Selective Gate-Length Biasing for Cost-Effective Runtime Leakage Control

Puneet Gupta†, Andrew B. Kahng‡†, Puneet Sharma†, Dennis Sylvester*
†ECE and ‡CSE Departments, University of California at San Diego
*EECS Department, University of Michigan, Ann Arbor
{puneet@ucsd.edu, abk@ucsd.edu, sharma@ucsd.edu, dennis@eecs.umich.edu}

ABSTRACT

With process scaling, leakage power reduction has become one of the most important design concerns. Multi-threshold techniques have been used to reduce runtime leakage power without sacrificing performance. In this paper, we propose *small* biases of transistor gate-length to further minimize power in a manufacturable manner. Unlike multi-V_{th} techniques, gate-length biasing requires no additional masks and may be performed at any stage in the design process.

Our results show that gate-length biasing effectively reduces leakage power by up to 25% with less than 4% delay penalty. We show the feasibility of the technique in terms of manufacturability and pin-compatibility for post-layout power optimization. We also show up to 54% reduction in leakage uncertainty due to inter-die process variation in circuits when biased gate-lengths, versus only unbiased one, are used. Circuits selectively biased show much less sensitivity to both intra and inter die variations.

Categories and Subject Descriptors: B.7.2 [Design Aids]:Layout
General Terms: Algorithms, design, performance
Keywords: Layout, lithography, OPC, leakage, power, manufacturability

1. INTRODUCTION

High power dissipation shortens battery life, reduces circuit performance and reliability, and has a large impact on packaging costs. CMOS circuit power consists of dynamic and static components. Leakage is becoming an ever-increasing component of total dissipated power with its contribution projected to increase from 18% at 130nm to 54% at the 65nm node [23]. This necessitates development of new methods to reduce leakage power.

A number of approaches have been proposed to reduce static leakage power when the system is in *standby* mode. [10] proposed the source biasing principle, where a positive bias is applied in the standby state to the source terminal of an "off" device. [11] suggested using transistor stacks to reduce standby leakage. [12, 13, 14, 15] proposed use of multi-threshold CMOS in which a high-V_{th} CMOS is used to disconnect power supply to a low V_{th} logic module during the standby state. Substrate-bias management for leakage reduction is also proposed in [16].

The only mainstream approach to reduce leakage power during *active*, or *runtime*, mode is the multi-V_{th} manufacturing process [4, 5, 17]. In this approach, cells in non-critical paths are assigned a high V_{th} while cells in critical paths are assigned a low V_{th}. The major drawback to this technique has traditionally been the rise in process costs due to additional steps and masks. The increased costs have been outweighed by the substantial leakage reductions they

DAC 2004, June 7–11, 2004, San Diego, California, USA.

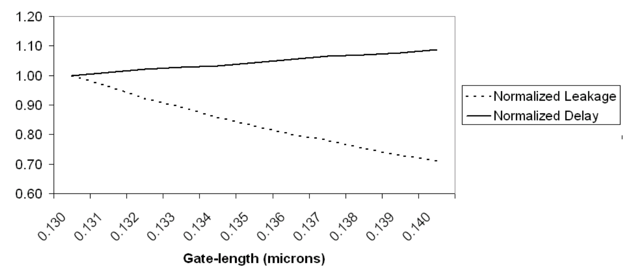

Figure 1: Variation of leakage and delay (each normalized to 1.00) for an NMOS device in an industrial 130nm technology.

provide, and multi-V_{th} processes are now standard. However, a new complication facing multi-V_{th} is the increased variability of V_{th} for low-V_{th} devices [24]. This occurs in part due to random doping fluctuations, as well as worsened DIBL (Drain Induced Barrier Lowering) and short-channel effects (SCE) in devices with lower channel doping. The larger variability in V_{th} degrades the achievable leakage reductions of multi-V_{th} and will only worsen with continued MOS scaling. Moreover, multi-V_{th} methodologies do not offer a smooth tradeoff between performance and leakage power. Devices with different V_{th} typically have a large separation in terms of performance and leakage, for instance a 15% speed penalty with a 10X reduction in leakage for high-V_{th} devices.

