EFFICIENT SUBTHRESHOLD LEAKAGE CURRENT OPTIMIZATION

Leakage Current Optimization and Layout Migration for 90- and 65-nm ASIC Libraries

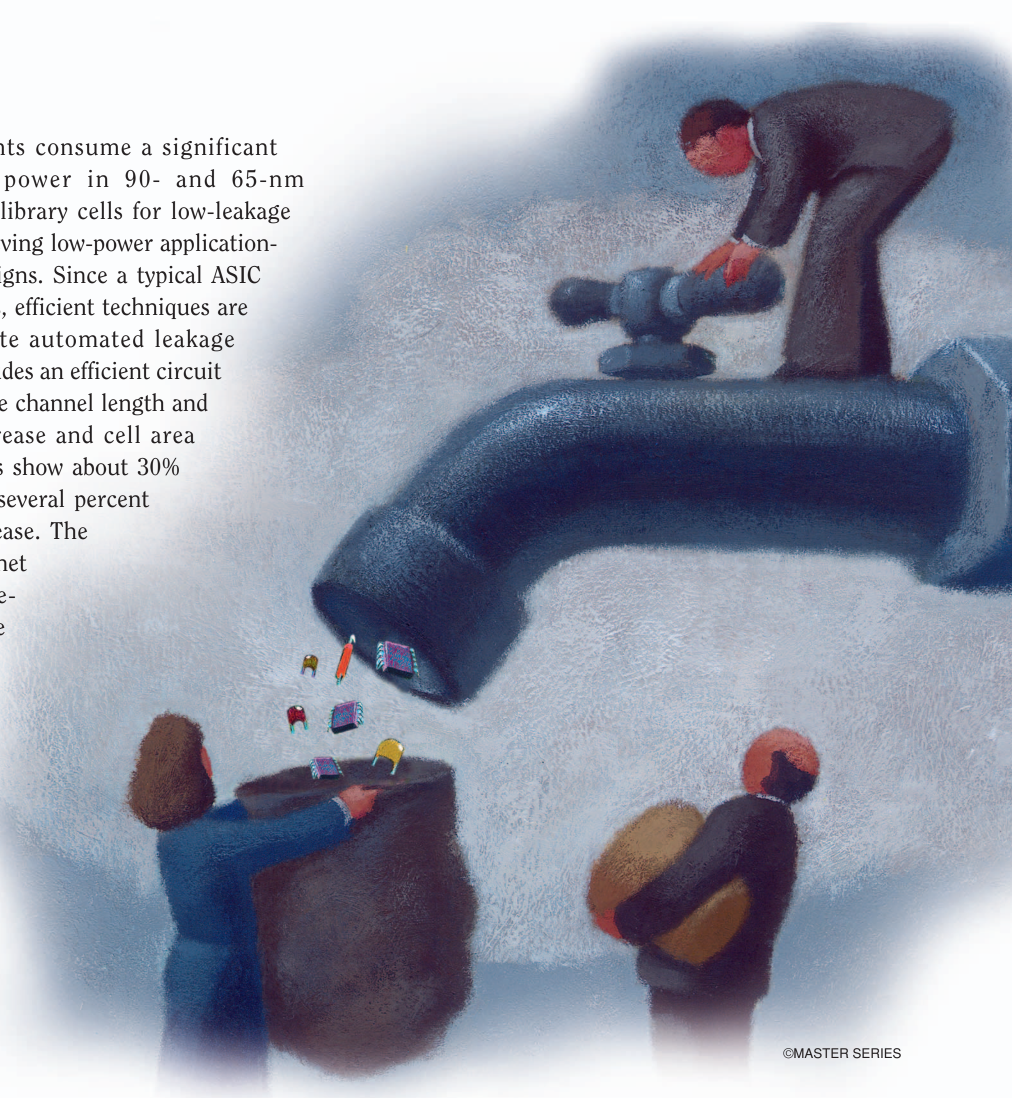

Subthreshold leakage currents consume a significant fraction of total circuit power in 90- and 65-nm technologies. Generation of library cells for low-leakage current is important for achieving low-power application-specific integrated circuit (ASIC) designs. Since a typical ASIC library may contain thousands of cells, efficient techniques are required. In this article, a complete automated leakage optimization flow is presented. It includes an efficient circuit optimization engine to optimize device channel length and width while keeping cell delay increase and cell area change minimal. Optimization results show about 30% leakage current reduction with only several percent active area and dynamic power increase. The optimization flow starts with SPICE net lists, including parasitic resistance-capacitance (RC) extracted from the existing library. After leakage optimization, revised cell layouts are generated and characterized based on the optimized net lists. In addition, investigations indicate that the optimization process has little impact on cell noise margins and that new cell layout variations

