

Clock Network Design for Ultra-Low Power Applications

Mingoo Seok, David Blaauw, Dennis Sylvester

EECS, University of Michigan, Ann Arbor, MI, USA

mgseok@umich.edu

ABSTRACT

Robust design is a critical concern in ultra-low voltage operation due to large sensitivity to process and environmental variations. In particular, clock networks need careful attention to ensure robust distribution of well-defined clock signals to avoid setup and hold time violations. In this paper, we investigate the design methodology of robust clock networks for ultra-low voltage applications. A case study shows that an optimally-chosen clock network improves skew variation by 36× and energy consumption by 49%, compared to a typical clock network. Additionally, the impact of supply voltage and technology scaling on the optimal clock network construction is investigated.

Categories and Subject Descriptors

B.7 [Integrated Circuits]: General

General Terms: Design

Keywords: ultra-low power, clock network, robust design

1. INTRODUCTION

Recently, energy-constrained applications such as biomedical and environmental sensing systems have gained a large amount of attentions [1][2][6]. Supply voltage scaling is one of the most well-known methods to achieve these Ultra-Low Power (ULP) systems due to the quadratic savings in dynamic energy. A number of authors [3][4] have suggested that energy-optimal designs should employ a scaled down supply voltage to near or below the threshold voltage (V_{th}) until the increase of leakage energy consumption offsets the dynamic energy savings. At these supply voltages, which we refer to as Ultra-Low Voltages (ULV) hereafter, 10-20X energy savings can be achieved [4].

However, the scaled supply voltage makes design less robust due to the reduced on-current to off-current ratio of transistors. The robustness can be further compromised by process and environmental variations since the subthreshold current varies exponentially with variations such as random V_{th} mismatch [14]. However, attempts to improve robustness often lead to higher energy consumption: for example, MOSFET upsizing to mitigate random V_{th} mismatch. In that sense, achieving low power and robustness together poses a challenge for designing ULV circuits.

In order to achieve an ultra-low power and robust system, clock network design is critical. With the highest switching activity, the clock network consumes up to ~40% of total dynamic power [5]. With similar trends in ULV regimes, clock networks make a large impact on total energy consumption, requiring additional design efforts. Along with the low power requirements, the clock network should be designed for robustness. As shown in EQ1, skew should be minimized and well-defined against process and environment variations, otherwise the design can have short paths

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

ISLPED'10, August 18–20, 2010, Austin, Texas, USA.

Copyright 2010 ACM 978-1-4503-0146-6/10/08...\$10.00.

and functional failures [6]. Additionally, slew needs to be well-controlled since it degrades the setup and hold time of registers.

$$T_{cq,reg}(T_{clk,slew}) + T_{min,logic} \geq T_{hold}(T_{clk,slew}) + T_{clk,uncertainty} \quad [EQ1]$$

There have been several works on designing clock networks in ULV regimes. In [7] and [8] the authors designed charge-pump based clock buffers to enhance robustness. Although the robustness is improved, both designs incur energy overhead and require custom clock buffers. The authors in [9] seek to tighten slew variations at ULV regimes by constraining slew differently at each clock tree level. However, they did not consider skew, a key metric for clock network design.

Therefore, in this paper we investigate a low power and robust clock network design methodology that avoids custom gates while considering energy, skew, slew and their variability at ULV regimes. We start by comparing various clock networks for a generic design. Several levels of buffered and un-buffered H-Trees and a simple signal-route clock network are studied. Then, device and interconnect process variations are analyzed for their impacts on clock networks. In addition, the impact of supply voltage and technology scaling on clock network is investigated.

From these studies, we find that the design methodology of clock network in ULV regimes should be radically different due to the negligible interconnect resistance. Typically, in super-threshold regimes designers add buffers to mitigate interconnect delay. However it becomes disadvantageous in ULV regimes since buffer delay varies with process, temperature, and supply voltage variations and degrades skew/slew robustness, while reducing already negligible skew contributions of interconnects. Therefore, we propose a different method using no buffers inside clock networks for minimizing skew/slew variations and energy consumption. As a case study, several clock networks for a 16b MSP430-compatible microprocessor [10] are implemented and simulated in SPICE. We confirm that an optimally-selected clock network greatly outperforms other typical clock networks in skew/slew variations and energy consumptions.