The use of longer gate-lengths (L_{Gate}) in transistors within non-critical gates was first described in [1]. In that work, very *large* gate-lengths were considered, resulting in heavy delay and dynamic power penalties. Moreover, cell layouts with larger gate-lengths are not layout-swappable with their nominal versions, which results in substantial ECO (Engineering Change Order) overheads during layout. The variation of delay and leakage with gate-length is shown in Figure 1 for an industrial 130nm process. Note that leakage current flattens out with gate-length beyond 140nm, making L_{Gate} biasing less desirable in that range.

The problem of rising leakage requires new, *manufacturable* techniques. Here, we propose a novel leakage reduction methodology based on small biases to the device gate-length. Contributions of our work include the following.

- A leakage reduction methodology based on less than 10% increase in drawn gate-length of devices.
- A thorough analysis of potential benefits and caveats of such a biasing methodology, including implications of lithography and process variability.
- Experiments and results showing potential benefits of an L_{Gate} biasing methodology in different design scenarios such as dual-V_{th}.

The organization of this paper is as follows. In the next section, we describe the proposed gate-length biasing methodology for leakage reduction. Section 3 gives experiments and results validating the proposed technique. Section 4 analyzes the potential manufacturing and process variation implications of biasing gate-lengths. Finally, Section 5 concludes with a brief description of ongoing research.

2. L$_{GATE}$ BIASING METHODOLOGY

In this section we describe the proposed gate-length biasing methodology. We characterize and then augment a standard-cell library, such that each master also has a *biased* L_{Gate} variant. A sizing tool is then used to incorporate slower but low-leakage cells into non-critical paths, while retaining faster, high-leakage cells in critical paths. Reflecting the experiments below, our discussion focuses on the introduction of a single biased variant for each cell in the library, and on an industry 130nm process technology. Of course, the approach also extends to multiple biased variants.

L$_{Gate}$ Biasing Granularity

Gate-length biasing can be performed at several levels of granularity, namely, technology-level, cell-level and device-level. Finer granularity leads to more difficult implementation but more flexibility in optimization and potentially larger leakage benefits. We have considered the following three levels of biasing granularity.

1. *Technology-Level.* All gates in the library have the same biased L_{Gate}. As a result there are *exactly* two distinct gate-lengths (in a dual-L_{Gate} approach) in the technology library. This approach is similar to the multiple V_{th} approach, in which a higher-V_{th} library is constructed. The benefits of technology-level biasing are ease of implementation, ease of library layout, and potential ease in process tuning.
2. *Cell-Level.* Every library cell master has its own specific biased gate-length. All devices within a given cell share this characteristic L_{Gate}, but different cell masters are allowed to have different biased gate-lengths. Potentially, there can be as many distinct biased gate-lengths as there are masters.
3. *Device-Level.* Ideally, a device-level gate-length biasing approach will allow independent biasing of every gate in the library. However, as this is computationally impractical within our characterization and search framework[1] we restrict ourselves to independent biasing of PMOS and NMOS devices within a cell. In other words, all PMOS devices within a given cell master have the same gate-length bias which is independent of the bias of NMOS devices in the cell as well as the PMOS devices in other cells. This simplification permits us to exhaustively search for the "optimal" biased L_{Gate} for devices. The rationale for this biasing approach is that in complementary MOS technologies, the NMOS devices in a cell typically have identical topology (e.g. series connected for NAND gates) and PMOS devices have identical topology (e.g. parallel connected for NAND gates). Leakage has a strong dependence on topology, with stacked devices leaking much less than unstacked ones [18].

Biased-L$_{Gate}$ Selection

The key question in our methodology is the value of L_{Gate} for each transistor in the cells. We consider less than 10% biasing of the gate-length. The reasons for such a small bias are as follows.

