

Impact of Gate-Length Biasing on Threshold-Voltage Selection

Andrew B. Kahng^{†‡} Swamy Muddu[†] Puneet Sharma[†]

abk@ucsd.edu

smuddu@ucsd.edu

sharma@ucsd.edu

UCSD [†]ECE and [‡]CSE Departments, La Jolla, CA

Abstract

Gate-length biasing is a runtime leakage reduction technique that leverages on the short-channel effect by marginally increasing the gate-length of MOS devices to significantly reduce their leakage current for a small delay penalty. The technique was shown to work effectively with multi-threshold-voltage assignment, the only mainstream approach for runtime leakage reduction. Typically, designers use threshold voltages selected by the foundries to optimize their designs. Higher threshold-voltage devices, that are less leaky but slow, are assigned to non-critical paths and lower threshold-voltage devices, that are fast but leaky, are assigned to critical paths.

In this paper we study the effect of modifying threshold voltages set by the foundry on leakage reduction achieved on three large, real-world designs. We assess the impact of the availability of gate-length biasing on threshold voltage selection. We achieve comparable leakage reductions when foundry-set dual threshold voltages are used with biasing than when foundry-set triple threshold voltages are used without biasing. Our results indicate that leakage reductions can be improved if threshold voltages are carefully chosen considering the availability of gate-length biasing. We also observe that foundry-set threshold voltages are not optimal for achieving best possible leakage reductions.

1. Introduction

Leakage power is one of the key challenges faced by the semiconductor industry today. Leakage variability is also of great concern and along with timing variability determines the yield. Leakage current has three main components: (1) subthreshold leakage, (2) gate leakage, and (3) reverse biased drain substrate and source-substrate junction band-to-band tunneling leakage [2]. Among these, subthreshold leakage is the dominant contributor to total leakage at 130nm and is forecast to remain so in the future [2], especially at operating temperatures that are almost always higher than room temperature.

A lot of research has been done towards leakage reduction and the proposed techniques can be classified into two classes depending on whether they reduce *standby* leakage or *runtime* leakage. Standby leakage reduction techniques are used for devices that are not in operation (i.e., are in standby mode) while runtime techniques reduce leakage of active devices. Several techniques have been proposed for standby leakage reduction such as body biasing or *VTMOS*, *MTC-MOS*, source biasing, use of transistor stacks, input vector

control, etc. However, very few approaches have been proposed for runtime leakage reduction and threshold-voltage (V_{th}) assignment is the only mainstream approach. Other recently proposed runtime leakage reduction techniques include gate-length biasing [5] and adaptive body biasing. Gate-length biasing has attracted a lot of interest and several IDMs have recently published papers on the use and behavior of biasing [4, 7, 8, 11]

Multi-threshold voltage designs offer high performance with low runtime and standby leakage power. Low V_{th} devices offer high performance but are leakier. While high V_{th} devices are slow but leak less. Circuit optimization techniques use low V_{th} devices (cells) on critical paths while high V_{th} is used for all other devices (cells). Multiple V_{th} 's are manufactured by altering the doping levels over the diffusion regions in MOS devices. Extra masks and lithography steps are required for achieving multiple V_{th} 's, therefore the process costs increase. The increased costs have been outweighed by the substantial leakage reductions they provide, and dual- and triple- V_{th} processes are now standard.

Gate-length biasing leverages on short-channel effect (SCE) to marginally increase the gate-length to alter the threshold voltage and reduce leakage significantly at the cost of small delay increase. Technologies below the 90nm node utilize super-halo doping which causes reverse short channel effect (RCSE) to mitigate SCE up to some extent. However, we have found gate-length biasing to significantly reduce leakage for all the two 130nm and two 90nm processes that we have seen so far (excluding BPTM). Recent reports from leading integrated design and manufacturing (IDM) houses indicate SCE to dominate the V_{th} roll-off curve at least until the 45nm technology node [10, 8, 4, 11]. We, however, note that the V_{th} roll-off curve must be understood to assess the feasibility of this approach.

In this paper we:

- study the simultaneous use of V_{th} assignment and gate-length biasing. Specifically, we empirically assess the impact of availability of gate-length biasing on V_{th} selection, a decision that is typically made by the foundries.
- evaluate the effectiveness of foundry V_{th} 's on leakage improvements of large designs by comparing against different synthesized V_{th} 's.
- analyze the impact of leakage improvement obtained due to gate-length biasing on top of different V_{th} 's.

