a method of expressing logic circuits using Markov Random Fields, permitting Boolean functions to be computed using devices susceptible to potentially-frequent transient faults [42]. We think it likely that the random background charge problem will ultimately be dealt with by a combination of improved fabrication technology, post-fabrication testing to identify and avoid a subset of the affected SETs, and run-time fault-tolerance via conventional structural redundancy or recent advances in probabilistic computation. IceFlex provides for regular structural redundancy and run-time error correction.

TABLE III
IMPACT OF MAJORITY VOTE LOGIC ON SELB FAULT PROBABILITY

SET fault probability Majority vote inputs	1/1,000			1/10,000			1/100,000		
	3	5	7	3	5	7	3	5	7
Raw fail prob.	6.20E-2	6.20E-2	6.20E-2	6.38E-3	6.38E-3	6.38E-3	6.40E-4	6.40E-4	6.40E-4
Best prob.	1.11E-2	2.17E-3	4.45E-4	1.22E-4	2.57E-6	5.71E-8	1.23E-6	2.62E-9	5.86E-12
SET MVL prob.	1.11E-2	2.18E-3	4.57E-4	1.22E-4	2.69E-6	1.77E-7	1.23E-6	3.82E-9	1.21E-9

We now consider the fault model for IceFlex SELBs. Every path from SELB input to output contains 64 SETs. In the third row of Table III, we show the SELB raw failure probabilities, i.e., the probability of a SELB producing an incorrect output. SELB failure probability is a function of the SET fault probability, for which Table III shows three values. 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. The resulting $1/f$ noise produces long-duration failure periods. Therefore, in this analysis, we (conservatively) assume that susceptible devices consistently fail. In reality, errors may not be consistent. We also consider the higher SET fault probability of 1/1,000 and the lower fault probability of 1/100,000. Advances in fabrication and detection of most SETs susceptible to random background offset charge effects by post-fabrication testing may permit reduction in run-time SET fault probability.

We have considered the effect of using no MVL (*Raw fail prob.*), fault-free MVL (*Best prob.*), and SET MVL. Using a given reliability configuration, it is not possible for MVL-based designs to produce lower SELB fault probabilities than those shown in the *Best prob.* row. SET MVLs are constructed from multi-gate SETs. We focus on the three-input SET MVL design to simplify depiction; the five-input, and seven-input SET MVL follows an analogous design style. This circuit has identical structure to the parity gate shown in the *SET parity circuit* portion of Figure 5. However, the separation of gates and island are adjusted such that the circuit traverses only 1/2 Coulomb oscillation period during use. The SET pull-up gates are separated sufficiently to require the majority of the gates to be high. The converse is true of the pull-down gates. For each SET depicted in the figure, four SETs are used in parallel in order to permit the failure of one SET while still producing correct results. We have computed the delay of the SET MVL by considering the worst-case scenario, in which a path that is 3/5 or 4/7 closed has a faulty driver SET and a path that is 2/4 or 3/7 closed has no faulty SETs.

As shown in Table III it is possible for a seven-input SET-only MVL with redundant SELBs to reduce the failure rate to 1/8,500,000, given a SET fault probability of 1/10,000, or 1/830,000,000, given a SET fault probability of 1/100,000. 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 (see Section II-B).

If a method of rapidly determining which SETs are susceptible to random background charge effects is ever developed, these effects can be avoided in the same way that fabrication defects are avoided: via the use of a regular computation structure in which operations are mapped only to fault-free devices. There has been some promising work on this topic,

in which illumination is used to produce ions, accelerating the onset of random background charge effects [43].

IV. EXPERIMENTAL RESULTS

In this section, we evaluate the suitability of using SETs in low-power embedded system design. 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 application.

A. Characterization of the IceFlex Architecture

Following the design parameters shown in Table II, we evaluate the performance and power consumption of IceFlex using HSPICE. For SET circuitry, the SPICE model and device parameters are described in Section II-C. For CMOS logic and metal wire, we use the 22nm Berkeley BSIM4 predictive technology model, which models the impact of temperature on MOS devices. We analyzed designs adhering to the $C_{\Sigma} = e^2/(40k_B T)$ constraint. We also analyzed designs with the less conservative $C_{\Sigma} = e^2/(10k_B T)$ constraint, although space constraints force us to omit detailed results tables for this setting. A low-power setting (targeting megahertz-range frequencies) and a high-performance setting (targeting gigahertz-range frequencies), are considered.

