A COMPARATIVE STUDY OF ULTRA-LOW VOLTAGE DIGITAL CIRCUIT DESIGN

Aaron Arthurs, Justin Roark, and Jia Di

Computer Engineering and Computer Science Department, University of Arkansas
Fayetteville, Arkansas, USA
jdi@uark.edu

ABSTRACT

Ultra-low voltage digital circuit design is an active research area, especially for portable applications such as wearable electronics, intelligent remote sensors, implantable medical devices, and energy-harvesting systems. Due to their application scenarios and circuit components, two major goals for these systems are minimizing energy consumption and improving compatibility with low-voltage power supplies and analog components. The most effective solution to achieve these goals is to reduce the supply voltage, which, however, raises the issue of operability. At ultra-low supply voltages, the integrity of digital signals degrades dramatically due to the indifference between active and leakage currents. In addition, the system timing becomes more unpredictable as the impact of process and supply voltage variations being more significant at lower voltages. This paper presents a comparative study among three techniques for designing digital circuits operating at ultra-low voltages, i.e., Schmitt-triggered gate structure, delay-insensitive asynchronous logic, and Fully-Depleted Silicon-on-Insulator technology. Results show that despite the tradeoffs, all eight combinations of these techniques are viable for designing ultra-low voltage circuits. For a given application, the optimum circuit design can be selected from these combinations based on the lowest voltage, the dynamic range, the power budget, the performance requirement, and the available semiconductor process node.

KEYWORDS

Ultra-Low Voltage, Asynchronous Logic, Delay-Insensitive, Schmitt-Triggered, Silicon-on-Insulator

1. INTRODUCTION

Digital circuits capable of operating reliably under ultra-low voltages are highly beneficial to portable electronic systems that have strict power constraints. There are trade-offs in battery-powered portable devices among power consumption, performance, and supply voltage. Ultra-low voltage designs significantly reduce active power consumption and are useful in devices that cannot maintain a steady, above-threshold supply voltage. Applications of ultra-low voltage circuits include intelligent remote sensors, energy-harvesting systems, implantable medical devices, and wearable electronics.

However, at such low supply voltages operability becomes an issue. In ultra-low voltage designs, the threshold voltage of the transistor is never reached, i.e., subthreshold operation. In this region, leakage current (I_{off}), rather than active current (I_{on}), drives the transistors. Therefore, the drive strength significantly degrades, which differentiates many circuit characteristics from those working in above-threshold, saturation region. One difference is the increased propagation delay in each logic gate due to the low drive strength, which makes ultra-low voltage circuits only suitable for low-speed applications. Another difference is the need of special gate structures in order to prevent the failure in pulling up/down the gate output. For example, logic gates for subthreshold operation are advised to have no more than two inputs.

DOI : 10.5121/vlsic.2012.3301

The architecture and gate design of a digital system have strong impacts on the scalability of supply voltage (V_{DD}). A Fast Fourier Transform circuit has been developed on a 0.18μm bulk-silicon process to operate at 180mV [1]. Setting the supply voltage for a CMOS inverter to *4kT/q* is described in [2]. Process, voltage, and temperature (PVT) variation induced timing fluctuations have a greater impact on subthreshold circuit operations than that on above-threshold operations. Techniques such as doping profile modification and adaptive body-biasing have been proposed as means to mitigate such timing fluctuations. Similar to other power reduction techniques such as dynamic voltage scaling, forced transistor stacking, and power gating [3-5], maintaining the integrity of the digital signal is critical to achieve reliability under ultra-low voltages.

Up to now, ultra-low voltage digital circuit design is still a relatively new research area and there is no universally adopted conclusion on which design methodology is the best for a given application. This paper is to provide a comparative study on the combination of process-, architecture-, and circuit-level techniques in order for researchers to evaluate the effectiveness and efficiency of each solution, as well as the tradeoffs among them. The process-level comparison is between a bulk-silicon process from IBM and a Fully-Depleted Silicon-on-Insulator (FD-SOI) process from MIT Lincoln Laboratory [6]. The architecture-level comparison is between the traditional clocked synchronous logic and a delay-insensitive asynchronous logic, NULL Convention Logic [7]. The circuit level comparison is between the regular static gate design and a Schmitt-triggered inverter based gate structure [8]. The metrics include active energy, leakage power, and performance.