Multiple foundry V_{th} 's increase manufacturing cost and impact wafer yield due to increase in number of masks and

processing steps respectively. In this paper, we show that a combination of gate-length biasing and dual V_{th} assignment provides cost-effective leakage reduction comparable to that of triple V_{th} assignment.

The organization of this paper is as follows. In the next section we present the advantages of multi- V_{th} and gate-length biasing and motivate the need for their use with each other. Section 3 describes our methodology to alter foundry-set V_{th} for use in our studies. In Section 4 we describe our experimental setup and present results. Section 5 concludes with a summary of our results.

2. Simultaneous Threshold Voltage Assignment and Gate Biasing

Simultaneous use of V_{th} assignment and gate biasing was proposed in [6]. The authors reported leakage savings of 2% to 30% on small testcases using a sensitivity-based greedy optimizer. In this section we compare the two techniques, present the advantages of their simultaneous use, and motivate the need for judicious selection of the different threshold voltages that is aware of the availability of gate-length biasing.

Multi- V_{th} (achieved through multiple doping concentrations) and gate-length biasing both alter V_{th} to tradeoff leakage power against delay. Both techniques can be applied post-layout when accurate timing information is available and do not require iterations with the synthesis, placement, and routing loop. However, there are certain differences between the two techniques. Multi- V_{th} increases manufacturing costs, both NRE and recurring, because extra masks and process steps are required per additional V_{th} . Gate-length biasing does not increase process costs but can potentially reduce the effectiveness of resolution enhancement techniques (RET) due to the presence of multiple gate-lengths (critical dimensions). If the gate-length increase is small ($< \sim 10\%$), biasing can be performed during the RET phase.

Gate-length biasing has several disadvantages over multi- V_{th} . Multi- V_{th} achieves a more favorable tradeoff between leakage and delay than biasing. Also, biasing increases input gate-capacitance marginally that can affect the delay of fanin cells due to increased loading. Increased gate-capacitance also increases dynamic power making biasing usable only when activity factors are small (< 0.1). Due to these shortcomings gate-length biasing cannot be used a replacement for higher-cost multi- V_{th} .

Advantages of gate-length biasing, in addition to lower process costs, include finer control over leakage-delay tradeoff and significant reduction in leakage variability. Figure 1 shows the tradeoff between delay and I_{off} as the gate-length is increased up to 10% for three foundry-set V_{th} . This finer control allows circuit-level leakage optimizers to reduce leakage more than what is possible with the coarse control provided by V_{th} assignment only. Essentially, a simple optimizer can utilize the residual timing slacks after V_{th} assignment using gate-length biasing to reduce leakage; a sophisticated optimizer of course can perform cleverer.

Leakage variability is another major challenge faced by manufacturers today. A number of sources of variation can cause fluctuations in gate-length, and hence in performance and leakage. Up to $20\times$ variation in leakage has been reported in production microprocessors [3]. Gate-length bias-




Figure 1: Off-current and delay of an NMOS device as its V_{TH0} is modified for HVT, SVT, and LVT.

ing has been shown to significantly reduce leakage variability. Up to 40% reduction in leakage variability was achieved in [6] due to gate-length biasing and [9] showed biasing to be the most effective leakage variability reduction technique. The primary reason for leakage variability reduction due to gate-length biasing is that due to the short channel effect, variations in gate-length translate to smaller leakage variations when the gate-length is larger. Due to the improvement in leakage and leakage variability provided by gate-length biasing at no additional manufacturing cost, we propose to use gate-length biasing with V_{th} assignment.

Table 1 presents a comparison of leakage reductions obtained on our testcases (details in Section 4) obtained with (a) two foundry-set V_{th} 's, LVT and SVT, (b) two foundry-set V_{th} 's, LVT and HVT, (c) three foundry-set V_{th} 's LVT, SVT and HVT, (d) LVT and SVT with biasing and (e) LVT and HVT with biasing. We can observe that foundry dual- V_{th} combined with biasing achieves reductions comparable to foundry triple- V_{th} . The main advantage of choosing biasing over dual- V_{th} is the reduced cost and yield impact during manufacturing.