Tables IV and V summarize the performance and power characterization of the logic components and interconnect fabric of the $C_{\Sigma} = e^2/(40k_B T)$ version of IceFlex, including multi-gate SET reconfigurable lookup table (LUT)¹, SET register (Register), SET and CMOS four-out-of-seven majority voting logic (MVL), multi-gate (MG) and CMOS-style (CS) exclusive-or, (CO) carry-out logic, and SET local interconnect (Single, Double, and Hex), hybrid SET/CMOS global interconnect (Global) and SET input switch fabric (ISF). From these results, we make the following observations.

First, IceFlex has high energy efficiency, good performance, and high flexibility in terms of performance and energy efficiency tradeoff. At the low-power setting, the power consumptions of SET-based logic components and local interconnect fabric are nano-Watts. The hybrid SET/CMOS global interconnect has the highest power consumption. This is a result of the high capacitance of global wires and high power consumption of the CMOS buffers. All components in the low-power version of IceFlex still have latencies in the range of nanoseconds. SETs have high junction resistance and low driving strength. Using the high-performance setting, by scaling the SET junction resistance down to 100k Ω , the latencies of the SET-based logic and local interconnect fabric are consistently lower than 100ps. Even though reducing

¹To allow comparison with Xilinx FPGAs, a 16-to-1 setting is used.

TABLE IV
CHARACTERIZATION OF ICEFLEX MICROARCHITECTURE FOR $C_{\Sigma} = e^2/(40k_B T)$

		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	10.04	7.86	7.09	6.80	5.57	5.03	4.75	0.08	0.06	0.05	0.05	0.05	0.04	0.04	
	Register	1.42	1.09	1.02	1.00	0.90	0.88	0.86	0.01	0.01	0.01	0.01	0.01	0.01	0.01	
	7-INPUT MVL	0.58	0.57	0.58	0.58	0.59	0.56	0.58	3.28E-03	3.18E-03	3.16E-03	3.20E-03	3.24E-03	2.99E-03	3.14E-03	
	SET-MVL	1.15	1.13	1.13	1.00	1.08	1.04	1.06	0.01	0.01	0.01	0.01	0.01	0.01	0.01	
	Arithmetic Logic	SUM	2.32	2.31	2.31	2.31	2.31	2.28	2.29	0.01	0.01	0.01	0.01	0.01	0.01	0.01
		MG	3.02	2.97	2.95	2.96	2.95	2.89	2.93	0.01	0.01	0.01	0.01	0.01	0.01	
		CS	1.15	1.13	1.13	1.00	1.08	1.04	1.06	0.01	0.01	0.01	0.01	0.01	0.01	
		CO	0.07	0.26	0.44	0.58	1.60	2.64	3.70	6.67	25.76	44.53	58.19	162.20	266.69	373.81
Power (nW)	LUT	0.08	0.30	0.53	0.72	1.99	3.14	4.48	8.02	29.88	53.16	72.12	199.64	315.21	450.34	
	Register	0.05	0.20	0.36	0.48	1.32	2.17	3.02	5.37	20.05	35.87	48.15	132.24	217.31	302.60	
	7 INPUT-MVL	0.01	0.03	0.06	0.08	0.21	0.34	0.48	0.94	3.51	6.26	8.44	23.24	37.58	52.90	
	SET-MVL	1.61E-03	0.01	0.01	0.01	0.04	0.07	0.09	0.22	0.80	1.44	1.91	5.19	8.88	12.04	
	Arithmetic Logic	SUM	0.01	0.04	0.07	0.09	0.25	0.40	0.57	1.04	3.87	6.90	9.30	25.60	41.51	58.35
		MG	0.01	0.03	0.06	0.08	0.21	0.34	0.48	0.94	3.51	6.26	8.44	23.24	37.58	
		CS	0.01	0.03	0.06	0.08	0.21	0.34	0.48	0.94	3.51	6.26	8.44	23.24	37.58	
		CO	0.01	0.03	0.06	0.08	0.21	0.34	0.48	0.94	3.51	6.26	8.44	23.24	37.58	