- An increase in drawn dimension that is less than the layout grid resolution (typically 10nm for 130nm technology) ensures pin-compatibility with the unsized version of the cell. This is very important to ensure that multi-L_{Gate} optimizations can be done post-placement or even after detailed routing without ECOs. In this way, we retain the layout transparency that has made multi-V_{th} optimization so adoptable within chip implementation flows. Biases smaller than the layout grid-pitch also ensure design-rule correctness for the biased cell layout, as long as the unbiased version is correct.
- The nominal gate-length of the technology is usually very close to or beyond the "knee" of the leakage vs. L_{Gate} curve. For large bias, the advantage of super-linear dependence of leakage on gate-length is lost. Moreover, dynamic power and delay both increase almost linearly with gate-length. Therefore, small bias gives more "bang for the buck".
- From a manufacturability point of view, having two prevalent pitches (which are not close enough) in the design can harm printability properties (i.e., size of process window). Note that we retain the same poly-pitch as the unbiased version of the cell. There is a small decrease in spacing between gate poly geometries but it is still well within minimum spacing required by the process.

Delay impact of L_{Gate} bias is measured by a simple canonical circuit wherein a minimum-sized inverter is placed as load to the gate under consideration. The gate-length of the inverter is matched to that of the driving gate. Delays for all possible input transitions are then measured, and we compute average rise and fall delays for the cell under test. Leakage is also measured and averaged over all possible input combinations. These average delay and leakage numbers are computed by detailed (HSPICE) circuit simulation for L_{Gate} bias selection.

For the 130nm process that we use in our studies, we search over (positive) bias values from 1nm to 9nm (layout grid pitch is 10nm) in steps of 1nm. For purposes of determining the appropriate gate-length bias we ignore dynamic power impact of biasing (i.e., we only consider the delay/leakage tradeoff). We compute the bias so as to restrict the delay penalty of biasing to a prescribed $delay_{penalty}$ percentage. Both fall and rise delays are not allowed to rise more than this predetermined delay penalty. The bias value is then chosen to minimize leakage power subject to the delay penalty constraint. Our bias selection uses $delay_{penalty} = 10\%$.[2].

Library Generation

An important component of the methodology is layout and characterization of the dual-L_{Gate} library. Since we investigate very small biases to the gate-length, the layout of the biased library cell does not need to change except for simple automatic scaling of dimensions. Moreover, since the bias is smaller than the minimum layout grid pitch, design rule violations are highly unlikely. Of course, after the slight modifications to layout, the biased versions of the cell are put through the standard extraction and power/timing characterization process.

3. EXPERIMENTS AND RESULTS

We now describe our test flow for validation of the L_{Gate} biasing methodology, and present experimental results. We consider up to two gate-lengths and two threshold voltages. We perform experiments for the following scenarios – Single-V_{th}, single-L_{Gate} (SVT-SGL); Dual-V_{th}, single L_{Gate} (DVT-SGL); Single-V_{th}, dual-L_{Gate} (SVT-DGL); Dual-V_{th}, dual L_{Gate} (DVT-DGL).

The dual-V_{th} flow uses high and low values of V_{th} while the single-V_{th} flow uses only the low value of V_{th}. The basic elements of our flow are a dual L_{Gate} library that captures the effects of L_{Gate} biasing on leakage, delay and input capacitance; and a tool to perform leakage-aware sizing.

Dual L$_{Gate}$ Library Characterization

We prune the TSMC 130nm library to contain only eight commonly used cells: INVX4, NANDX4, BUFX4, ANDX6, NORX4, ORX6, AO22X4 and OA22X4. To get the delay and leakage number, HSPICE [25] simulations are run using TSMC 130nm netlists and STMicroelectronics 130nm spice models. [3]

Dual L$_{Gate}$ Optimization

We use a sizer similar to *Duet* proposed in [4]. All cells are sorted in decreasing order of $\Delta leakage \times slack$ where $\Delta leakage$ is the improvement in leakage after a cell is replaced with its less leaky variant, and $slack$ is its timing slack after the replacement has been made. We use *Design Compiler v2003.06-SP1 (DC)* [27] for final validation of all timing and power results as well as computation of dynamic power. [4]

[1] We use exhaustive search to find the best biased gate-length values. When every device L_{Gate} is allowed to vary independently, the search space becomes too large; effective identification of biased variants in this ideal framework is an open direction.