We now study the impact of availability of gate-length biasing on V_{th} selection. Typically foundries select V_{th} 's for each process such that high leakage reductions are achieved for all designs that use them. To achieve high leakage reductions: (1) a large number of cells must be convertible to higher V_{th} 's, and (2) per-cell leakage reduction due to higher V_{th} assignment should be large. In real-world designs, a large percentage of paths are timing critical and most have small timing slacks. Thus higher V_{th} 's must have a small delay penalty to allow a lot of cells to get converted to higher V_{th} . Unfortunately, small delay penalty yields small per-cell leakage reduction on higher V_{th} assignment. Increasing the higher V_{th} increases the per-cell leakage reduction, but decreases the number of cells that can be converted to higher V_{th} . This occurs because of the larger delay penalty associated with the increased higher V_{th} . Similarly, lowering the higher V_{th} would allow more cells to get higher V_{th} assigned but the per-cell leakage reduction would decrease. Therefore selection of a good set of V_{th} 's is dependent on the leakage-delay tradeoff due to V_{th} , and slack distribution and structure of the design's netlist. Assessing the impact of availability of biasing on V_{th} assignment is complex and difficult, if not impossible, to generalize for different testcases. We empirically study the effect on real-world testcases and

Table 1: Leakage reduction with: (a) two foundry-set V_{th} 's, LVT and SVT (b) two foundry-set V_{th} 's, LVT and HVT (c) three foundry-set V_{th} 's LVT, SVT and HVT (d) V_{th} 's, LVT and SVT with biasing (e) V_{th} 's, LVT and HVT with biasing.

Circuit	Leakage Reduction (%)				
	LVT + SVT	LVT + HVT	LVT + SVT + HVT	LVT + SVT + Biasing	LVT + HVT + Biasing
AES	65.89	64.03	74.41	70.15	67.94
OR1200	59.86	70.83	71.46	65.89	73.04
DES3	70.07	72.43	80.28	74.41	75.55

industry SPICE models.

3. Threshold Voltage Customization

We use TSMC 100nm process that has three foundry-set V_{th} 's. Our experiments require more V_{th} 's because we study the impact of changing the available set of V_{th} 's on leakage reduction. We artificially generate V_{th} 's by modifying the V_{TH0} parameter in the SVT SPICE device model. To test the accuracy of this method, we modify the V_{TH0} of (foundry-set) SVT gradually until I_{off} and I_{on} (or delay) characteristics are similar to those of (foundry-set) HVT and LVT. Figures 2 and 3 plot the delay and I_{off} of the NMOS and PMOS devices in an inverter cell ($INVX4$) respectively. The NMOS and PMOS widths are $1.16\mu m$ and $1.64\mu m$ respectively and the inverter is loaded with $5fF$ capacitance. The figures also show the I_{off} and delay tradeoffs when V_{TH0} is changed for HVT and LVT SPICE models. As can be seen, we are able to match the delay and I_{off} of HVT and LVT very accurately by changing V_{TH0} of SVT.

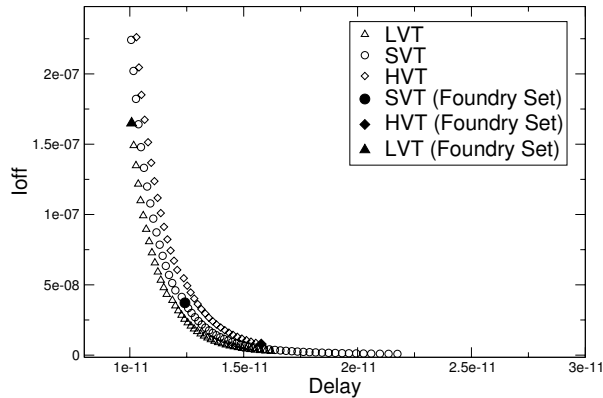


Figure 2: Off-current and delay of the NMOS device in $INVX4$ as V_{TH0} is modified for HVT, SVT, and LVT.