Xiaoning Qi, Sam C. Lo, Alex Gyure, Yansheng Luo, Mahmoud Shahram, Kishore Singhal, and Don B. MacMillen

 8755-3996/06/$20.00 ©2006 IEEE

from placement, routing, and compaction have little effect on the results. This efficient automated layout to layout cell leakage optimization flow is most suitable for leakage reduction and library migration for 90- and 65-nm ASIC designs and beyond.

Due to aggressive process technology scaling, it is now possible to integrate several billions of transistors onto a single die. Smaller feature sizes make transistor density higher, and reduction in oxide thickness enables faster transistors. To achieve low active power in designs, circuit supply voltage and corresponding transistor threshold voltage must be scaled down. Unfortunately, the reduction in threshold voltage has also resulted in an increase in channel leakage currents that can increase total power consumption substantially. Physical gate lengths will continue to scale down to 8 nm by 2017, as predicted by the 2005 International Technology Roadmap for Semiconductors (ITRS'05) [1], notwithstanding technical challenges such as leakage currents and process variations. Figure 1 plots subthreshold leakage current for different technologies, which shows that the subthreshold leakage (I_{off}) increases exponentially as process technologies scale, including the projected leakage at the 45-nm technology node. To ensure that designs can benefit from higher density and performance in nanometer technologies, leakage optimization and design-for-manufacturing (DFM) have become essential elements of nanometer library design.

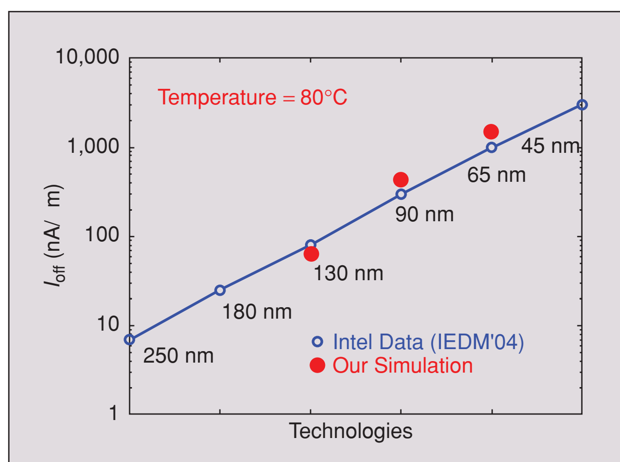

1. Sub-threshold leakage currents increase exponentially with technologies [2]. Our simulations were based on third-party foundry technologies.

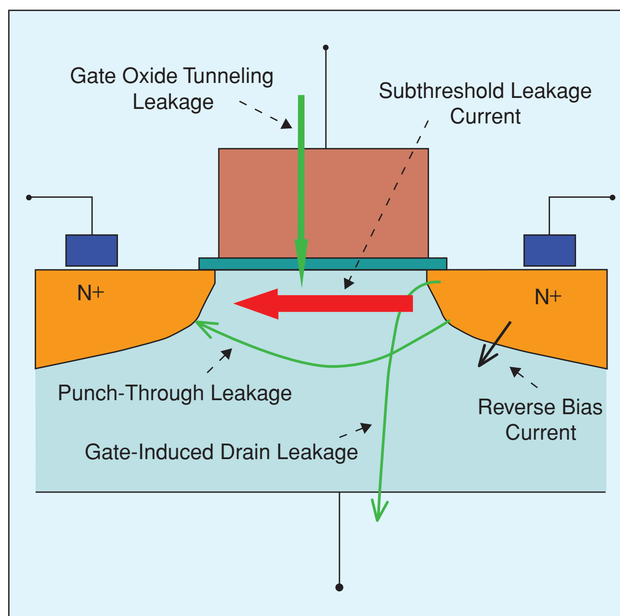

2. Leakage currents include source drain sub-threshold leakage current and gate leakage current. In 90- and 65-nm technologies, subthreshold leakage current dominates.