[2] The number 10% is determined empirically. Larger bias can lead to larger per-cell leakage saving at a higher performance cost. However, in a resizing setup (described below) with a delay constraint, the leakage benefit over the whole design can decrease as the number of instances which can be replaced by their biased version (slower but less leaky) is reduced.

[3] The library contains nominal (NMOS: 0.187V, PMOS: -0.16825V) and low (NMOS: 0.107V, PMOS: -0.08825V) V_{th} devices. The nominal gate-length is 130nm.

[4] There is a small mismatch between static timing engines of *DC* and *Duet*. We report results from *DC* only.

Table 1: **Optimum biased L_{Gate} values (in nanometers) at device-level granularity given 10% limit on delay penalty delay. ΔLeakage denotes the (percentage) leakage savings over the corresponding unbiased cell with gate-length equal to 130nm.**

Cell	Low V_{th} P	Low V_{th} N	Low V_{th} % Δ Leakage	Nominal V_{th} P	Nominal V_{th} N	Nominal V_{th} % Δ Leakage
INVX4	137	139	29.66	137	139	34.31
NANDX4	137	139	26.41	136	139	17.10
BUFX4	137	138	27.96	136	139	32.16
ANDX6	139	131	31.60	138	134	27.42
NORX4	137	139	28.69	137	139	18.16
ORX6	136	139	26.15	136	139	25.75
AO22X4	137	138	27.52	136	139	22.82
OA22X4	139	130	28.93	137	136	22.60

Table 2: Normalized critical-path delay, leakage power, and dynamic power results for various V_{th} and gate-length scenarios. The second gate-length is determined by technology-level or device-level L_{Gate} selection.

	SVT-SGL			DVT-SGL		
Test Case	Delay	Leakage	Dynamic	Delay	Leakage	Dynamic
c5315	1	1	1	1.034	0.325	0.974
c6288	1	1	1	1.027	0.557	0.984
c7552	1	1	1	1.009	0.202	0.968
alu128	1	1	1	1.020	0.248	0.971
	SVT-DGL-tech			DVT-DGL-tech		
Test Case	Delay	Leakage	Dynamic	Delay	Leakage	Dynamic
c5315	1.017	0.779	1.034	1.017	0.296	1.004
c6288	1.038	0.857	1.023	1.033	0.534	0.993
c7552	1.018	0.743	1.044	1.028	0.171	1.010
alu128	1.040	0.741	1.044	1.040	0.218	1.004
	SVT-DGL-device			DVT-DGL-device		
Test Case	Delay	Leakage	Dynamic	Delay	Leakage	Dynamic
c5315	1.017	0.789	1.032	1.034	0.299	1.002
c6288	1.022	0.876	1.020	1.027	0.538	0.991
c7552	1.009	0.752	1.018	1.02	0.171	1.010
alu128	1.03	0.753	1.042	1.03	0.221	1.001

Test Cases

We use simple combinational circuits drawn from the ISCAS85 benchmark suite and Opencores [28] as test cases. The four test cases synthesize to 2069 (c5315), 4070 (c6288), 2360 (c7552) and 13279 (alu128) gates. In our results, we report leakage, dynamic power and circuit delay. We do not assume any wire-load models, as a result of which the dynamic power and delay are underestimated.