Increasing the number of available V_{th} 's improves the granularity of selection of best V_{th} 's but is time consuming because characterization, a computationally intensive process, must be performed for all available V_{th} 's for each cell in our library. Therefore we strike a middle ground and increase the number of V_{th} 's from three to seven. We create four more V_{th} 's in addition to the foundry-set V_{th} 's such that all V_{th} values are approximately equally spaced. Table 2 gives threshold voltages, I_{off} and delay values for different V_{th} 's

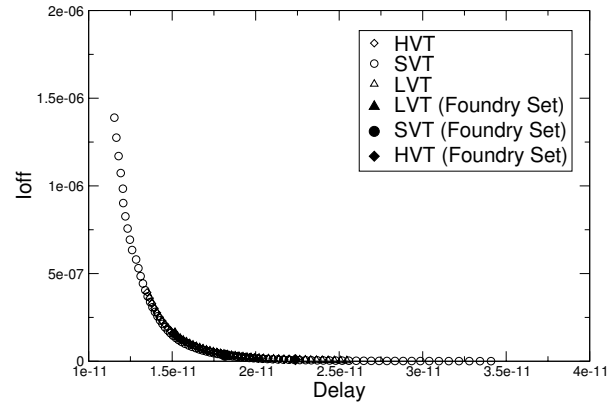


Figure 3: Off-current and delay of the PMOS device in $INVX4$ as V_{TH0} is modified for HVT, SVT, and LVT.

used in our study.

4. Experiments and Results

In this section we describe our experimental setup, results and analysis.

4.1 Experimental Setup

We alter V_{TH0} in TSMC 100nm SPICE device models (as explained in Section 3) to get SPICE models for our seven V_{th} 's. Library characterization is performed using Cadence SignalStorm v04.10 and Synopsys HSPICE on Artisan TSMC standard cell SPICE netlists for 50 combinational cells. We do not create cell variants for different V_{th} 's and gate-biases for the 13 sequential cells in our library to reduce characterization runtime. Therefore, sequential cells are not optimized for leakage in our experiments.

We use three large testcases available in the open domain [1]. Larger SoC chips are formed of several smaller modules like our testcases and each module is optimized separately. So we consider our testcase sizes fairly representative. V_{th} assignment and gate-length biasing tradeoff timing slacks for leakage reduction. Since V_{th} assignment and gate-length biasing are typically performed post-layout, the circuits have already undergone an array of optimizations and have a tight slack distribution. To obtain tight slack distributions for our testcases, we iteratively perform synthesis with Synopsys DesignCompiler v2004.12, each time decreasing the target clock cycle time by a small step, until a tight slack distribution

Table 2: The seven threshold-voltages used in our experiments.

Name	Description	NMOS			PMOS		
		V_{th} value	$I_{off}(nA)$	Delay (ps)	V_{th} value	$I_{off}(nA)$	Delay (ps)
HVT	Foundry-set high- V_{th}	0.402	7.5	14.86	-0.300	9.4	21.65
SVT	Foundry-set standard- V_{th}	0.327	37.2	12.42	-0.235	34.2	18.16
LVT	Foundry set low- V_{th}	0.257	164.2	10.38	-0.155	160.6	14.89
HSVT	V_{th} midway between HVT and SVT	0.362	17.6	13.48	-0.265	18.9	19.62
HHVT	Ultra-high V_{th} (V_{th} of HVT is midway between HHVT and HSVT)	0.437	3.7	16.37	-0.330	5.1	23.69
SLVT	V_{th} midway between SVT and LVT	0.292	78.5	11.31	-0.195	74.8	16.41
LLVT	Ultra-low V_{th} (V_{th} of LVT is midway between SLVT and LLVT)	0.222	338.0	9.66	-0.115	337.5	13.68

is attained. Details about our testcases are summarized in Table 3. As can be seen, synthesis with LLVT (LVT) library reduces the clock cycle time but dramatically increases the leakage in comparison to the LVT(SLVT) library. Histogram in Figure 4(a) shows the path slack distribution for testcase *AES* after synthesis.

Our leakage reduction results are obtained from an industry leakage optimizer that is tuned for leakage reduction using V_{th} assignment and gate-length biasing. It can optimize multi-million gate designs with high quality of results in tractable runtime. In all our experiments, we do not allow the clock cycle time to increase during optimization. The slack distribution of *AES* is shown in Figure 4(b) after dual- V_{th} assignment optimization. While the leakage reduces significantly, the number of critical paths dramatically increases.