Leakage mechanisms include subthreshold leakage, gate oxide tunneling leakage, junction leakage, hot-carrier injection leakage, gate-induced drain leakage, and punch-through leakage currents [2]. In Figure 2, some of these leakage components are illustrated. At the 180-nm technology node, subthreshold leakage currents emerge as an issue. At the 90-nm technology node, subthreshold leakage and gate leakage become a problem. At even smaller gate lengths, all leakage components will be evident. However, studies find that subthreshold leakage current is the dominant mechanism in 90- and 65-nm technologies. Subthreshold current, also known as weak inversion conduction current, between source and drain in a MOS transistor occurs when the gate voltage is below the transistor threshold voltage (V_{th}). Although threshold voltage is decreasing with technology scaling, the subthreshold slope, a measure of transistor switch-off characteristics or transitions between weak and strong inversion regions are mainly controlled by material properties which have not improved drastically with scaling.

Circuit simulations reveal that in the 65-nm technology, subthreshold leakage power can be as high as 60% of total chip power as shown in Figure 3, where several process corner, supply voltage, and temperature (PVT) combinations were simulated for an 800-MHz high-end ASIC. In the simulation, assumptions were made that 95% of gates were quiet, and 5% of gates were switching [3]. It is evident that temperature, as well as process corners and supply voltage, has significant impact on leakage currents. As an increasing number of designs move into the 90- and 65-nm technologies, it is important to reduce subthreshold leakage current in creating ASIC cell libraries or to migrate an existing library to a low-leakage library to minimize total chip power.

Several techniques have been proposed to reduce subthreshold leakage currents [2], [4]. These include device and process optimization [5] and circuit design techniques to reduce and control subthreshold leakage [6], [7]. More specifically, the techniques involve supply voltage optimization [8], reverse body biasing [9], and multiple-threshold voltage assignment and state assignment [10], [11]. Moreover, Yuan proposed a gate replacement method [12] that required leakage-optimized gate libraries. Lichtensteiger [13] discussed leakage modeling to estimate leakage currents in ASIC libraries, but leakage reduction and optimization techniques were not proposed. Sirisantana [14] proposed the use of multiple large

channel lengths to reduce leakage current with some delay penalty. Techniques for optimizing a library systematically and efficiently with multiple gate lengths are highly desirable. In addition, the layout realization and possible impact of parasitics on the final optimization were yet to be investigated. Another popular method is to use dual-threshold-voltage devices in designs to reduce leakage currents. Both dual-threshold-voltage and multiple-channel- length techniques can reduce leakage currents by replacing cells in noncritical paths with high V_{th} devices or multiple large channel length devices. At the cost of delay degradation and/or possible area increase, leakage currents of the circuits can be reduced. Compared to the multiple-channel-length method, the dual-threshold-voltage method leads to greater leakage current reduction at the cost of larger delay degradation. It also requires extra wafer processing steps and is thus more costly. The multiple-channel-length method, on the other hand, has freedom to tradeoff leakage with delay and area, and thus has a finer tuning capability. Also, a larger channel length can potentially reduce process variability, and ultimately increase yield. However, a new cell layout is needed after cell optimization.

In this article, an efficient subthreshold leakage current reduction and optimization method using multiple-channel-lengths is presented, and results are given for 90- and 65-nm ASIC libraries. The flow starts with extracted SPICE net lists that include RC parasitics from an existing library. While using multiple channel lengths to reduce subthreshold leakage current, cell delay changes are kept small, and cell active area (transistor length × transistor width) increase is limited to 10%. The SPICE simulation-based optimization engine can quickly characterize circuit performance and evaluate well-defined objective functions (e.g., leakage, delay) for each candidate. It generates new sets of gate lengths and widths for all transistors in a cell which satisfies leakage, timing, and area constraints. With newly optimized transistor net lists, a cell generation tool can lay out or migrate an exiting layout into a new cell library with significantly reduced subthreshold leakage currents and total power. With effective compaction, the 10% active area increase only results in fractional percentage increase in total cell area. Finally, the library is recharacterized to generate new timing models. In the rest of the article, a subthreshold leakage reduction and optimization methodology is first presented, followed by discussions of the constraints in the optimization and the impact on noise margins and layout generation. Second, the complete working flow and optimization results are described. Conclusions are drawn in the final section.