2. BACKGROUND

2.1. Asynchronous Logic and NULL Convention Logic (NCL)

Asynchronous circuits are by definition circuits that do not require an external mechanism (e.g., clocks) to control the flow of data processing. Delay-insensitive (DI) asynchronous logic is a subset of asynchronous paradigm, which is *correct-by-construction* and does not require timing analysis. NCL is a quasi-delay-insensitive style of asynchronous logic which, for the mostly negligible compromise of the interconnect isochronic fork assumption, has practical use in designing modern and future digital circuits [8]. The benefits of NCL include high energy efficiency, robust circuit operation, flexible timing requirement, PVT variation tolerance, and low noise/emission.

In addition to DATA 0 and 1, NCL introduces a third signal state called NULL, which indicates that data is not ready yet. A common implementation of NCL is dual-rail logic where each signal is composed of two wires that indicate either NULL or a single bit of data. NCL circuits operate by alternating DATA and NULL wavefronts, transitioning between data validity and invalidity to separate multiple sets of data. In order to ensure the validity of circuit outputs, NCL requires that all inputs transition to DATA (NULL) before processing the subsequent NULL (DATA).

NCL specifies a pipeline framework that is similar to that of synchronous logic, as shown in Figure 1. In an NCL pipeline, a circuit is divided into multiple stages, and each stage is sandwiched between two DI registers. A DI register, unlike a synchronous register, is controlled by input data from the previous stage (or primary input) and a completion signal from the next stage (or primary output). When the DI register receives a completion signal from the next stage, it latches either DATA or NULL from its input and propagates the completion signal to the previous stage. In a cycle of operation of an NCL pipeline, each stage 1) requests for DATA from their previous stage; 2) processes DATA from the previous stage and outputs results to the next stage; 3) waits for the next stage to acknowledge the results; 4) requests for NULL from the previous stage; 5) propagates NULL to the next stage; and 6) waits for the next stage to

acknowledge NULL. In short, DATA/NULL propagates forward followed by an acknowledgment of receiving DATA/NULL.

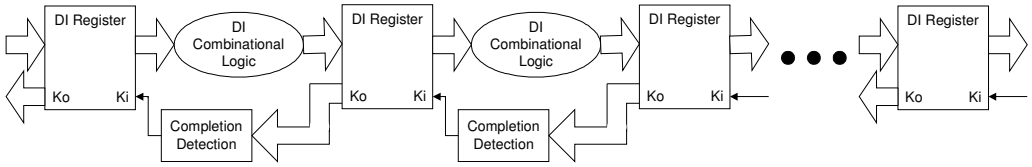

Figure 1. NCL Architecture.

The basic building blocks of NCL circuits are 27 threshold gates [6]. A threshold gate is a state-holding circuit that specifies a Boolean expression, which triggers output assertion and keeps the output asserted until all gate inputs are unasserted. As shown in Figure 2, the generic threshold gate is THmn gate, where $1 \leq m \leq n$. THmn gates have n inputs; at least m of the n inputs must be asserted before the output will become asserted. Once a gate's output is asserted, all n inputs must be unasserted for the output to be unasserted. The gate holds its current state through hysteresis, a feedback function that ensures complete transitions to either the DATA (asserted) or NULL (unasserted) state. For example, a TH34w2 gate has four inputs and a threshold of three; the first input has a weight of two, and the remaining three inputs each has a weight of one. The Boolean expression that asserts the TH34w2 gate's output is *ab+ac+ad+bcd*. Once asserted, the output can only be unasserted if all inputs are unasserted. Static CMOS threshold gates are composed of networks that either change the output or hold the output. As shown in Figure 3, the *RESET* block forces the output to be logic 0 and the *SET* block forces the output to be logic 1. The *Hold0* and *Hold1* blocks maintain the output unchanged when the number of logic 1's among the inputs is between 0 and the gate's threshold.

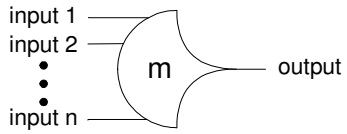

Figure 2. THmn Threshold Gate.