TABLE V
CHARACTERIZATION OF ICEFLEX INTERCONNECT FABRIC FOR $C_{\Sigma} = e^2/(40k_B T)$

		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)	ISF	6.696	5.238	4.727	4.537	3.712	3.351	3.169	0.050	0.039	0.037	0.036	0.030	0.028	0.027
	Single	0.728	0.699	0.694	0.697	0.799	0.770	0.784	0.006	0.006	0.006	0.007	0.005	0.005	0.005
	Double	0.704	0.687	0.685	0.689	0.794	0.766	0.781	0.006	0.006	0.006	0.007	0.005	0.005	0.005
	Hex	0.692	0.680	0.680	0.684	0.791	0.763	0.779	0.006	0.006	0.006	0.007	0.005	0.005	0.005
	Global	2.996	4.523	4.657	4.237	4.572	4.520	6.785	0.163	0.110	0.092	0.086	0.074	0.073	0.099
Power (nW)	ISF	0.219	0.844	1.457	1.903	5.302	8.727	12.226	22.022	85.034	146.920	191.957	535.072	879.837	1233.147
	Single	0.008	0.032	0.057	0.076	0.210	0.342	0.479	0.959	3.387	6.193	7.977	24.992	34.101	53.581
	Double	0.017	0.063	0.113	0.152	0.420	0.684	0.958	1.917	6.775	12.386	15.955	49.984	68.202	107.160
	Hex	0.034	0.127	0.226	0.305	0.840	1.368	1.917	3.835	13.549	24.771	31.909	99.967	136.400	214.320
	Global	271.780	23.912	6.668	4.460	3.555	4.513	5.857	6674.800	5146.700	5560.900	5824.100	5318.200	4856.100	4745.700

resistance 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.

Second, these results demonstrate the impact of temperature on SET performance and power consumption – as the temperature increases, performance increases and the power efficiency decreases. This is a result of the impact of thermal energy on tunneling events and therefore circuit behavior, which is described in Section II. The number of electrons with sufficient energy to overcome the Coulomb blockade effect increases with temperature, thereby increasing tunneling rate, performance, and power consumption.

The $C_{\Sigma} = e^2/(40k_B T)$ setting enables greater resistance to shot noise than the $C_{\Sigma} = e^2/(10k_B T)$ setting. However, it also imposes performance and power consumption penalties. For SET circuitry, the required supply voltage is inversely proportional to gate capacitance. Compared to the $C_{\Sigma} = e^2/(10k_B T)$ setting, $C_{\Sigma} = e^2/(40k_B T)$ requires a further reduction of SET gate capacitance and an increase in supply voltage. Note that the driven capacitance of a SET circuit is dominated by the metal wires. Therefore, decreased gate capacitance has negligible impact on power consumption. The increased supply voltage, on the other hand, increases circuit dynamic power consumption. Moreover, the increased voltage range increases the duration of signal swing, thereby increasing latency. If the less conservative $C_{\Sigma} = e^2/(10k_B T)$ design versions were able to provide adequate noise immunity, the latencies reported in Table IV would be halved at all temperatures, as would power consumptions at higher temperatures.

1) *SET Multi-Gate Multiplexer Tree*: As described in Section III-B1, multi-gate SETs improve the performance, power

consumption, and area efficiency of the multiplexer tree design. This section characterizes the impact of thermal energy on the proposed multi-gate design.

As described in Section III-B1, at the high-performance $C_{\Sigma} = e^2/(10k_B T)$ setting, the dual-gate design is used for temperatures at or below 200 K. For these settings only single-gate design is feasible at temperatures greater than 250 K due to high static current at these temperatures. As a result, circuit power consumption is increased at high temperatures. From 200 K to 250 K, both latency and power consumption increase. In addition, when using the same design, we observe that both the circuit performance and power consumption increase with temperature. The same trend was described in Section IV-A. Using the low-power design of IceFlex, only the single-gate design is feasible (see Section III-B1). The results are summarized in Table IV. Using $e^2/C_{\Sigma} \geq 40k_B T$, SET circuitry is less susceptible to thermal energy thanks to the increased charging energy. Therefore, both low-power and high-performance dual-gate multiplexer tree designs become feasible across the entire temperature range. As shown in Figure 7, using the high-performance $C_{\Sigma} = e^2/(40k_B T)$ setting, the performance and power consumption of the multi-gate multiplexer tree design increase consistently with temperature. A similar trend can be shown for the corresponding low-power design case.