METHODOLOGIES

There are several ways to reduce subthreshold leakage current as introduced in the first section. The multiple-channel-length method is effective because leakage current is a strong function of channel length for short channel devices. At the same time, it is necessary to ensure that other circuit metrics, such as delay and noise margins, are not severely degraded. In addition, the layout impact needs to be investigated so that the leakage reduction results from circuit optimization will not lead to excessive area penalty and topology change. In this article, three optimization schemes are proposed for the multiple-channel-length method. In the first scheme, the transistor length, width and cell height can all be changed to

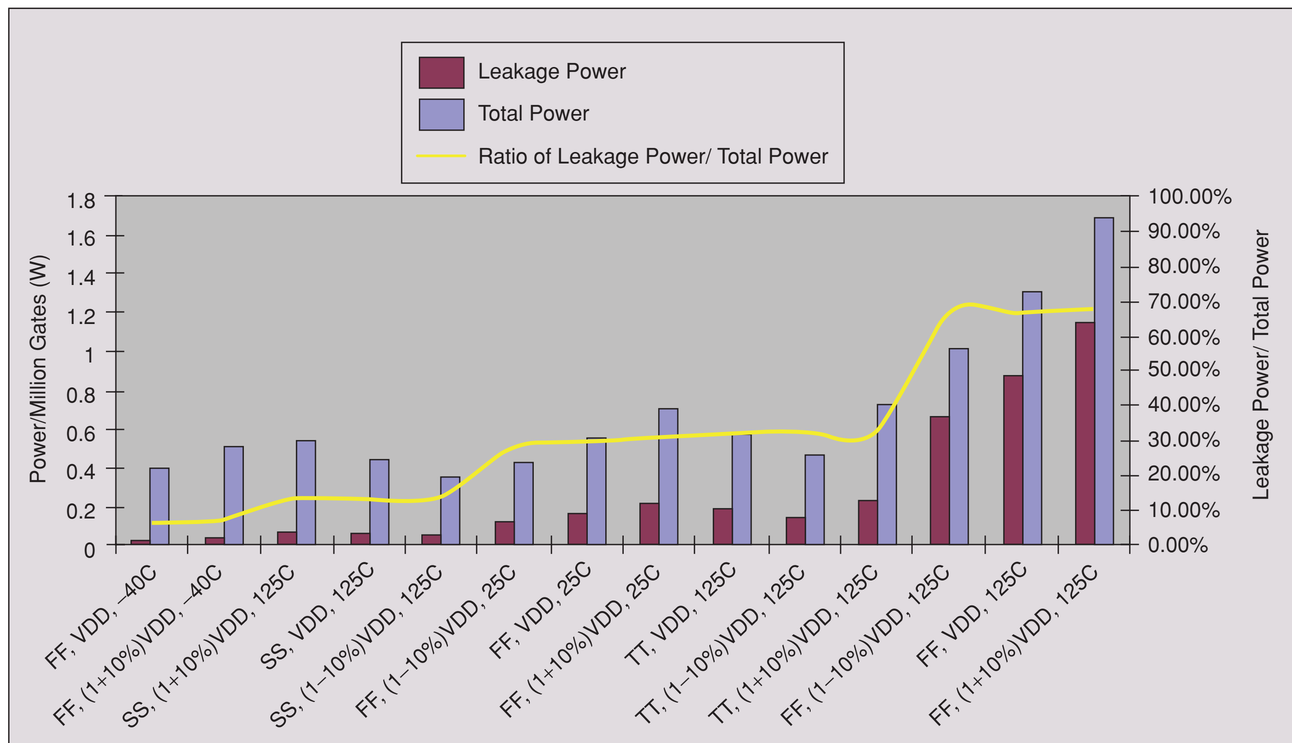

3. Leakage power, total power, and their ratios were analyzed in a 65-nm ASIC. There are 27 PVT combinations, and the highest 14 ratios/cases are plotted. Temperature is a key factor in leakage power.