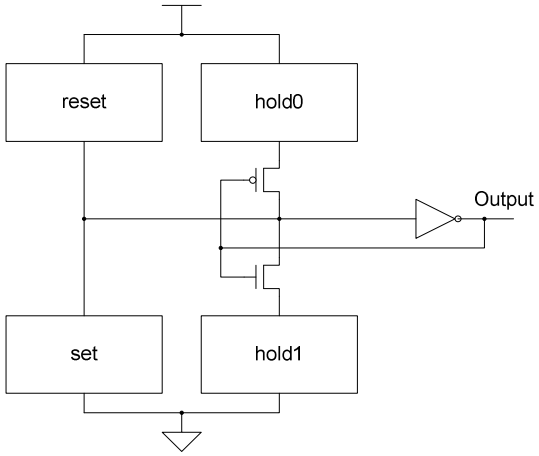

Figure 3. Static Threshold Gate Structure.

2.2. Schmitt-Triggered Gate Structure

The Schmitt-triggered gate design, introduced in [8], actively suppresses leakage current through either the pull-up or pull-down network. Following the static CMOS implementation of a Schmitt trigger, the pull-up and pull-down networks are duplicated and stacked in series. The node in between the original and duplicate networks becomes a virtual supply, which is also connected to the opposite supply voltage through a power-gating transistor. The gate output is fed back to control each virtual supply's power-gating transistor. While the gate inputs select which network is active, the gate output feedback drives the inactive network's supply to the same voltage as the gate output, i.e., positive feedback. Figure 4 shows a general schematic of a Schmitt-triggered gate.

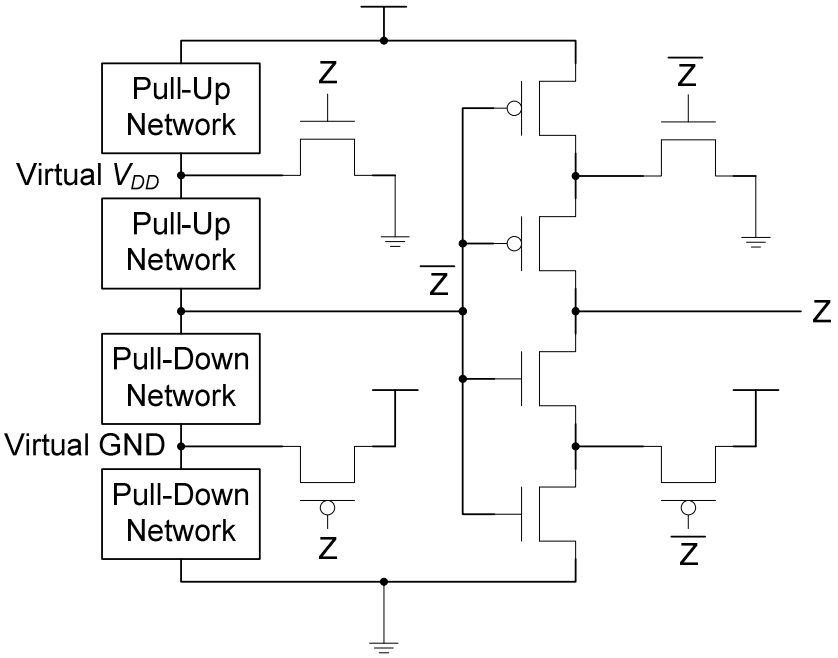

Figure 4. Schmitt-Triggered Gate Structure.

Correct operation of a Schmitt-triggered gate requires that the active network overpowers the power-gating transistor. Conversely, the power-gating transistor must suppress leakage current paths through the inactive network. Weak power-gating transistors are utilized in order for the active network to override the gate output feedback; otherwise, the gate output will not switch state.

Ideally, each virtual supply swings between the respective network supply voltage and the opposite supply voltage, depending on which network is active. However, in reality each virtual supply alternates between their network supply voltage and a voltage level in between supplies. For example, the virtual ground in the pull-down network alternates between ground and a voltage level higher than ground. Experiments show that each virtual supply is biased towards their network supply voltage due to weak power-gating transistors.

2.3. Fully-Depleted Silicon-On-Insulator

SOI technologies fall under two categories: partially-depleted and fully-depleted (PD-SOI and FD-SOI). The difference between PD-SOI and FD-SOI processes is the thickness of the SOI and buried oxide layers. FD-SOI uses thin SOI and buried oxide layers such that the depletion region

extends throughout the transistor's floating body and partially into the substrate underneath the buried oxide layer. For digital circuits, full depletion is preferred over partial depletion because full depletion removes all free charged carriers from the body when forming the inversion layer (channel). FD-SOI requires little or no doping in the transistor's body. The combination of a neutrally charged body and the buried oxide significantly reduces the drain-to-substrate parasitic capacitance, requiring less energy to switch the transistor. Since the body is lightly and uniformly doped, it takes little effort to form the inversion layer, and threshold variation is reduced because light doping minimizes random dopant fluctuation. When an FD-SOI transistor is inactive, the drain-to-body leakage current is negligible due to the body being electrically isolated and having little intrinsic charge.