2. Clock Network Comparison at Low V_{dd}

2.1 Comparison Frameworks

Figure 1 shows clock networks for a simplified design where 4096 master-slave flip-flops or sinks are placed regularly in 1.4 x 1.4mm² area. (Only 64 sinks are shown in Figure 1 for clarity) These are used in the simulations throughout Sections 2 and 3. The candidate networks for comparison are signal-route, and 1-4 level un-buffered/buffered H-Trees (3- and 4-level H-Trees are not shown in Figure 1). The signal-route clock network routes the clock like an ordinary signal with no balancing attempted. At the bottom of the H-Tree, sinks are also routed as signals. The signal-route network can be considered as a 0-level un-buffered H-tree to simplify plotting. Grid and grid-tree hybrid clock networks are not considered in this paper since they often incur large power penalties. The chosen sink density is based on a survey of two microprocessors: 32b ARM Cortex M3 microcontroller [11] and




Figure 1. Clock network topologies

16b MSP430-like microcontroller, which is used in the case study of Section 4.

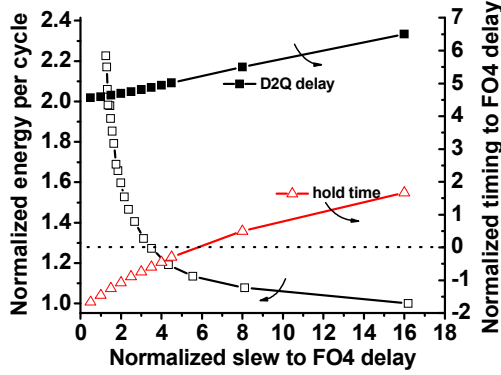


Figure 2. Tradeoff between slew and clock network energy.




Figure 3. Wire length and required buffer.

We assume 4FO4 as slew constraint (i.e., clock signal transitions from 10% to 90% in 4 fanout-4 delays) since this represents a balance between energy consumption of clock buffers and slew of clock signals as shown in Figure 2. At this slew rate, the D-Q delay and hold time remain in reasonable ranges. However other slew rates can be chosen based on application requirements. Clock drivers are sized up to achieve the same slew rates at sinks for all topologies considered. For the signal-route and un-buffered H-Trees, a large central driver (usually consisting of several cascaded buffers of increasing driving strengths) is sized to switch the entire clock network at the slew constraint. For the buffered H-Trees, many small buffers, sized for the iso-slew constraint, are distributed inside the clock network. Here, we assume a fanout-4 ratio to drive buffers. Higher level H-Trees require more buffers

although individual buffers are smaller at the iso-slew constraint (Figure 3).

Since interconnect resistance is negligible in ULV designs, minimum width metal is used for clock networks, reducing energy consumption. The clock net is shielded by supply voltage nets to minimize crosstalk noise. Since the clock network is shielded, the wire capacitance can be well-defined regardless of surrounding wires and their switching activities.

We use a 0.3V supply voltage and a 0.18 μ m CMOS technology, which is a typical technology and supply voltage combination for energy-optimal ULP designs [12][13]. However, we also consider the impact of higher supply voltages and a scaled technology later in Section 3.

2.2 Comparison at Nominal Conditions

Given the framework of Section 2.1, we compare the energy consumption and global skew for the clock networks with SPICE simulations. In this paper we consider energy dissipated in clock buffers and interconnects. Energy consumed internal to registers including local clock drivers that sharpen clock signal edges are not included as these will be constant across network topologies. Figure 4 shows that higher level H-trees consume more energy due to longer interconnect. For higher level trees, un-buffered networks consume slightly more energy than buffered counterparts due to the iso-slew constraint. Since wire RC increases quadratically with the length of the wire, distributed buffers in the buffered H-Trees are more energy-efficient for achieving the same slew than central drivers in un-buffered H-Trees.