reduce leakage current while maintaining the circuit delay. This is suitable for generating a new cell library where there is flexibility in cell height. The second scheme is to change transistor length and width, but not the cell height, which leads to moderate delay degradation and area increase. The third scheme is to increase transistor length, but not transistor width or the cell height. This scheme leads to minimal cell area increase with relatively larger delay degradation. Both the second and third schemes are suitable for migrating an exiting library to a new low-leakage library. While the transistor length and width optimization techniques apply to all three schemes, this article will focus on the second and third schemes in terms of layout generation.

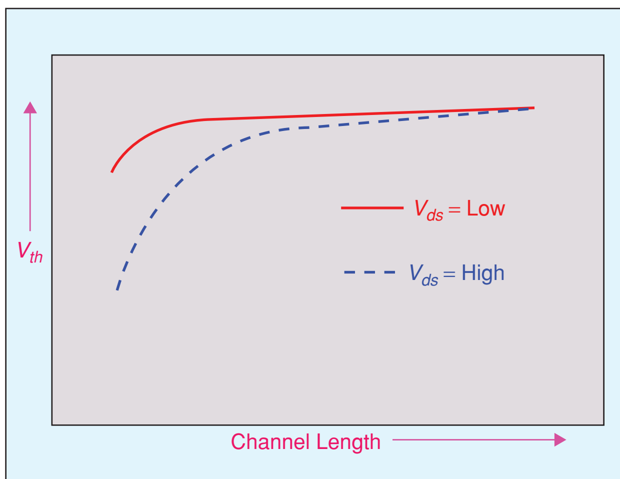

4. Threshold voltage increases dramatically when channel lengths increase from the smallest feature sizes. The increase in threshold voltage helps reduce subthreshold leakage current [13].

Channel Length and Leakage Currents

The subthreshold current between the drain and source occurs in a MOS transistor when the gate voltage is below V_{th} [15]. There are several factors which impact threshold voltage and subthreshold current. They include drain-induced barrier lowing (DIBL) effect, body effect, channel width, short channel effect and temperature effects. Figure 4 qualitatively shows that the channel length decrease results in a sharp roll-off for transistor threshold voltage. This is especially evident for transistors in technologies below 250 nm where HALO techniques must be used to suppress the short channel effect [2]. Mathematically, subthreshold leakage current can be modeled as [16], [17]

$$I_{\text{off}} = \mu_{\text{eff}} C_{ox} \frac{W}{L_{\text{eff}}} \left(\frac{kT}{q}\right)^2 \exp(1.8) \exp\left(\frac{-V_{th} + \eta V_{DD}}{n\frac{kT}{q}}\right), \quad (1)$$

where μ_{eff} is the temperature dependent mobility, C_{ox} is gate-oxide capacitance, W and L_{eff} are transistor width and effective channel length, T is the temperature, k is the Boltzmann constant, V_{th} is the threshold voltage, V_{DD} is the supply voltage, η is the DIBL coefficient, q is the electron charge, and n is the transistor subthreshold swing coefficient. It can be seen that a channel length increase will not only directly reduce $I_{\text{off},}$ but also increase V_{th}, which further reduces I_{off}. Therefore, it is a key parameter to be optimized for subthreshold leakage current. In Figure 5, subthreshold leakage currents in the NMOS and PMOS transistors of an inverter are plotted against channel length based on a 65-nm technology, where the minimum channel length is 60 nm. The leakage currents in both transistors are substantially reduced as the lengths are increased. The main challenge is how to optimize transistors in a cell automatically while minimizing the change in other metrics (e.g., delays, noise margins, and layout area). A complete automated and efficient flow of the optimization and layout generation is critical to leakage reduction and library migration for hundreds and even thousands of cells found in a typical ASIC library.