3. TEST CIRCUITS AND SIMULATION SETUP

The test circuit selected is an IEEE single-precision (32-bit) floating-point coprocessor that supports addition, subtraction, and multiplication. Figure 5 shows its block diagram. This coprocessor is implemented in eight versions, corresponding to all eight combinations of the three techniques, i.e., bulk-Si static synchronous, bulk-Si Schmitt-triggered synchronous, bulk-Si static NCL, bulk-Si Schmitt-triggered NCL, FD-SOI static synchronous, FD-SOI Schmitt-triggered synchronous, FD-SOI static NCL, and FD-SOI Schmitt-triggered NCL.

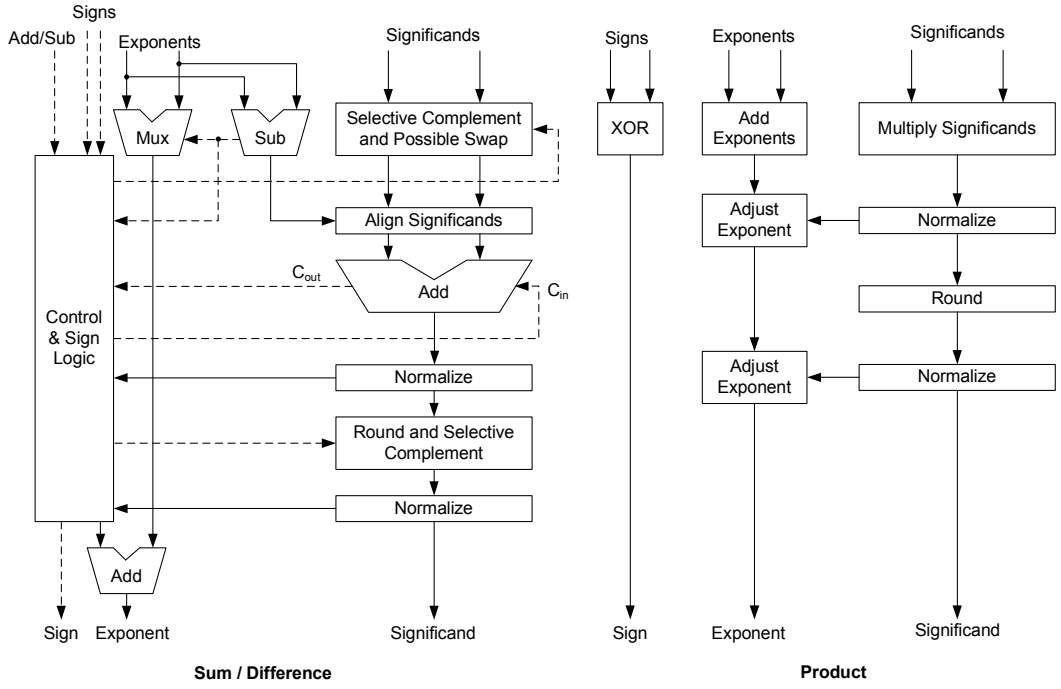

Figure 5. Block Diagram of the Test Circuit.

The methodology to size the transistors in a Schmitt-triggered gate is currently a trial-and-error process that depends on the target supply voltage, the fan-out of each gate, and each transistor's input (gate) capacitance. The transistor sizes in each network are counterbalanced against the feedback transistors until the gate output can swing sufficiently, at least between 20 and 80 percent of the supply voltage. Based on the results from [8], 65 mV is chosen as the target supply voltage for each Schmitt-triggered gate. On the other hand, static CMOS gates are sized to balance output rise and fall times at an above-threshold supply of 1V.