Results

As described in Section 2, we choose the L_{Gate} bias at the technology-level and at the device-level. We do not present results for cell-level gate-length biasing as this offers no advantage over device-level biasing in terms of quality, nor over technology-level biasing in terms of ease of implementation. The nominal gate-length for the technology is 130nm. The technology-level biased L_{Gate} is calculated to be 136nm based on an allowable 10% delay penalty. Table 1 shows the optimum device-level biased L_{Gate} values and the corresponding leakage power benefit with a delay penalty constraint set to 10%. We see 9%-36% leakage power benefit with less than 10% delay overhead. This strongly supports our hypothesis that small biases in L_{Gate}, intelligently applied, can afford significant leakage savings with virtually no performance impact.

The timing constraint we give to the synthesis tool is very close to the minimum achievable by any combination of threshold voltages and gate-length. Synthesis is performed using low-V_{th}, nominal-L_{Gate} library. For introduction of a V_{th} or L_{Gate} we relax the timing constraint by 2% to give sizing more room to recover power. Results for this delay-constrained sizing for leakage recovery are shown in Tables 2. Adding a gate-length to single V_{th} designs can save 14.3% to 25.9 % leakage power with less than 4% delay penalty. For dual V_{th} implementations the leakage benefit is less than 12%.The dynamic power penalty is less than 3.3% in all cases.

4. PROCESS EFFECTS

In this section, we investigate certain manufacturability and process variability implications of our L_{Gate} biasing approach.

Lithography: Manufacturability

As our method relies on biasing of drawn gate-length, it is important to correlate it with actual printed gate-length on the wafer. This is even more important as the bias we introduce in gate-length is of the same order as typical critical dimension (CD) tolerance in manufacturing processes. Moreover, we expect larger gate-lengths to have better printability properties leading to less CD - and hence leakage - variability. To validate our multiple gate-length approach in a post-manufacturing setup, we follow a reticle enhancement technology (RET) and process simulation flow for an example cell master.

We use the layout of the AND2X6 from TSMC 0.13 μm and perform model-based optical proximity correction (OPC) on it using *Calibre v9.3_2.5* [26].[5] The printed image of the cell is then calculated using *printimage* simulation in Calibre. We measure the gate-length for every device in the cell, for both biased and unbiased versions. The results for the printed gate dimensions are shown in Table 3. As expected, biased and unbiased gate-lengths track each other well. There are some outliers which may be due to simplicity of the OPC model being used. High correlation between *printed* dimensions of biased and unbiased versions of the cells shows that benefits of biasing estimated using *drawn* dimensions will not be lost during the RET and manufacturing flows.

Another potentially valuable benefit of even slightly larger gate-lengths is possible improved printability. Poly spacing is much larger than poly gate-length, so that the process window (which is constrained by the minimum resolvable dimension) tends to be larger as gate-length increases. For example, the depth of focus for various values of exposure latitude with the same illumination system as above for 130nm and 136nm lines is shown in Table 4.[6]

Process Variability

A number of sources of variation can cause fluctuations in gate-length, and hence in performance and leakage. This has been a subject of much discussion in the recent literature (e.g., [21, 20]). Up to 20X variation in leakage has been reported in practice [19]. For leakage, the reduction in variation post-biasing is likely to be substantial as the larger gate-length is closer to the "flatter" region of the leakage vs. L_{Gate} curve. To validate these intuitions, we study the impact of gate-length variation on leakage and performance both pre- and post-biasing using a simple worst-case approach. We assume the CD variation budget to be $\pm 10nm$. The performance and leakage of the test case circuits is measured at the worst-case, nominal and best-case process corners which consider just gate-length variation. This is done for the technology-level L_{Gate} biasing approach as an example. The results are shown in Table 5. For the four test cases, we see a 39% to 54% reduction in leakage power uncertainty caused by linewidth variation. Such huge reductions in uncertainty can potentially outweigh benefits of alternative leakage control techniques. We note that the corner case analysis just models the inter-die component of variation, which typically constitutes half of the total CD variation.

To assess the impact of both within-die (WID) and die-to-die (DTD) components of variation, we run 2000 Monte-Carlo simulations with $\sigma_{WID} = \sigma_{DTD} = 3.33nm$. The variations are assumed

[5] Model-based OPC is performed using annular optical illumination with $\lambda = 248nm$ and $NA = 0.7$.