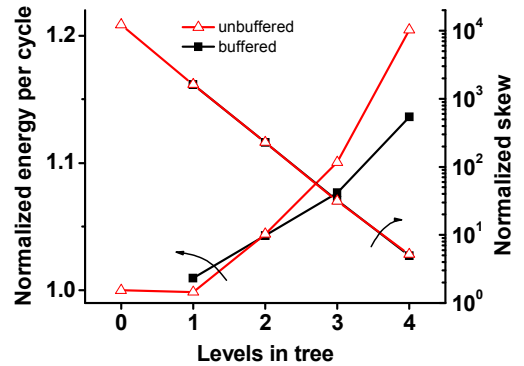


Figure 4. Energy and skew comparison of clock networks.

Skew is improved exponentially as we increase the tree level since the area of the subsection, which is proportional to the amount of skew, becomes 4 \times smaller per level (Figure 1(b) and (c)). Theoretically, a 6-level H-Tree eliminates any path mismatch for

the 4096 sinks. The signal-route or 0-level H-tree exhibits the largest skew due to the longest path mismatch as expected.

2.3 Impact of MOSFET Process Variations

It is well known that MOSFET parameter variations, such as random V_{th} mismatch, have an exponential effect on gate delay at ULV regimes [14]. In a clock network, delay variation degrades skew and slew from the expected values, causing both performance degradation and functional failure. Although clock buffers use relatively large MOSFETs, they still show considerable delay variations from random V_{th} mismatch due to the high sensitivity of subthreshold current. Therefore, it is critical to consider the effect of process variations on clock networks for robust operation at ULV regimes. We do not include the effects of temperature and supply voltage variations for simplicity since these affects skew and slew in the similar fashion as process variations do, only worsening skew and slew variations further.

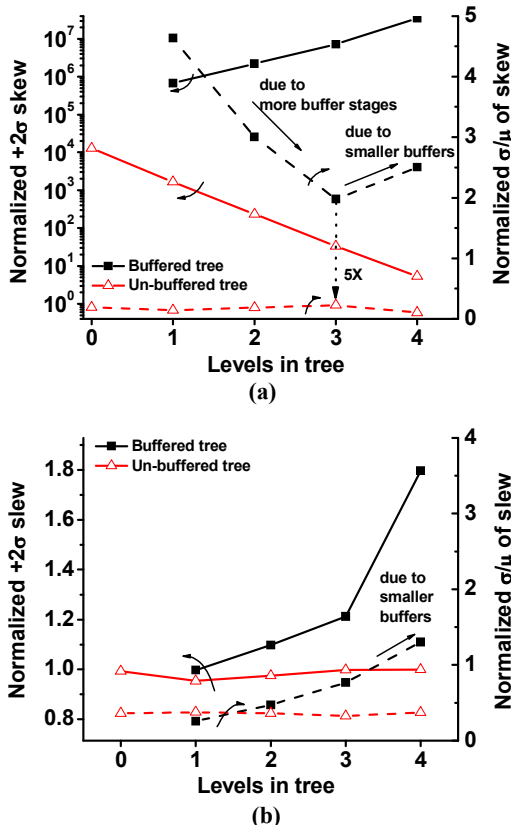


Figure 5. (a) Skew and (b) Slew with MOSFET variations.

In this section, we consider the impact of MOSFET process variations on clock network designs. Monte Carlo simulations with random MOSFET mismatch on the clock networks are performed. We use SPICE model with embedded statistical data from foundries. Global variation is ignored since it has a negligible impact on skew and can be tuned out using global parameters such as body biasing and voltage scaling at a reasonable overhead [1][20][21][22].

Figure 5(a) shows the $+2\sigma$ value of skew across different clock network topologies. Compared to the case with no process variation, buffered trees exhibit several orders of magnitude larger degradation in skew. This is because the buffer delay which used

to be cancelled among buffers starts to contribute to skew.

Another interesting observation is that the $+2\sigma$ skew increases for higher level buffered-H-Tree while the opposite trends are observed without process variations. It implies that adding buffers in ULV regimes has no contribution in mitigating path RC mismatch but only degrades the total path delay. We will discuss the issue of driving interconnects in Section 2.5. The σ/μ of skew for the buffered H-Trees is also at least 5X worse than un-buffered topologies.

Figure 5(b) shows the slew having similar trends to the skew. The un-buffered topologies show a good robustness on slew control while buffered trees have degraded and more variable slew as we increase tree level.