2) *Power and Performance of Interconnect Design*: Power consumption, performance, and the tradeoff between them are of central importance in interconnect design. We considered both SET-only and SET/CMOS hybrid interconnect driver designs. The relative static power benefit of the SET/CMOS hybrid design over the SET-only design increases as the wireload increases. This is mainly due to an increase in the static power consumption of the SET-only design as more SET

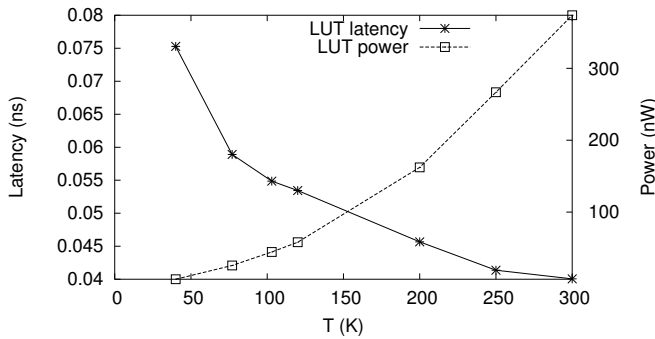


Fig. 7. Power and performance of the multi-gate SET multiplexer tree for high performance, $C_\Sigma = e^2/(40k_B T)$.

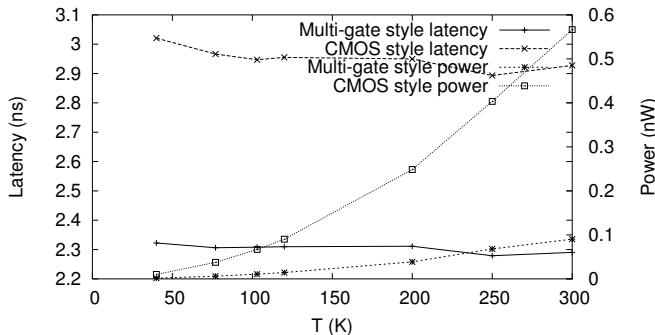


Fig. 8. Performance and power characterization of exclusive-or logic for low power for $C_\Sigma = e^2/(40k_B T)$.

buffers are used to meet the driving strength requirements. The SET-only design has superior power efficiency. As the wire length increases, the proportion of capacitance contributed by CMOS buffer gates becomes less significant relative to wire capacitance. Therefore, compared to the SET-only design, the dynamic power consumption of the SET/CMOS hybrid design also improves, but is still inferior to that of the SET-only design. At 300 K, for both the $C_\Sigma = e^2/(40k_B T)$ and $C_\Sigma = e^2/(10k_B T)$ settings, we found that SET-only designs had better energy efficiencies for wires shorter than approximately 1 mm, and SET/CMOS hybrid designs were better for longer wires. As temperature increases, the thermal energy impact increases. As a result, the static power consumption of SETs increases. Therefore, the wire length at which the SET/CMOS design begins to outperform the SET-only design decreases as temperature increases.

Table V illustrate two interesting trends for global interconnect. The power consumption of both the low-power and the high-performance $C_\Sigma \leq e^2/(40k_B T)$ hybrid SET/CMOS designs decrease with increasing temperature. At low temperatures, the output voltage ranges and driving currents for the SETs are small, increasing CMOS buffer static power consumption.

3) *Performance and Power Characterization of SET Non-Unate Logic*: SETs support the efficient implementation of some non-unate arithmetic functions. We evaluate the power consumption and performance of an exclusive-or gate, a non-unate Boolean function widely used in arithmetic logic, e.g., in addition and multiplication. We compared the two different implementations described in Section III-B3, the proposed SET-based design and the CMOS-style SET implementation.

TABLE VI
LATENCY AND ENERGY IMPROVEMENT FOR EXCLUSIVE-OR DESIGN

Performance setting	C_Σ constraint (F)	Performance improvement (%)	Energy improvement (%)
Battery	$e^2/(10k_B T)$	40.8	64.1
Battery	$e^2/(40k_B T)$	22.0	87.1
High	$e^2/(10k_B T)$	32.1	84.6
High	$e^2/(40k_B T)$	25.2	84.4

Figure 8 shows the power and performance characterization of these two designs for the low-power and $C_\Sigma = e^2/(40k_B T)$ settings. These results demonstrate the superior power consumption and performance of this design style, which is not possible using BJTs, CMOS, or threshold logic. Compared to the CMOS-style SET implementation, the design that exploits the periodic I-V curve of SETs achieves the latency and power consumption reductions indicated in Table VI, i.e., approximately a 25% reduction in latency and 75% reduction in energy consumption.

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

This section characterizes the performance and power consumption of IceFlex when used to implement numerous general-purpose and application-specific processor cores. We evaluate the suitability of IceFlex for use in both portable battery-powered and high-performance embedded systems by determining its performance and energy efficiency when used to implement the processor cores described below. We have divided the cores into *battery-powered* and *high-performance* categories for the convenience of the reader.