[6] The process simulation was performed using *Prolith v8.0* [29].

Table 3: Comparison of printed dimensions of unbiased and biased versions of AND2X6. The unbiased nominal gate-length is 130nm while the biased nominal is 136nm. Note the high correlation between unbiased and biased versions.

Device Number	Gate Length (nm) PMOS Unbiased	PMOS Biased	PMOS Diff.	NMOS Unbiased	NMOS Biased	NMOS Diff.
1	125	132	+7	126	132	+6
2	124	126	+2	126	129	+3
3	124	126	+2	126	129	+3
4	121	127	+6	124	130	+6
5	121	127	+6	122	128	+6
6	122	128	+6	122	128	+6
7	125	131	+6	124	131	+7

Table 4: Process window improvement with gate-length biasing. The CD tolerance is kept at 13nm. ELAT=Exposure latitude; DOF=Depth of Focus.

DOF (μm)	ELAT (%) for 130nm	ELAT (%) for 136nm
0.09	7.66	7.71
0.33	6.97	7.04
0.5	5.98	6.23
0.67	4.67	5.02
1	2.06	2.71

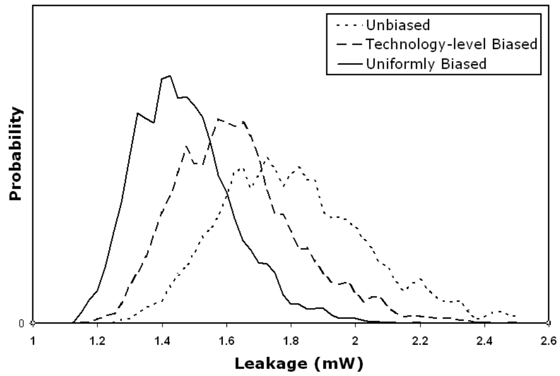

Figure 2: Leakage distributions for unbiased, uniform-biased and technology-level selectively-biased alu128. Note the "left-shift" of the distribution with the introduction of biased devices in the design.

Table 5: Reduction in performance and leakage power uncertainty with biased gate length in presence of inter-die variations. The uncertainty spread is specified as a percentage of nominal. The results are given for nominal V_{th}. Uniform bias is 6nm.

	Circuit Delay (ns)						
	Unbiased			Uniform Bias			% Spread
Test Case	BC	WC	NOM	BC	WC	NOM	Reduction
c5315	0.58	0.76	0.66	0.63	0.81	0.72	0
c6288	1.80	2.35	2.05	1.95	2.52	2.26	-3.6
c7552	1.02	1.35	1.18	1.11	1.46	1.29	-5.7
alu128	0.95	1.25	1.10	1.04	1.35	1.20	-3.3
	Leakage (mW)						
	Unbiased			Uniform Bias			% Spread
Test Case	BC	WC	NOM	BC	WC	NOM	Reduction
c5315	0.289	0.137	0.181	0.214	0.122	0.151	+39.4
c6288	0.579	0.276	0.364	0.430	0.247	0.305	+53.5
c7552	0.322	0.156	0.200	0.240	0.140	0.171	+39.7
alu128	1.936	0.930	1.230	1.440	0.833	1.023	+39.6

to follow a Gaussian distribution with no correlations. We compare the results for three single V_{th} scenarios: unbiased, technology-level biasing and uniform biasing of the entire design by 6nm. Leakage distributions for the test case $alu128$ are shown in Figure 2.