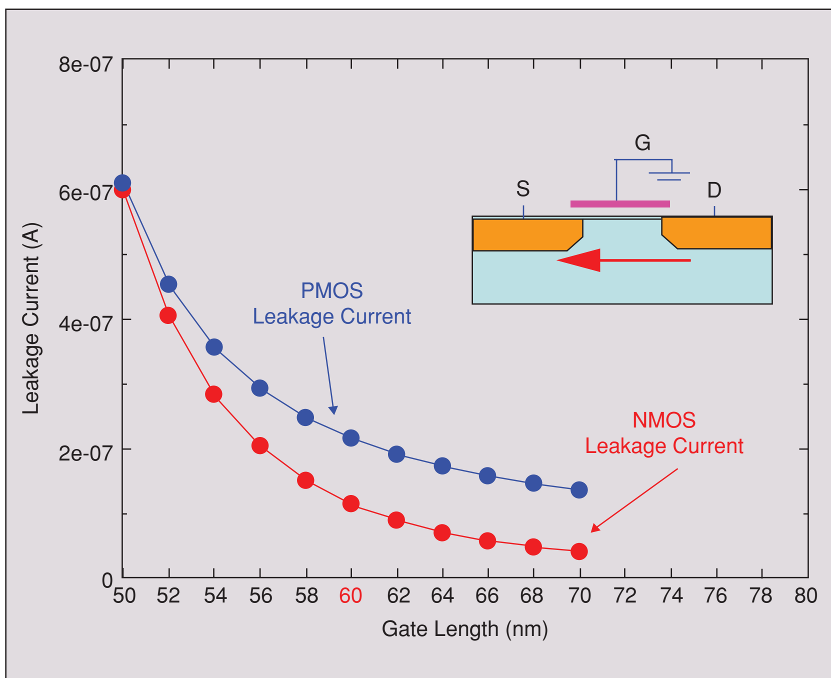

5. Subthreshold leakage currents from NMOS/PMOS transistors in an inverter in a 65-nm technology. The gate delays are kept constant in the simulation.

Leakage Optimization via Transistor Resizing

In the migration flow, the optimization begins with SPICE net lists from an existing cell library, where the initial layout and parasitic RC are extracted. The channel length of every transistor can be considered an independent variable, or they can be grouped together to speed up optimization. For example, all gate lengths of NMOS transistors can use one channel length *Ln,* while all channel lengths of

PMOS transistors are Lp. The goal is to optimize channel lengths so that leakage currents in various states are reduced significantly. In addition, transistor widths must generally be increased to maintain the input to output cell delays. This is usually required for library migration. If some delay degradation is permitted, transistor widths would increase only minimally or even remain unchanged. To tradeoff performance and area, constraints on the active area (channel length×channel width) increase are set to be within a user-specified tolerance, such as 10%. Due to multiple input states, there are several different leakage measurements for NMOS and PMOS devices. For example, a two-input NAND gate has four input combinations and thus four leakage measurements (one dominated by PMOS and three dominated by NMOS). In the optimization, each of the leakage currents can be a goal, or the average of the four leakage currents can be a goal. Likewise, two inputs can have four cases to trigger output to transit from low to high or high to low. All these delays should be in the delay constraints during the optimization. The mathematical formulation of the optimization problem is:

$$
\begin{aligned}
&\text{Minimize} && f_i(x) = I_{\text{off}}(i) \\
&\text{Subject to} && \left|\text{delay}(j) - D(j)\right| \le \delta \\
& && \sum_s L_s W_s \le (1+\alpha) \sum_s A_s,
\end{aligned}
$$

where $x = \{L_1, L_2, \ldots, W_1, W_2, \ldots\}$ are the gate geometries, i refers to an input state, j refers to an output transition, $D(j)$ is the original delay of the jth transition, and δ is a slack (ideally, $\delta = 0$). L_s and W_s are length and width for transistor s, and A_s is the original active area for transistor s. α is the allowed relative area increase.

The optimization engine runs SPICE simulation to compute the objective functions and constraints, generates circuit candidates, and evaluates the objective functions again. The advantage of circuit simulation-based optimization against simple formulae or model-based optimization is superior accuracy. In addition, any temperature dependence of leakage is easily incorporated. To speed up the optimization, parallel computing schemes are adopted with multiple SPICE runs on multiple machines. Figure 6 shows optimization results for a NAND gate in a 65-nm technology. Large leakage reduction for the four input states are shown, where ~30% leakage reduction was achieved at the cost of 4% active area increase. Since active area is usually a small fraction of the total cell area, and usually there is white (free) space in the original cell, the resulting total area increase is minimal. Details of the cell area optimization will be presented in a later section.