Battery-powered: AES (Rijndael) IP core (AES), AT-Mega103 microcontroller (AVR), coordinate rotation computer (CORDIC), ECC core (ECC), 32-bit IEEE 754 floating-point unit (FPU), Reed-Solomon encoder (RS), USB 2.0 function (USB), and video compression systems (VC).

High-performance: Power-efficient RISC CPU (ARM7), synchronous / DLX core (ASPIDA DLX), five-stage pipeline RISC CPU (Jam RISC), entire SPARC V8 processor (LEON2 SPARC), RISC CPU (Microblaze), MIPS I clone (miniMIPS), MIPS processor (MIPS) supporting most MIP I opcodes (Plasma), MIPS I integer only clone (UCore), and MIPS I clone (YACC).

The Xilinx Virtex-II XC2V2000 FPGA is used as a base case 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% [44] and a 65 nm feature size. Then, we scale the FPGA synthesis results into a 22 nm process based on HSPICE predictive technology model simulation results for the two technologies [45]. 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 SELBs along the longest combinational path by the delay of an IceFlex SELB plus the delay of a local interconnect. IceFlex power consumption was computed by taking the sum of the power consumptions of all components

TABLE VII
ICEFLEX PERFORMANCE AND POWER CONSUMPTION AT ROOM
TEMPERATURE FOR $C_{\Sigma} = e^2/(40k_B T)$

Benchmarks	FPGA		IceFlex			
	22 nm CMOS technology*		Battery-powered		High-performance	
	Freq (MHz)	Energy (J/cycle)	Freq (MHz)	Energy (J/cycle)	Freq (MHz)	Energy (J/cycle)
ARM7	26.3	2.96e-09	2.0	5.47e-11	224.0	4.79e-11
ASPIDA DLX	125.7	8.86e-10	11.5	6.37e-12	1333.3	5.58e-12
Jam RISC	95.9	8.92e-10	12.8	3.65e-12	1481.5	3.19e-12
LEON2 SPARC	85.9	1.88e-09	8.8	2.39e-11	1025.6	2.09e-11
Microblaze RISC	115.1	7.28e-10	16.4	2.01e-12	1904.8	1.76e-12
miniMIPS	88.0	4.87e-10	9.6	9.78e-12	1111.1	8.56e-12
MIPS	80.4	1.02e-09	10.5	4.34e-12	1212.1	3.80e-12
Plasma	75.4	1.13e-09	8.8	6.91e-12	1025.6	6.05e-12
UCore	136.4	8.19e-10	12.8	5.45e-12	1481.5	4.78e-12
YACC	72.1	1.18e-09	19.2	3.08e-12	2222.2	2.69e-12
AES	205.3	3.43e-10	28.7	2.34e-12	3333.3	2.05e-12
AVR	71.9	2.67e-10	9.6	5.34e-12	1111.1	4.67e-12
CORDIC	271.8	1.37e-10	114.9	2.05e-13	13333.3	1.79e-13
ECC	39.1	4.91e-10	11.5	6.92e-12	1333.3	6.05e-12
FPU	28.4	1.00e-09	2.6	8.02e-11	296.3	7.02e-11
RS	496.7	1.28e-11	57.5	4.61e-14	6666.7	4.05e-14
USB	171.6	3.24e-10	38.3	1.53e-12	4444.4	1.34e-12
VC	114.16	1.24e-09	23.0	1.04e-11	2666.8	9.10e-12
Avg. energy Improvement			68.58×		78.46×	

at the maximum operating frequency. Note that, since Xilinx ISE does not report use of global interconnect for any of the processors we synthesized, we exclude the hybrid global interconnect from IceFlex power analysis. In designs that use primarily local interconnect (i.e., single, double, and hex interconnect), the reported power consumption results will be accurate. However, for designs in which global hybrid SET-CMOS interconnect dominates, the power consumption may approach that of global interconnect in a corresponding 22 nm CMOS design.

Table VII shows the operating frequencies and energy efficiency in Joules per clock cycle of the CMOS FPGA and IceFlex variants for each benchmark application. As described in Section III-A5, recent progress in fabrication is reducing the severity of the random background charge problem. Due to space constraints, we show the characteristics of non-spatially-redundant versions of IceFlex.

We have also analyzed the performance and power consumption of a configuration with seven-way SELB spatial redundancy. The MVL uses internal fine-grained four-way SET parallelism. The MVL fault-tolerance hardware delays are added to the delay of each SELB stage. The fault-tolerant versions of the processors have similar maximum frequencies to those listed in Table VII. However, they only improve on the energy efficiency of 22 nm CMOS FPGAs by averages of 8.64× (battery-powered) and 10.55× (high-performance).