5. CONCLUSIONS AND ONGOING WORK

We have presented a novel methodology that uses selective, *small* L_{Gate} biases to achieve an *easily manufacturable* approach to run-time leakage reduction. For our test cases we have observed the following. (1)The gate-length bias we propose is always less than the pitch of the layout grid. This avoids design rule violations. Moreover, it implies that the biased and unbiased cell layouts are completely pin-compatible and hence layout-swappable. This allows biasing-based leakage optimization to be possible at any point in design flow unlike sizing-based methods. (2) With simple uniform technology-level biasing applied to the entire design 12%-28% leakage improvement can be achieved at the cost of 8%-12% delay penalty and 3%-6% dynamic power penalty. (3) Using simple sizing techniques, we are able to achieve up to 25% leakage savings with just 4% timing and 5% dynamic power overhead. With dual-L_{Gate} libraries constructed with a smaller $delay_{penalty}$ and multiple versions of frequently used cells, the improvements can be much better. (4) The devices with biased gate-length are *more* manufacturable and have a larger process margin than the nominal devices. Biasing does not require any extra process steps unlike multiple-threshold based leakage optimization methods. (5) L_{Gate} biasing leads to more process-insensitive designs with respect to leakage current. Biased designs have up to 54% less leakage worst-case variability in presence of inter-die variations as compared to nominal gate-length designs (6) In presence of both inter- and intra-die CD variations, selective L_{Gate} biasing can yield designs less sensitive to variations.

The variability benefits of gate-length biasing are obvious. Gate-length biasing can offer large leakage reductions for cost-sensitive low volume ASICs where mask and process costs account for a large part of total design cost. Biases which give less than 5% delay penalty can be explored to give a more fine-grained control of the leakage-delay tradeoff. Our ongoing work is along the following directions- (1) Construction of effective biasing-based leakage optimization heuristics. We are also investigating the use of more than two gate-lengths for more frequently used and leaky cells in the library such as inverters and buffers. (2) L_{Gate} selection at true device-level granularity. Here intelligent search methods will be needed as brute force exhaustive search is computationally infeasible. Moreover, such device-level leakage optimization can be easily extended to incorporate choice of V_{th} to yield an integrated $L_{Gate} - V_{th}$ leakage reduction flow. (3) Evaluating the impact of biasing on leakage at future technology nodes wherein leakage is a much bigger issue than in 130nm.

6. ACKNOWLEDGEMENTS

Some of our early leakage vs. delay simulation studies were performed by P. Gupta during a summer internship at IBM T. J. Watson research lab. We would like to thank F.-L. Heng and R. Puri for helpful discussions.