The σ/μ for skew and slew shows different trends with tree level. Figure 5(a) shows that the skew variability first reduces since clock signals travel through more stages of buffers and thus delay variations are averaged. However it starts to increase at level 3 due to the smaller and thus more process-sensitive buffers. However, slew variability is mostly determined by the final buffers which directly drive sinks. Therefore, it has no averaging effect, different from the skew case.

2.4 Impact of Interconnect Process Variations

Interconnect variation is another source of performance variability in scaled CMOS technologies. However, it can be considered as a secondary effect in ULV regimes since its impact on delay is roughly linear, while device variations have exponential effects. Therefore, we apply the worst case interconnect variation to the studied clock networks, and evaluate whether their skew contribution is significant compared to the contribution of MOSFET variations.

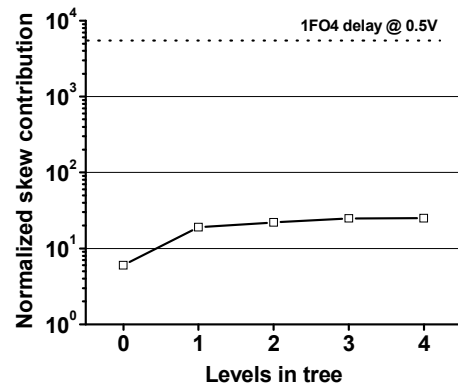


Figure 6. Skew contribution from interconnect variations.

Finding the worst case corner for interconnects is difficult and requires detailed information from physical design since a fixed process variation (e.g., thinner interconnect) might cause two opposite effect on delay, depending on whether the particular wire delay is capacitance dominated or resistance dominated [15]. However, at ULV regimes, the worst case for interconnects is better defined since interconnect delays are always capacitance dominated due to the negligible interconnect resistance. The worst interconnect corner for skew can be defined between two non-overlapping paths experiencing min and max interconnect capacitance (provided by the foundry design kit). For example, in Figure 1(b), if the path from n0 to n1 has max capacitance and the

path from n0 to n64 has min capacitance, two sinks at n1 and n64 will experience the largest skew.

With the worst interconnect corner, we run SPICE simulations to evaluate the contribution of interconnect variations on skew, compared to MOSFET variations. Simulations show that it takes only 10-15% of total skew across the 1-4 level buffered H-Trees. For un-buffered trees, the interconnect variations might seem to be contributing non-negligible skew. However, this is mainly because the large central drivers are little affected by process variations. As shown in Figure 6, the absolute amount of skew contribution from interconnect variations is much smaller than gate delays in ULV regimes. Additionally, the worst case corner for interconnects is highly pessimistic. Therefore, we can simply ignore interconnect variations without too much loss of accuracy.

2.5 Driving Interconnects at ULV Regimes

At super-threshold regimes, repeaters are commonly added in the middle of a long interconnect, which gives better performance [16]. The benefit comes from shorter interconnect segments (i.e. quadratically smaller wire RC) and sharper slew rate to the input of a following buffer. As shown in Figure 7, adding one buffer in the middle of a long interconnect improve performance at $V_{dd}=1.8V$ for wires $> 3mm$.

However, these benefits are no longer valid at ULV regimes. First the delay penalty of adding buffers is often much larger than the reduction of wire RC. Figure 7 shows that adding buffers cannot reduce delay even for interconnects longer than 20mm. EQ2 using the results of [17] can easily verify the results. Slew rate is also negligibly affected by interconnects since the total resistance ($R_{fet}+R_{wire}$) is dominated by MOSFET resistance.

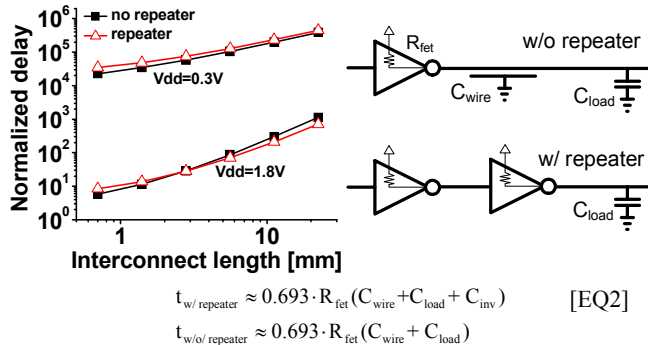


Figure 7. Driving a long interconnect without repeaters

Technically, adding buffers to drive a long interconnect is only harmful at ULV regimes since they act as another source of variation. It also consumes more energy.