1) *Ultra-Low-Power Applications:* The data in Table VII indicate that the non-redundant, room temperature, low-power version of IceFlex is suitable for use in applications such as sensor network nodes, if they can be fabricated with sufficiently small island capacitances. In the following analysis, we shall focus on the AVR core, which is representative of a commonly-used sensor network node processor. Alkaline AA batteries typically have 2,800 mAH of energy and nominal operating voltages of 1.5 V, i.e., they can deliver approximately 15,000 J. Using the conservative $C_{\Sigma} \leq e^2/(40k_B T)$

constraint, a low-power IceFlex AVR implementation running at 4 MHz consumes approximately 200 μ W, permitting it to run for 20 years on one AA battery, i.e., longer than the shelf life of most such batteries. When the less conservative $C_{\Sigma} \leq e^2/(10k_B T)$ constraint is used, the average energy consumption improvements increase to 95.60× (non-redundant battery powered), 115.65× (non-redundant high performance), 12.27× (redundant battery powered), and 15.27× (redundant high performance).

This power consumption is also low enough to permit an AVR processor to operate on energy scavenged from the environment. If we assume an energy scavenging volume of 5 cm³ and use Roundy's power densities of 4 μ W/cm³ for indoor solar energy, 200 μ W/cm³ for vibrations, 10 μ W/cm³ for daily temperature variation, and 0.003 μ W/cm³ for acoustic noise at 75 dB [46], we find that one sensor network node is capable of scavenging enough energy to power an IceFlex AVR processor running at the maximum clock frequency from vibrations or daily temperature variation, at 3.7 MHz from indoor solar energy, and at 2.8 kHz from 75 dB acoustic noise. However, SET circuits that operate at room temperature and adhere to the $C_{\Sigma} \leq e^2/(40k_B T)$ constraint will rely on features with sizes approaching (but not crossing) physical limits. Although the use of SETs in battery-powered applications has potential, it depends on the solution of formidable fabrication challenges or the development of compact, low-power cooling methods.

2) *Energy-Efficient High-Performance Applications:* We can draw the following general conclusions from Table VII. For a wide range of processor cores, the SET-based IceFlex architecture is capable of achieving energy efficiencies two orders of magnitude better than 22 nm CMOS-based FPGAs. Peak frequencies ranging from 200 MHz to 2 GHz are maintained for all processors.

One might expect the high-performance version of IceFlex to consistently achieve higher frequency but lower energy efficiency than the low-power version of IceFlex. However, its energy efficiency is typically better, as well. Operating at higher frequencies can permit reduced static energy consumption, and therefore better energy efficiency, especially at room temperature where static power consumption is high (see Figure 2). Therefore, for SET-based architectures that are operated at room temperature and have low performance requirements, it will generally be more energy efficient to operate the device at high frequency and periodically 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 125 mW and $C_{\Sigma} \leq e^2/(40k_B T)$, one could use one LEON2 SPARC implemented with an FPGA and running at 85 MHz or 5 LEON2 SPARCs implemented with the high-performance variant of IceFlex and operating at 1,025 MHz. This implies an overall performance 60× higher than that of the FPGA version. Taken to its logical extreme, assuming a power budget of 100 W and one instruction per cycle, one could execute 4.8 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 $268\times$ improvement in energy efficiency compared with an FPGA implementation.

V. CONCLUSIONS

In this article, we have analyzed the impact of using SETs in architecture and circuit design; proposed IceFlex, a fault-tolerant, reconfigurable, hybrid SET/CMOS 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. Our results indicate that using SETs for computation poses many design challenges, some of which can be solved with the proposed architecture and circuit design techniques. In addition, we find that SETs have unique properties that permit significant improvements in circuit efficiency when compared with BJT, CMOS, and threshold logic based design. In summary, we find that a hybrid SETs/CMOS architecture has the potential to improve energy efficiency in battery-powered high-performance applications by two orders of magnitude compared with 22 nm CMOS while permitting operating frequencies that are as high, or higher. Although they hold great promise, the practical use of SETs will require additional research into fault tolerance techniques, processing technologies, and novel circuit designs. In particular, the use of SET-based designs in portable applications will either require the fabrication of features with sizes approaching physical limits or the development of compact, energy-efficient technologies permitting operation below ambient temperature. We hope this article provides a starting point for additional research in the area and reveals the potential advantages and challenges of SET-based architectures.