4.2 Experiments

We perform the following three analyses:

1. Comparison of leakage reductions from triple- V_{th} assignment, and dual- V_{th} assignment with biasing. I.e., we explore whether dual- V_{th} with biasing can be used as a replacement of triple- V_{th} .
2. Assess the choice of foundry-selected V_{th} 's by changing foundry-set V_{th} 's and comparing leakage reductions.
3. Impact of availability of gate-length biasing on the selection of V_{th} 's.

Leakage reductions from dual- V_{th} with biasing were compared with triple- V_{th} in Table 1. The clock cycle time is constant over all optimization runs and bias values of 4nm, 6nm, 8nm, and 10nm were used. We observe that triple- V_{th} reduces leakage with respect to dual- V_{th} (without biasing) by $\sim 8\%$ on average. However, when biasing is available, dual- V_{th} can yield reductions comparable to triple- V_{th} that involves higher process costs.

We now assess the choice of foundry-selected V_{th} 's when biasing is not available. We perform dual- V_{th} assignment with no increase in clock cycle time. The high V_{th} and low V_{th} are changed and leakage is measured. Table 4 presents the results for our testcases. For *AES* and *DES3* HSVT yields the best reduction while for *OR1200* HHVT is the best. In *OR1200* a smaller percentage of paths is critical (i.e., the slack distribution is relatively loose) and this causes a lot of cells to be assigned the highest V_{th} . HHVT has the highest leakage reduction per cell and consequently it yields the best leakage reduction. Therefore, foundry-set V_{th} 's may

not yield the best leakage reduction and the slack distribution, netlist structure, and leakage-delay tradeoff must be understood prior to V_{th} selection.

It is also clear from Table 4 that reducing the low V_{th} does not yield good leakage reductions. One could expect that using lower low- V_{th} would loosen the design and make it amenable to optimization. For instance, a gate that drives several gates can be assigned lower V_{th} to decrease its delay so that the gained slack can be used to reduce the leakage of the many gates driven by it. However, due to the extremely high leakage cost associated with lower low- V_{th} its usage for reducing leakage is questionable. Therefore, low- V_{th} selection should be governed by timing only and the maximum V_{th} that allows the circuit to meet timing must be used as the low- V_{th} .

To assess the impact of availability of gate-length biasing on V_{th} selection, we perform leakage optimization for three different low V_{th} 's for each of our testcases. The clock cycle time and absolute leakage values differ for the three low V_{th} 's as shown in Table 3. For each testcase and low V_{th} combination, we change the set of available gate-length biases and run experiments to identify the optimum (i.e., the highest leakage reduction yielding) high V_{th} . Table 5 shows the leakage reduction achieved for testcase *AES* synthesized under three low V_{th} for different available gate-bias values and high V_{th} .

Following observations may be made from the results:

- Leakage reduction increases as more biases are allowed for optimization. However, the benefit progressively diminishes as the number of biases becomes large.
- As the low V_{th} is increased, the optimum higher V_{th} also increases. This indicates that V_{th} 's should neither be spaced wide apart nor placed too close so that the leakage-delay tradeoff is effectively covered.
- Availability of gate-length biasing has small impact on optimum high V_{th} . For LVT as the lower V_{th} , we observe that the optimum high V_{th} shifts from HSVT to SVT when biasing becomes available. Similar trends were observed for *DES3*.

Table 6 shows the optimum high V_{th} for all three of our testcases.

5. Conclusions

In this paper we have shown that simultaneous use of gate-length biasing and threshold-voltage assignment to be effective for runtime leakage reduction. Our experiments with

Table 3: Testcases used in our experiments and their details.

Testcase	Source	#Cells	LLVT		LVT		SLVT	
			Delay (ns)	Leakage (mW)	Delay (ns)	Leakage (mW)	Delay (ns)	Leakage (mW)
AES	opencores.org	22,000	1.134	9.46	1.214	4.61	1.294	2.24
OR1200	opencores.org	37,000	2.860	24.01	2.960	13.08	3.110	7.69
DES	opencores.org	86,000	1.081	36.31	1.106	18.08	1.160	9.08