3. Impact of Voltage and Technology Scaling

In Section 2, we considered 0.3V supply voltage and a $0.18\mu m$ technology. While this represents the optimal choice [12] for most energy-constrained systems, other application spaces may require higher performance and therefore prefer different supply voltages and technologies. In this section we discuss the impact of the supply voltage and technology on the optimal selection of clock networks.

3.1 Supply Voltage Scaling

Figure 8 shows the results of Monte-Carlo iterations with random MOSFET variations on 1-level buffered and un-buffered H-Trees over supply voltages. One interesting observation is that there is a crossover voltage at $\sim 0.85V$ in Figure 8(a). At $V_{dd} < 0.85V$, the un-

buffered network outperforms in $+2\sigma$ skew and σ/μ of skew. However, the buffered tree performs better at $V_{dd} > 0.85V$. This is because the buffers in the buffered H-Tree become less sensitive to process variations at higher supply voltage. Additionally, buffers start to drive interconnects strong enough to mitigate some of path RC mismatches, resulting in improved skew-related metrics.

Slew also has a crossover voltage at 0.6V in Figure 8(b). At $V_{dd} > 0.6V$, a degradation in $+2\sigma$ slew is observed for the un-buffered H-Tree since interconnect resistance is no longer negligible compared to the MOSFET resistance of the clock drivers. However, the buffered H-Tree maintains the similar slew across the supply voltages due to shorter interconnect.

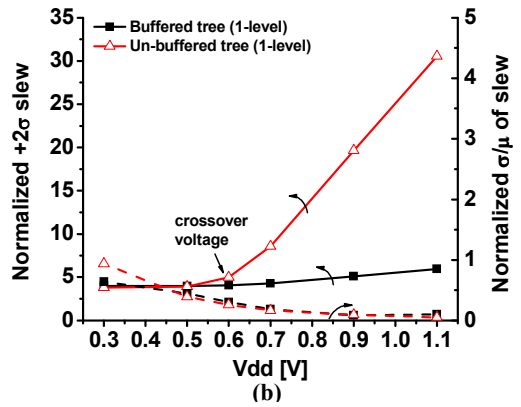



Figure 8. Impact of voltage scaling on skew and slew.

3.2 Technology Scaling

In Section 3.1, we observed crossover voltages in skew and slew. Technology scaling also acts in the similar way since scaled technology has more resistive interconnects and less resistive MOSFETs with lower V_{th} .

Figure 9 shows the MOSFET and interconnect resistance trends in two different technologies. We assume that the interconnect length is scaled with the channel length of technology for the same design. Still, increase in wire resistance is observed. The difference between wire and device resistance reduces from 16000X at $V_{dd}=0.3V$ and $0.18\mu m$ technology to only 17X at $V_{dd}=0.5V$ and 65nm technology.

We additionally run the Monte-Carlo simulations on the 1-level buffered and un-buffered H-Trees to identify the crossover

voltages of $+2\sigma$ skew and slew in a 65nm General Purpose (GP). GP process is chosen as a more pessimistic option for un-buffered topology, compared to Low Power (LP) process CMOS technology. We use the statistical data supplied by the foundry design kits.

Figure 10 shows the trends of crossover voltages over two technology nodes. Both skew and slew crossover voltages appear at lower voltage for scaled technology due to the reduced difference of resistance between devices and interconnects. Slew might be the limiter for un-buffered clock networks since its crossover voltage appears earlier than the skew counterpart. One might want to move the slew crossover voltage to higher V_{dd} regime to exploit less skew and skew variability from un-buffered clock networks. Since we use minimum width interconnects for low power, thick top-level metals and wider metals can be considered as an option to improve slew. However it might have energy overhead, requiring a careful evaluation.




Figure 9. Resistance scaling across technologies.