Cadence UltraSim is chosen to run transistor-level simulations because it is optimized for large-scale mixed-signal circuits. UltraSim employs circuit partitioning and simplified MOSFET models in order to improve simulation performance. There are several levels of accuracy that range from BSIM MOSFET models with no partitioning to digital table-lookup models with hierarchical circuit partitioning. Tuning the simulator is a trade-off between simulation performance and accuracy that depends on the type of verification required. In this paper, functioning testing and power analysis are the two main objectives of simulation. While functional testing can be performed at any accuracy level, power analysis at minimum requires the analog table-lookup models for MOSFETs. UltraSim must therefore be set to 'mixed-signal' mode that uses analog table-lookup models and circuit partitioning.

The IBM 8RF-DM 130nm process is chosen for bulk-Si, and the MIT Lincoln Lab 150nm xLP process is selected for FD-SOI. The base MOSFET models used between the two processes differ. UltraSim is set to 'mixed-signal' mode for the 130nm bulk-silicon process and 'digital-accurate' mode for the 150nm xLP process. In 'digital-accurate' mode, UltraSim uses digital table-lookup models that are not suitable for power analysis. The reason for using 'digital-accurate' mode is that the xLP MOSFET devices have irresolvable convergence issues with UltraSim's analog table-lookup models. Another option considered is simulating the xLP process with the BSIMSOI MOSFET models, but simulation performance degrades significantly and becomes impractical for a large digital circuit such as the IEEE floating-point coprocessor.

Each NCL coprocessor connects to a Verilog-A NCL controller that supplies input vectors, performs handshaking, and checks the results. The NCL controller also measures the amount of time taken to process each set of inputs, including the NULL cycle. Each synchronous coprocessor connects to a Verilog-A synchronous controller. The synchronous controller drives a digital clock at 50 percent duty cycle at an adjustable frequency, supplies input vectors on each negative clock edge, and checks the results on falling clock edges after the coprocessor pipeline is primed. Between the coprocessor-under-test and Verilog-A controller are input buffers, which act as input transition filters to ensure correct input timing at any given supply voltage. Both the NCL and synchronous controllers provide the same set of input values that consist of three additions, three subtractions, and three multiplications as shown in Table 1.

Table 1. Input/Output IEEE Single-Precision Floating-Point Vectors

Vector	Operation	X (input)	Y (input)	Z (output)
1	Addition	0x447CC000	0x448AE000	0x4504A000
2	Addition	0xC070B7FE	0xC12E057D	0xC16A337D
3	Addition	0x490D7523	0xD011CC0C	0xD011C9D6
4	Subtraction	0x447CC000	0x448AE000	0xC2C80000
5	Subtraction	0xC070B7FE	0xC12E057D	0x40E3AEFA
6	Subtraction	0x490D7523	0xD011CC0C	0x5011CE42
7	Multiplication	0x447CC000	0x448AE000	0x49891CA8
8	Multiplication	0xC070B7FE	0xC12E057D	0x4223A238
9	Multiplication	0x490D7523	0xD011CC0C	0xD9A12032

All coprocessors undergo a discrete supply voltage sweep, from 50mV to 1V, to determine their operating supply range. At certain supply voltages, the minimum required clock period and throughput are recorded. Active energy and leakage power are measured by integrating the supply current over the appropriate period of time. Equations 1 and 2 measure active energy and leakage

power respectively. Active energy is measured during circuit operation while leakage power is calculated over a one-second period after circuit operation.

$$ActiveEnergy = Vdd\int_{t1}^{t2} i(t)dt \quad (1)$$

$$LeakagePower = \frac{Vdd\int_{t2}^{t3} i(t)dt}{t3-t2} \quad (2)$$

4. RESULTS AND ANALYSIS

During simulations, it is discovered that all four Schmitt-triggered designs operate within a much narrower V_{DD} range compared to the static counterparts, with the highest operational V_{DD} below the lowest operational V_{DD} of static designs. In other words, there is no overlap in terms of operational V_{DD} points between the two design groups. Therefore, the results for these two groups are presented separately. In addition, all four bulk-Si designs are simulated using UltraSim's 'mixed-signal' mode, for which MOSFETs are modelled with pre-calculated tables to speed up BSIM equation calculations. However, as stated in the previous section, convergence problems occur when simulating the four FD-SOI designs using 'mixed-signal' mode. A less accurate mode, 'digital-accurate', which bypasses BSIM equations using simplified equations and digital look-up tables, has to be used to simulate these designs. Therefore, the results from these simulations are presented separately and are not meant to be compared with those of the other four circuits.