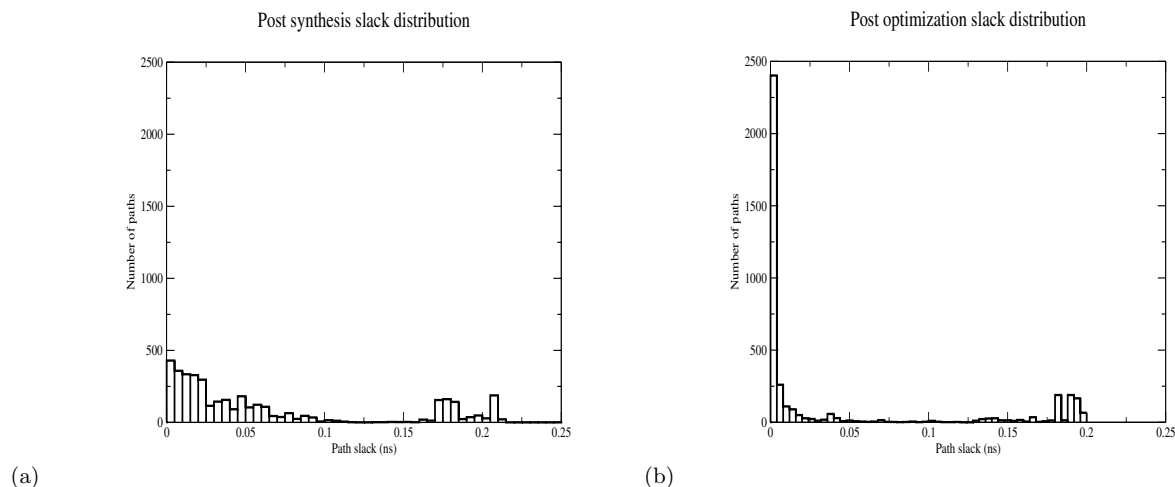


Figure 4: Distribution of timing slacks for top 4000 paths of AES (a) after synthesis and (b) after optimization

real-world testcases and an industry leakage optimizer indicate that dual- V_{th} assignment when used with gate-length produces leakage reductions that are comparable to triple- V_{th} assignment (without biasing). Therefore, gate-length biasing may be used with dual- V_{th} assignment as a substitute for the higher-cost triple- V_{th} process.

Our experiments with perturbation of foundry V_{th} 's show that foundry-selected V_{th} 's may not yield the best leakage reductions. The slack distribution, netlist structure, and leakage-delay tradeoff due to V_{th} assignment must be analyzed to judiciously choose (for fabless design houses) or customize (for IDM's) the V_{th} 's for improved leakage reduction. We also note that reducing low V_{th} deteriorates leakage reduction despite loosening the slack distribution of the circuit and making it more amenable to leakage optimization. Hence, the highest V_{th} for which the circuit meets timing must be used as the low V_{th} .

Availability of gate-length biasing as an optimization knob has a small impact on optimal high- V_{th} (i.e., high- V_{th} that minimizes leakage). For well-constrained designs with tight slack distribution, we found the optimal high V_{th} to lower in presence of gate-length biasing.

6. References

- [1] Opencores.org. <http://www.opencores.org/projects/>.
- [2] A. Agarwal, C. H. Kim, S. Mukhopadhyay and K. Roy. Leakage in Nano-Scale Technologies: Mechanisms, Impact and Design Considerations. In *Proc. ACM/IEEE Design Automation Conference*, pages 6–11, 2004.
- [3] S. Borkar, T. Karnik, S. Narendra, J. Tschanz, A. Keshavarzi and V. De. Parameter Variations and Impact on Circuits and Microarchitecture. In *Proc. ACM/IEEE Design Automation Conference*, pages 338–342, 2003.
- [4] F. Boeuf et al. A conventional 45nm CMOS Node Low-Cost Platform for General Purpose and Low Power applications. In *IEEE International Electron Devices Meeting*, pages 425–428, 2004.
- [5] P. Gupta, A. B. Kahng, P. Sharma and D. Sylvester. Selective Gate-Length Biasing for Cost-Effective Runtime Leakage Control. In *Proc. ACM/IEEE Design Automation Conference*, pages 327–330, 2004.
- [6] P. Gupta, A. B. Kahng, P. Sharma and D. Sylvester. Gate-Length Biasing for Runtime Leakage Control. PP(99):–, 2005.
- [7] P. Royannez et al. 90nm low leakage SoC design techniques for wireless applications. In *IEEE International Solid-State Circuits Conference*, pages 138–140, 2005.
- [8] S. Nakai et al. A 65 nm CMOS Technology with a High-Performance and Low-Leakage Transistor, a $0.55\mu m^2$ 6T-SRAM Cell and Robust Hybrid-ULK/Cu Interconnects for Mobile Multimedia Applications. In *IEEE International Electron Devices Meeting*, pages 11.3.1–11.3.4, 2003.
- [9] Y.-F. Tsai, N. Vijaykrishnan, Y. Xie and M. J. Irwin. Influence of leakage reduction techniques on delay/leakage uncertainty. In *Proc. of VLSI Design*, pages 374–379, Jan 2005.
- [10] Y. Nakahara et al. A robust 65-nm node cmos technology for wide-range vdd operation. In *IEEE International Electron Devices Meeting*, pages 11.2.1–11.2.4, 2003.
- [11] Z. Luo et al. High Performance and Low Power Transistors Integrated in 65nm Bulk CMOS Technology. In *IEEE International Electron Devices Meeting*, pages 661–664, 2004.