4. Clock Network Design for a 16b MSP430-compatible Microprocessor at ULV regimes

In Sections 2 and 3, the simplified design, where sinks are regularly placed, is used to study on designing robust yet low power clock networks. In this section, we will continue our investigations on clock networks using more practical design, a 16b MSP430-compatible microprocessor.

We first characterize standard cells at $V_{dd}=0.3V$ in a $0.18\mu m$ CMOS technology with 6 metal stack. The core of the microcontroller is synthesized and APR-ed (Automatic Placement & Route) with industrial EDA tools. Then, 7 different clock networks including signal-route and 1-3 level buffered/un-buffered H-Trees clock networks are implemented. Fourth and fifth metal layers are used to implement clock networks. It is shielded with V_{dd} net. One example APR-ed design employing 3-level buffered H-Tree is shown in Figure 11. The traces for H-Trees are highlighted for visibility. The total footprint is $0.6 \times 0.6 mm^2$. Interconnects, buffers and flip-flops in clock networks are extracted with parasitic capacitance and resistance in a SPICE format for simulations. Mismatch Monte-Carlo iterations are performed to evaluate skew, slew and energy for each clock network.

As shown in Figure 12, 1-level un-buffered H-Trees can improve 4 orders of magnitude in $+2\sigma$ skew and $\sim 36X$ in σ/μ of skew, compared to the worst case clock network. The 1-3 level buffered clock networks can produce up to 5X worse skew from the values of design phase, which can cause functional failures after

fabrications. Note that the worst clock network in the comparison, which is the 3-level buffered H-Tree, might be chosen as an optimal network in super-threshold regimes [18][19], confirming the importance of the clock network selection at ULV regimes.

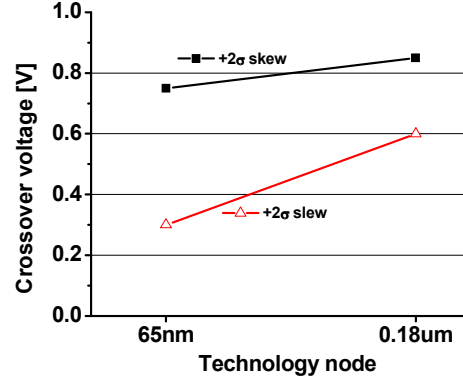


Figure 10. Crossover voltages for $+2\sigma$ skew and slew.




Figure 11. Layout view for the APR-ed microprocessor with 3-level buffered H-Tree clock network.

The $+2\sigma$ slew and its variability are plotted in Figure 12(b), which has the similar trends to skew. The 3-level buffered H-Tree can have 28% higher slew from variation than the design values, resulting in less robust design.

Energy consumption for each clock network is also compared in Figure 12(c). Higher level trees consume more energy. Although the 1-level un-buffered consumes the second least energy after the signal-route clock network (Signal-route clock network consumes 4% less energy than 1-level un-buffered H-Tree), it consumes $\sim 49\%$ less energy than the 3-level buffered H-Tree. The 3-level buffered H-Tree consumes relatively larger energy than expected from the simplified analysis in Section 2 since the individual buffer strength scales more slowly than the simplified design at the slew constraint.

In Section 3, we observed that a clock network which used to be optimal at low V_{dd} loses optimality when supply voltage goes up to a certain point, which we define as a crossover voltage. Here we also sweep the supply voltage to find the crossover voltages. As shown in Figure 13, the 1-level un-buffered H-Tree is skew-optimal choice at $V_{dd}=0.3-1.0V$. At $V_{dd}>1.0V$, the 1-level buffered H-Tree becomes skew-optimal. The crossover voltage




Figure 12. Skew, slew and energy comparison for the clock networks of the 16b microcontroller.




Figure 13. Optimal clock network across supply voltages.

for +2 σ slew appears at V_{dd}=0.6V, which is lower than the skew crossover voltage. Thick metal layers or non-minimum width interconnects might be considered to alleviate slew degradation at the cost of energy consumption

5. Conclusions

In this paper, we investigate on designing a low power yet robust clock network at ULV regimes. After comparing several clock networks in energy consumption, skew, and slew in the contexts of both simplified and practical designs, we find that a radically different methodology of using no buffers inside clock network is beneficial at ULV design. A case study with a 16b microcontroller shows that the optimally-chosen clock network at an energy-optimal operating point can improve +2 σ skew by 4 orders of magnitude, skew variability by 36X, and energy consumption by 49%, compared to the clock network which can be considered a typical practice in super-threshold voltage designs. Impact of process variation and supply voltage and technology scaling on ULV clock network design are also investigated.