Table 2 shows the comparison between the two bulk-Si static designs. Several observations can be made: 1) both designs function properly from 1V V_{DD} to 125mV, which is a wide dynamic range. This is due to the reliable static gate structures; 2) at all V_{DD} points the synchronous design has better result in active energy and leakage power. This is due to the smaller area of synchronous design. The average performance data of the two designs are similar. At ultra-low voltages, the performance penalty of NCL's four-phase handshaking protocol is less critical since the data processing time inside logic gates becomes significantly longer. In fact, it is expected that if process variation factors are taken into account, the NCL design should have substantial advantage in performance over the synchronous design. The clock period of synchronous design needs to account for the worst-case pipeline stage delay, but the NCL handshaking protocol absorbs delay variations across all stages and exhibits average-case performance.

Table 2. Bulk-Si Static Design Comparison

V_{DD} (mV)	Active Energy (pJ)		Leakage Power (nW)		Average Performance (µS)	
	Synchronous	**NCL**	**Synchronous**	**NCL**	**Synchronous**	**NCL**
500	48	119	697	1121	1	0.6
300	59	195	304	516	10	11.8
200	246	255	176	296	100	85.6
150	341	469	114	207	200	234.7
125	417	602	99	168	300	377.5
115	Failure	Failure	Failure	Failure	Failure	Failure

Table 3 shows the comparison between the two bulk-Si Schmitt-triggered designs. Several observations can be made: 1) both designs operate at lower V_{DD} points than static counterparts. This is due to the Schmitt-triggered gate structure; 2) the NCL design has a wider dynamic range

than the synchronous design because of its robust timing requirement; 3) the active energy, leakage power, and the average performance data of the two designs are comparable, while the NCL design is slightly better. This means under ultra-low V_{DD}, energy and delay caused by the Schmitt-triggered structure inside each gate are dominant factors compared to the logic portion.

Table 3. Bulk-Si Schmitt Design Comparison

V_{DD} (mV)	Active Energy (pJ)		Leakage Power (nW)		Average Performance (µS)	
	Synchronous	**NCL**	**Synchronous**	**NCL**	**Synchronous**	**NCL**
125	Failure	Failure	Failure	Failure	Failure	Failure
115	Failure	12680	Failure	293	Failure	4500
100	Failure	15160	Failure	252	Failure	6303
90	17070	14060	257	223	5000	6630
80	13137	12980	227	196	10000	7013
75	Failure	12440	Failure	182	Failure	7249
70	Failure	Failure	Failure	Failure	Failure	Failure

Table 4 shows the comparison between the two FD-SOI static designs. Several observations can be made: 1) both designs function properly from 1V V_{DD} to 150mV, which is a wide dynamic range. This is again due to the reliable static gate structures; 2) the active energy, leakage power, and the average performance data of the two designs are comparable at 500mV V_{DD}, while the NCL design is much better at 150mV V_{DD}. However, as stated before, these data are less accurate (e.g., the active energy data for the Boolean design at 150mV V_{DD}) due to the use of a less-accurate simulation mode.

Table 4. FD-SOI Static Design Comparison

V_{DD} (mV)	Active Energy (pJ)		Leakage Power (nW)		Average Performance (µS)	
	Synchronous	**NCL**	**Synchronous**	**NCL**	**Synchronous**	**NCL**
500	162	183	398	358	1	0.7
150	557	114	33	22	1000	347.8
115	Failure	Failure	Failure	Failure	Failure	Failure

Table 5 shows the comparison between the two FD-SOI Schmitt-triggered designs. The two designs have different dynamic ranges, which only overlap at 125mV V_{DD}. As mentioned previously, data from the FD-SOI process are less accurate (e.g., the average performance data for the NCL design at 150mV V_{DD}) due to the use of a less-accurate simulation mode.