Table 4: Post-optimization leakage (in mW) for two low V_{th} 's and different high V_{th} 's.

Circuit	Low V_{th}	High V_{th}			
		HHVT	HVT	HSVT	SVT
AES	LLVT	3.100	2.689	2.182	1.812
	LVT	1.785	1.658	1.529	1.572
OR1200	LLVT	4.149	4.264	4.649	5.557
	LVT	3.699	3.818	4.219	5.254
DES3	LLVT	8.732	7.224	5.906	5.902
	LVT	5.600	4.986	4.614	5.4112

Table 5: Leakage reduction for different high- V_{th} with different maximum gate-length biasings for AES. Best leakage reduction is shown in bold.

Lower V_{th}	Max. bias	Higher V_{th}			
		HHVT	HVT	HSVT	SVT
LLVT	no biasing	47.94	51.37	56.78	61.58
	4nm	54.07	57.24	62.09	64.88
	6nm	54.97	57.90	62.84	65.50
	8nm	55.35	58.21	63.10	65.79
	10nm	55.87	58.49	63.5	66.06
LVT	no biasing	61.27	64.03	66.83	65.89
	4nm	64.89	66.68	68.23	68.35
	6nm	65.66	67.31	68.85	69.19
	8nm	65.97	67.70	69.12	69.72
	10nm	66.21	67.90	69.46	70.15
SLVT	no biasing	64.42	65.33	63.23	49.08
	4nm	67.10	67.93	66.43	54.10
	6nm	67.74	68.72	67.25	56.10
	8nm	68.02	69.10	67.87	57.55
	10nm	68.21	69.41	68.34	58.63

Table 6: Best high V_{th} for three low V_{th} 's and maximum bias of 4nm, 6nm, 8nm, and 10nm

Circuit	Low V_{th}	No biasing		4nm		4, 6nm		4, 6, 8nm		4, 6, 8, 10nm	
		High V_{th}	Reduction (%)	High V_{th}	Reduction (%)	High V_{th}	Reduction (%)	High V_{th}	Reduction (%)	High V_{th}	Reduction (%)
AES	LLVT	SVT	61.58	SVT	64.88	SVT	65.50	SVT	65.79	SVT	66.03
	LVT	HSVT	66.83	SVT	68.36	SVT	69.20	SVT	69.72	SVT	70.15
	SLVT	HVT	65.33	HVT	67.93	HVT	68.72	HVT	69.10	HVT	69.41
OR1200	LLVT	HVT	82.24	HVT	83.26	HVT	83.62	HVT	83.90	HVT	83.90
	LVT	HHVT	71.74	HHVT	72.74	HHVT	73.02	HHVT	73.09	HHVT	73.52
	SLVT	HHVT	55.49	HHVT	57.98	HHVT	58.87	HHVT	59.56	HHVT	60.14
DES3	LLVT	SHVT	80.12	SVT	82.97	SVT	83.36	SVT	83.67	SVT	83.90
	LVT	SHVT	74.48	SHVT	76.65	SHVT	77.03	SHVT	77.28	SHVT	77.49
	SLVT	HHVT	66.14	HVT	69.57	HVT	70.10	HVT	70.42	HVT	70.61