Acknowledgement

The authors acknowledge the support of the Gigascale Systems Research Center, two of the five research centers funded under the Focus Center Research Program, a Semiconductor Research Corporation program.

References

- [1] G. Chen, et al, "Millimeter-Scale Nearly Perpetual Sensor System with Stacked Battery and Solar Cells," *ISSCC*, pp.288-289, 2010
- [2] M. Seok et al, "The Phoenix Processor: A 30pW Platform for Sensor Applications," *Symp. on VLSI Circuits*, pp. 188-189, 2008.
- [3] B. Zhai, et al, "Theoretical and Practical Limits on Dynamic Voltage Scaling," *Design Automation Conference*, 2004
- [4] B. Calhoun, et al, "Characterizing and Modeling Minimum Energy Operation for Subthreshold Circuits," *ISLPED*, 2004
- [5] N. Magen, et al, "Interconnect power dissipation in a Microprocessor," *Int. Workshop on SLIP*, 2004
- [6] J. Kwong, et al, "A 65nm Sub-Vt Microcontroller with Integrated SRAM and Switched-Capacitor DC-DC Converter," *ISSCC*, 2008
- [7] J. Kil, et al, "A High-Speed Variation-Tolerant Interconnect Technique for Sub-Threshold Circuits Using Capacitive Boosting," *ISLPED*, pp.67-72, Aug. 2006
- [8] S. Lin, et al, "A Capacitive Boosted Buffer Technique for High Speed Process Variation Tolerant Interconnect in UDVS application," *ASP-DAC*, pp. 304-309, 2008
- [9] J. Tolbert, et al, "Slew-Aware Clock Tree Design for Reliable Subthreshold Circuits," *ISLPED*, pp.15-20, Aug. 2009
- [10] OpenCores, <http://www.opencores.org>
- [11] ARM.com, <http://www.arm.com/products/CPUs>
- [12] M. Seok, et al, "Optimal Technology Selection for minimizing energy and variability in low voltage applications," *ISLPED*, pp.9-14, Aug. 2008
- [13] D. Bol, et al, "Technology Flavor Selection and Adaptive Techniques for Timing-Constrained 45nm Subthreshold Circuits," *ISLPED*, pp.21-26, Aug. 2009
- [14] B. Zhai, et al, "Analysis and Mitigation of Variability in Subthreshold Design," *ISLPED*, 2005
- [15] Y. Liu, et al, "Impact of Interconnect Variations on the Clock Skew of a Gigahertz Microprocessor," *DAC*, pp.168-171, 2000
- [16] A. Chandrakasan, et al, "Design of High Performance Microprocessor Circuits," *IEEE Press*, pp.352-375, 2000
- [17] T. Sakurai, "Closed-form Expressions for Interconnect Delay, Coupling, and Crosstalk in VLSI's," *Trans. on Electron Devices*, vol.40, no.1, Jan, 1993
- [18] C. Yeh, et al, "Clock Distribution Architectures: A Comparative Study," *ISQED*, 2006
- [19] P. Restle, et al, "Designing the Best Clock Distribution Network," *Symp. on VLSI Circuits*, pp.2-5, 1998
- [20] Y. Ramadass, et al, "Minimum Energy Tracking Loop with Embedded DC-DC Converter Delivering Voltages down to 250mV in 65nm CMOS," *ISSCC*, pp.64-65, 2007
- [21] M. Wiecekowsky, et al, "A Hybrid DC-DC Converter for Sub-Microwatt Sub-1V Implantable Applications," *Int. Symp. on VLSI circuits*, pp.166-167, 2009
- [22] M. Seok, et al, "A 0.5V 2.2pW 2-Transistor Voltage Reference," *Custom Integrated Circuits Conference*, pp. 577-580, 2009