REFERENCES

- [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] "International Technology Roadmap for Semiconductors," 2006, <http://public.itrs.net>.
- [5] D. V. Averin and K. K. Likharev, "Coulomb blockade of tunneling and coherent oscillations in small tunnel junctions," *J. Low Temperature Physics*, vol. 62, pp. 345–372, Feb. 1986.
- [6] T. A. Fulton and G. J. Dolan, "Observation of single-electron charging effects in small tunnel junctions," *Physics Review Letters*, vol. 59, pp. 109–112, July 1987.
- [7] M. H. Devoret and R. J. Schoelkopf, "Amplifying quantum signals with the single-electron transistor," *Nature*, vol. 406, pp. 1039–1046, Aug. 2000.
- [8] Y. Nakamura, C. D. Chen, and J. S. Tsai, "100-K operation of Al-based single-electron transistors," *Japan Journal Applied Physics*, vol. 35, pp. 1465–1467, Nov. 1996.
- [9] X. Tang, et al., "An SOI single-electron transistor," in *Proc. Silicon-on-Insulator Conf.*, Oct. 1999, pp. 46–47.
- [10] M. Ahlskog, et al., "Single-electron transistor made of two crossing multiwalled carbon nanotubes and its noise properties," *Applied Physics Letters*, vol. 77, pp. 4037–4039, Dec. 2000.
- [11] K. Matsumoto, et al., "Room temperature operation of a single electron transistor made by the scanning tunneling microscope nanooxidation process for the TiO_x/Ti system," *Applied Physics Letters*, vol. 68, no. 1, pp. 34–36, Jan. 1996.
- [12] J.-I. Shirakashi, et al., "Single-electron charging effects in Nb/Nb oxide-based single-electron transistors at room temperature," *Applied Physics Letters*, vol. 72, no. 15, pp. 1893–1895, Apr. 1998.
- [13] Y. A. Pashkin, Y. Nakamura, and J. S. Tsai, "Room-temperature Al single-electron transistor made by electron-beam lithography," *Applied Physics Letters*, vol. 76, no. 16, pp. 2256–2258, Apr. 2000.
- [14] J. R. Tucker, "Complementary digital logic based on the Coulomb blockade," *J. Applied Physics*, vol. 72, no. 99, pp. 4399–4413, 1992.
- [15] 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.
- [16] F. Nakajima, et al., "Single-electron AND/NAND logic circuits based on a self-organized dot network," *Applied Physics Letters*, vol. 83, no. 13, pp. 2680–2682, Sept. 2003.
- [17] 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.
- [18] K. Yano, et al., "Room-temperature single-electron memory," *IEEE Trans. Electron Devices*, vol. 41, pp. 1628–1638, Sept. 1994.
- [19] C. Wasshuber, H. Kosina, and S. Selberherr, "A comparative study of single electron memories," *IEEE Trans. Electron Devices*, vol. 45, pp. 2365–2371, Nov. 1998.
- [20] 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.
- [21] 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.
- [22] 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.
- [23] 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.
- [24] 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.
- [25] S. Mahapatra, et al., "Analytical modelling of single electron transistor (SET) for hybrid CMOS-SET analog IC design," *IEEE Trans. Electron Devices*, vol. 51, no. 11, pp. 1772–1782, June 2004.
- [26] J. R. Heath and M. A. Ratner, "Molecular electronics," *Physics Today*, vol. 56, pp. 43–49, May 2003.
- [27] A. K. Geim and K. S. Novoselov, "The rise of graphene," *Nature Materials*, vol. 6, pp. 183–191, Mar. 2007.
- [28] S. Vanapalli, et al., "120 Hz pulse tube cryocooler for fast cooldown to 50 K," *Applied Physics Letters*, vol. 90, pp. 072504-1–072504-3, Feb. 2007.
- [29] D. K. Ferry and S. M. Goodnick, *Transport in Nanostructures*. Cambridge University Press, 1997.
- [30] Y. Ono, et al., "Si complementary single-electron inverter," *IEDM Technology Dig.*, pp. 367–370, 1999.
- [31] C. P. Heij, P. Hadley, and J. E. Mooij, "Single-electron inverter," *Applied Physics Letters*, vol. 78, pp. 1140–1142, 2001.
- [32] H. Wolf, et al., "Investigation of the offset charge noise in single electron tunneling devices," *Trans. on Instrumentation and Measurement*, vol. 46, no. 2, pp. 303–306, Apr. 1997.