7. REFERENCES

[1] N. Sirisantana, L. Wei and K. Roy, "High-Performance Low-Power CMOS Circuits Using Multiple Channel Length and Multiple Oxide Thickness", *Proc. ICCD*, 2000, pp. 227-232.
[2] D. Lee and D. Blaauw, "Static Leakage Reduction Through Simultaneous Threshold Voltage and State Assignment", *Proc. DAC*, 2003, pp. 192-194.
[3] L. Wei, K. Roy and C. K. Koh, "Power Minimization by Simultaneous Dual-V_{th} Assignment and Gate-sizing", in *Proc. CICC*, 2000, pp. 413-416.
[4] S. Sirichotiyakul, T. Edwards, C. Oh, R. Panda and D. Blaauw., "Duet: An Accurate Leakage Estimation and Optimization Tool for Dual-V_{th} Circuits", in *TVLSI*, Vol. 10, No. 2, April 2002, pp. 79-90.
[5] L. Wei, Z Chen, M. Johnson, K. Roy and V. De, "Design and Optimization of Low Voltage High Performance Dual Threshold CMOS Circuits", in *Proc. DAC*, 1998, pp. 489-494.
[6] L. Wei, Z. Chen, K. Roy, M. Johnson, Y. Ye and V. K. De, "Design and Optimization of Dual-Threshold Circuits for Low-Voltage Low-Power Applications", in *TVLSI*, Vol. 7, No. 1, March 1999, pp. 16-24.
[7] V. Sundararajan and K. Parhi, "Low Power Synthesis of Dual Threshold Voltage CMOS VLSI Circuits", in *Proc. ISLPED*, 1999, pp. 139-144.
[8] S. Kim, S. Kosonocky, W. Hwang, and Y. Shin, "Long-Term Power Minimization of Dual-V_{th} CMOS circuits", in *IEEE Proc. ASIC/SOC*, 2002, pp. 323-327.
[9] K. S. Khouri and N. K. Jha, "Leakage Power Analysis and Reduction During Behavioral Synthesis", in *Proc. ICCD*, 2000, pp.561-564.
[10] M. Horiguchi, T. Sakata and K. Itoh, "Switched-Source-Impedance CMOS Circuit for Low Standby Sub-threshold Current Giga-Scale LSI's", in *JSSC*, 1993, pp. 1131-1135.
[11] Y. Ye, S. Borkar and V. De, "A New Technique for Standby Leakage Reduction in High-Performance Circuits", in *Proc. Symp. on VLSI Circuits*, 1998, pp. 40-41.
[12] S. Mutoh, T. Douseki, Y. Matsuya, T. Aoki and S. Shigematsu, "1-V Power Supply High-Speed Digital Circuit Technology with Multithreshold-Voltage CMOS", in *JSSC*, 1995, Vol. 30, No. 8, pp. 847-854.
[13] J. Kao, S. Narendra and A. Chandrakasan, "MTCMOS Hierarchical Sizing Based on Mutual Exclusive Discharge Patterns", in *Proc. DAC*, 1998, pp. 495-500.
[14] S. Mutoh, S. Shigematsu, Y. Matsuya, H Fukada, T. Kaneko and J. Yamada, "1V Multithreshold-Voltage CMOS Digital Signal Processor for Mobile Phone Application", in *JSSC*, 1996, pp. 1795-1802
[15] S. Shigematsu, S. Mutoh, Y. Matsuya, Y. Tabae and J. Yamada, "A 1-V High-Speed MTCMOS Circuit Scheme for Power-Down Application Circuits", in *JSSC*, 1997, pp. 861-869.
[16] Y. Oowaki et al., "A sub-0.1μm Circuit Design with Substrate-Over-Biasing", in *ISSCC*, 1998, pp. 88-89.
[17] M. Ketkar and S. Saptnekar, "Standby Power Optimization via Transistor Sizing and Dual Threshold Voltage Assignment", in *Proc. ICCAD*, 2002, pp. 375-378.
[18] S. Mukhopadhyay, C. Neau, R.T. Cakici, A. Agarwal, C.H. Kim and K. Roy, "Gate Leakage Reduction for Scaled Devices Using Transistor Stacking", *TVLSI*, 11(4), 2003, pp. 716-730.
[19] S. Borkar, T. Karnik, S. Narendra, J. Tschanz, A. Keshavarzi and V. De, "Parameter Variations and Impact on Circuits and Microarchitecture", in *Proc. DAC*, 2003, pp. 338-342.
[20] Y. Cao, P. Gupta, A.B. Kahng, D. Sylvester and J. Yang, "Design Sensitivities to Variability: Extrapolations and Assessments in Nanometer VLSI", *Proc. ASIC/SOC* , 2002, pp. 411-415.
[21] R. Rao, A. Srivastava, D. Blaauw and D. Sylvester, "Statistical Estimation of Leakage Current Considering Inter- and Intra-Die Process Variation", *Proc. ISLPED*, 2003, pp. 84-89.
[22] ITRS 2003, *http://public.itrs.net*
[23] S. Narendra, D. Blaauw, A. Devgan and F. Najm, "Leakage Issues in IC Design: Trends, Estimation and Avoidance", *Proc. ICCAD* (tutorial), 2003.
[24] R. Puri, IBM Corporation, *Personal Communication*.
[25] http://www.synopsys.com/products/mixedsignal/hspice/hspice.html
[26] http://mentor.com/calibre/datasheets/opc/html/
[27] http://www.synopsys.com/products/logic/design_compiler.html
[28] http://www.opencores.org/projects/
[29] http://www.kla-tencor.com