Table 5. FD-SOI Schmitt Design Comparison

V_{DD} (mV)	Active Energy (pJ)		Leakage Power (nW)		Average Performance (µS)	
	Synchronous	**NCL**	**Synchronous**	**NCL**	**Synchronous**	**NCL**
200	Failure	Failure	Failure	Failure	Failure	Failure
150	Failure	17950	Failure	151	Failure	12571
125	4655	12130	70	150	5000	8674
100	6515	Failure	50	Failure	10000	Failure
90	5239	Failure	41	Failure	10000	Failure
80	Failure	Failure	Failure	Failure	Failure	Failure

3. CONCLUSIONS

This paper presents a comparative study of three potential circuit design techniques for ultra-low voltage digital circuits, i.e., Fully-Depleted Silicon-on-Insulator, delay-insensitive asynchronous logic (specifically, the NULL Convention Logic), and Schmitt-triggered gate structure. Eight IEEE single-precision floating-point coprocessors have been designed applying different combinations of these three techniques. IBM 8RF 130nm process is used as the example of bulk-Si process, and MIT Lincoln Lab's xLP FD-SOI 150nm process is used as the example of FD-SOI process.

Cadence UltraSim simulation results show that 1) the static designs have much wider dynamic ranges than the Schmitt-triggered designs; 2) the Schmitt-triggered designs are capable of operating at much lower V_{DD} points than the static designs; 3) the bulk-Si static synchronous design has better energy and power figures than the NCL counterpart, while the performance of the two is comparable; 4) the bulk-Si Schmitt-triggered NCL design has wider dynamic range than the synchronous counterpart, while the energy/power/performance figures of the two are comparable; and 5) due to convergence problem, the FD-SOI designs are simulated using a less-accurate mode, thereby the data presented can only be used for reference purpose.

Overall, all these combinations are viable for designing ultra-low voltage digital circuits. For a given application, the optimum circuit design methodology can be selected from these combinations based on the lowest voltage, the dynamic range, the power budget, the performance requirement, and the available semiconductor process node.

REFERENCES

[1] Wang, A.; Chandrakasan, A.; "A 180-mV subthreshold FFT processor using a minimum energy design methodology," IEEE Journal of Solid-State Circuits, vol.40, no.1, pp. 310- 319, Jan. 2005.

[2] Bryant, A.; Brown, J.; Cottrell, P.; Ketchen, M.; Ellis-Monaghan, J.; Nowak, E.J., "Low-power CMOS at Vdd=4kT/q," Device Research Conference, 2001, vol., no., pp.22-23, 2001.

[3] L. H. Chandrasena and M. J. Liebelt, "Energy minimization in dynamic supply voltage scaling systems using data dependent voltage level selection," Proceedings. ISCAS 2000 Geneva, The 2000 IEEE International Symposium on Circuits and Systems, vol.2, 2000, pages 525-528.

[4] M. C. Johnson, D. Somasekhar, and K. Roy, "Leakage control with efficient use of transistor stacks in single threshold CMOS," Proceedings of Design Automation Conference, 36th, 1999, pages 442-445.

[5] V. Khandelwal and A. Srivastava, "Leakage control through fine-grained placement and sizing of sleep transistors," ICCAD, 2004, pages 533-536.

[6] S. Vitale, P. Wyatt, N. Checka, J. Kedzierski, and C. L. Keast, "FDSOI Process Technology for Subthreshold-Operation Ultralow-Power Electronics," Proceedings of IEEE, Vol. 98, No. 2, pp. 333-342, February 2010.

[7] K. Fant and S. A. Bradt, "NULL Convention Logic: a complete and consistent logic for asynchronous digital circuit synthesis". Proceedings of International Conference on Application Specific Systems, 1996, pages 261-273.

[8] N. Lotze and Y. Manoli, "A 62mV 0.13µm CMOS Standard-Cell-Based Design Technique using Schmitt-Trigger Logic." International Solid-State Circuits Conference (ISSCC), 2011.

Authors

Aaron Arthurs received B.S. and M.S. degrees in Computer Engineering from the University of Arkansas in 2004 and 2007, respectively. He is currently a Ph.D. student in Computer Engineering at the same university. His research interests include digital IC design for ultra-low power, extreme environment, and embedded systems.

Justin Roark received B.S. degree in Electrical Engineering from the University of Arkansas in 2011. He is currently a M.S. student in Computer Engineering at the same university. His research interests include delay-insensitive asynchronous IC design for ultra-low power and radiation hardness.

Dr. Jia Di received his B.S. and M.S. degrees from Tsinghua University, Beijing, China, in 1997 and 2000, respectively, and his Ph.D. in Electrical Engineering from the University of Central Florida, USA, in the year of 2004. He is currently an Associate Professor in the Computer Science and Computer Engineering department of the University of Arkansas. His research interests include asynchronous logic circuit design for ultra-low power, extreme temperature, radiation-hardening, and hardware security.