- [33] M. Furlan and S. V. Lotkhov, "Electrometry on charge traps with a single-electron transistor," *Physics Rev. B*, vol. 67, p. 205313, 2003.
- [34] 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, p. 287, Dec. 1999.
- [35] N. M. Zimmerman, et al., "Excellent charge offset stability in Si-based SET transistors," in *Proc. Precision Electromagnetic Measurements*, Nov. 2002, pp. 124–125.
- [36] N. S. Zimmerman, et al., "Excellent charge offset stability in a Si-based single-electron tunneling transistor," *Applied Physics Ltrs.*, vol. 79, pp. 3186–3190, 2002.
- [37] Y. S. Yu, S. W. Hwang, and D. Ahn, "Transient modelling of single-electron transistors for efficient circuit simulation by SPICE," *Electronics Ltrs.*, vol. 152, no. 6, pp. 691–696, Dec. 2005.
- [38] 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, pp. 2028–2032, Apr. 1999.
- [39] J. M. Rabaey, *Digital Integrated Circuits*. Prentice-Hall, NJ, 1998.
- [40] A. DeHon, "Array-based architecture for FET-based nanoscale electronics," *IEEE Trans. Nanotechnology*, vol. 2, no. 1, pp. 23–32, Mar. 2003.
- [41] S. C. Goldstein and M. Budi, "Nanofabrics: spatial computing using molecular electronics," in *Proc. Int. Symp. Computer Architecture*, June 2001, pp. 178–189.
- [42] 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.
- [43] 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, pp. 213 118–1–213 118–3, 2006 May.
- [44] "Xilinx XPower," <http://www.xilinx.com>.
- [45] W. Zhao and Y. Cao, "New generation of predictive technology model for sub-45nm design exploration," in *Proc. Int. Symp. Quality of Electronic Design*, Mar. 2006, pp. 585–590.
- [46] 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, pp. 1131–1144, Oct. 2003.



Changyun Zhu (S'06) received his B.E. and M.E. degrees from Tsinghua University in 2002 and 2005. He is currently a Ph.D. student at Queen's University's Department of Electrical and Computer Engineering. His research interests include computer-aided design of integrated circuits, reliability modeling and optimization, and design for nanotechnologies.



Zhenyu (Peter) Gu (S'04) received his Ph.D. degree from Northwestern University in 2007, and B.S. and M.S. degrees from Fudan University, China in 2000 and 2003, respectively. Gu has published in the areas of behavioral synthesis and thermal analysis of integrated circuits. He is currently with Synopsys Inc, CA.



Robert P. Dick Robert Dick (S'95-M'02) is an Assistant Professor of Electrical Engineering and Computer Science, Northwestern University. He received his Ph.D. degree from Princeton University and his B.S. degree from Clarkson University. He worked as a Visiting Professor at Tsinghua University's Department of Electronic Engineering and as a Visiting Researcher at NEC Labs America. Robert received an NSF CAREER award and won his department's Best Teacher of the Year award in 2004. His technology won a Computerworld

Horizon Award and his paper was selected by DATE as one of the 30 most influential in the past 10 years in 2007. He served as a technical program subcommittee chair for CODES-ISSS and is an Associate Editor of IEEE Transactions on VLSI Systems and serves on the technical program committees of several embedded systems and CAD/VLSI conferences.



Li Shang Li Shang (S'99-M'04) is an Assistant Professor at the Department of Electrical and Computer Engineering, University of Colorado at Boulder. Before that, he was with the Department of Electrical and Computer Engineering, Queen's University. He received his Ph.D. degree from Princeton University, and his B.E. degree with honors from Tsinghua University. His work on hybrid SET/CMOS reconfigurable architecture was nominated for the Best Paper Award at DAC 2007. His work on thermal-aware incremental design flow was nominated for

the Best Paper Award at ASP-DAC 2006. His work on temperature-aware on-chip network has been selected for publication in MICRO Top Picks 2006. He also won the Best Paper Award at PDCS 2002. He is currently serving as an Associate Editor of IEEE Transactions on VLSI Systems and serves on the technical program committees of several design automation conferences. He won his department's Best Teaching Award in 2006. He is the Walter F. Light Scholar.



Robert G. Knobel Robert Knobel received his MSc degree in 1995 from the University of British Columbia and his PhD degree from the Pennsylvania State University in 2000, both in physics, studying high temperature superconductors and diluted magnetic semiconductors, respectively. From 2000–2003 he was a postdoctoral scholar at the University of California, Santa Barbara working with Andrew Cleland on nanomechanical systems and single electron transistors. Since 2003 he has been an Assistant Professor in the Department of Physics, Engineering

Physics and Astronomy at Queen's University, where he studies the quantum limits of measurement in nanoscale devices.