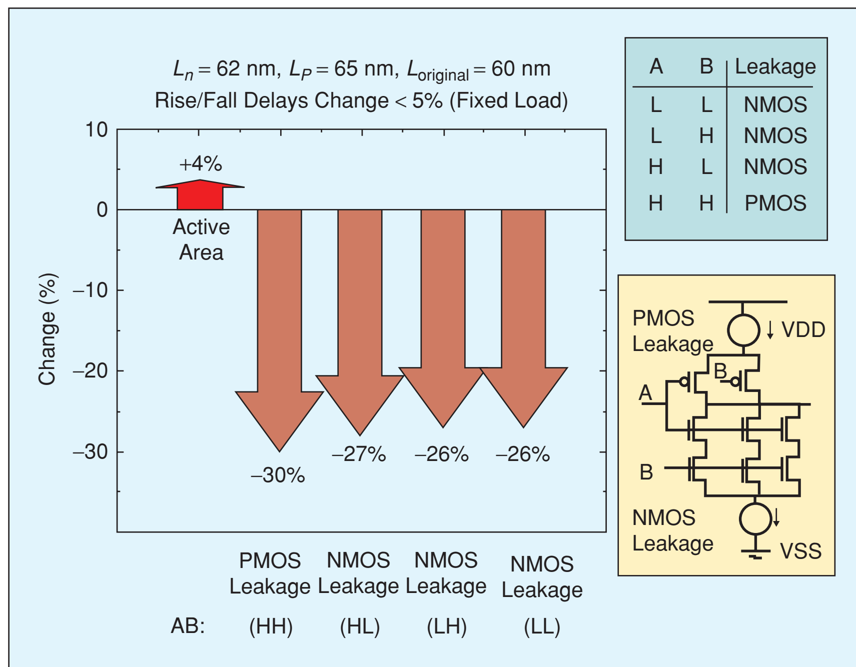

6. A NAND gate subthreshold leakage reduction result in a 65-nm technology. The optimization is performed under the conditions: 125C, 1.1 × VDD and typical process corner.

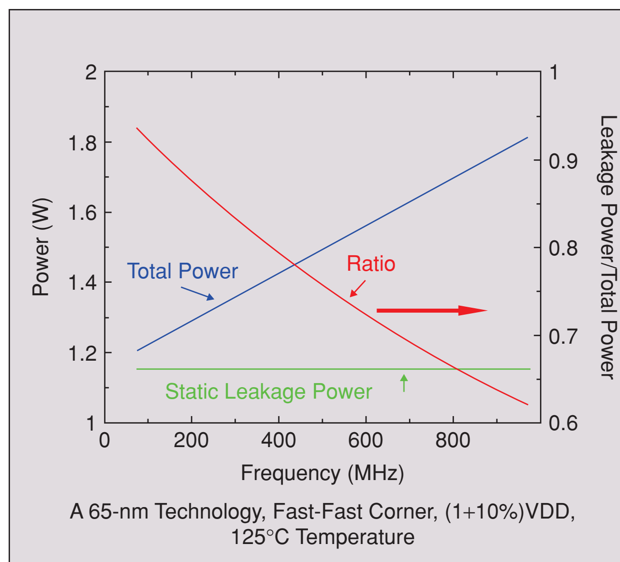

7. Leakage power is a significant fraction of total power at low frequencies. The data is for 1 million gates with 95% gates quiet and 5% gates switching.

Dynamic Power and Noise Margins

A consequence of this method is that gate input capacitance and dynamic power increase as transistor geometries are increased to reduce leakage power. Gate capacitance consists of area capacitance and fringing capacitance (a function of width), which in nanometer technologies, are roughly equal. In fact, gate capacitance tends to be dominated by fringing components as length decreases. More specifically, a possible 10% gate area increase allowed in optimization can result in 5–8% gate capacitance increase

in 90- and 65-nm technologies. Since dynamic power is a near linear function of gate/interconnect capacitance, the upper limit of the dynamic power increase is bounded by 8%. Considering a 25% leakage power reduction and ~65% of leakage power to total power ratio for an 800 MHz ASIC in a 65-nm technology, the net total power reduction is about 14%. In addition, if an ASIC is running at lower frequencies, the leakage power to total power ratio is even higher as shown in Figure 7, and the net total power reduction is more significant. Considering an ASIC with operating frequency of 500 MHz, a 25% leakage power reduction results in 17% total power reduction.

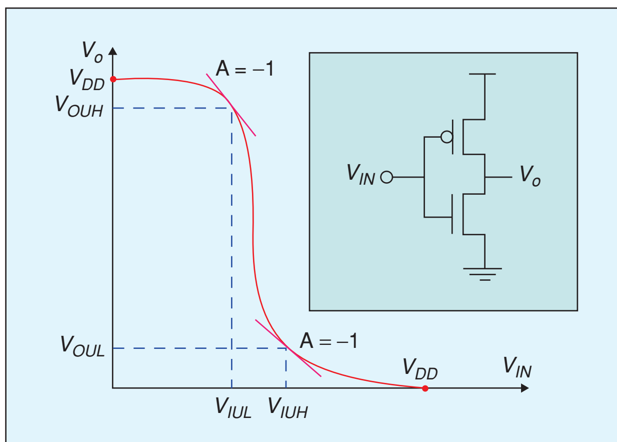

8. Voltage transfer function and CMOS unity gain margins for an inverter. The definition of V_{IUL} V_{IUH}, V_{OUH}, and V_{OUL}, are defined by the input and output voltages at the unity gain point ($A = -1$).

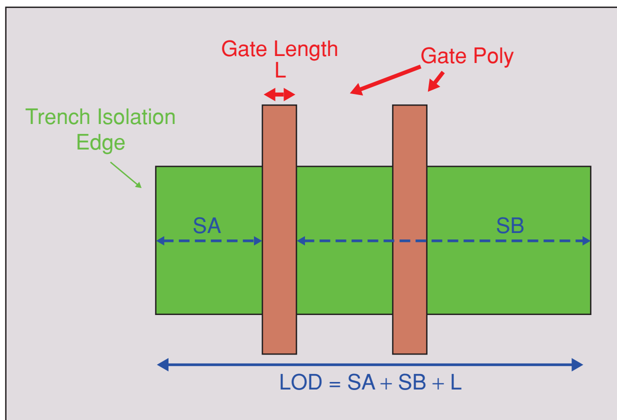

9. A layout of two-stacked MOSFET and the definitions of SA and SB.

While leakage currents are reduced by using multiple gate lengths and cell delay changes are kept minimal, it is necessary to investigate other important cell characteristics to see whether they have been significantly affected. One important characteristic of library cells is noise margin, and any leakage optimization should not degrade the cell's noise margin. Figure 8 shows the input and output transfer curves for a CMOS inverter, where unit gain points are defined for the output gain $A \equiv {}^{V_o}/_{V_{in}} = -1$. The unit gain noise margins are defined by

$$UGM_H \equiv V_{OUH} - V_{IUH} \quad (2)$$

$$UGM_L \equiv V_{IUL} - V_{OUL}, \quad (3)$$

where V_{OUH}, V_{IUH}, V_{IUL}, and V_{OUL} are defined by the input and output voltages at the unit gain points. Derivation of complete analytical formulae of noise margins for all gates is nontrivial, but SPICE simulations of the four quantities (V_{OUH}, V_{IUH}, V_{IUL}, and V_{OUL}) showed little change (<3%) before and after leakage optimization (see Table 1), which means that the unit gain noise margins will not suffer significant degradation.

Layout Generation and Migration

Since the optimized net lists will be used to generate new cell layouts, it is important to make sure that the new layouts will not alter the optimized leakage results. Figure 9 shows a typical top view of a layout for two-stacked transistors. The *SA*, *SB*, device area, and perimeters are the parameters that may be affected by the increased gate lengths. To investigate the impact from those layout parameters on leakage currents and other performance metrics, circuit simulations using the Monte Carlo method were performed for an inverter and a NAND gate. It is assumed that there is ±10% change of *SA*, *SB,* and device perimeters at 3-σ. Both NMOS/PMOS leakage current and gate delays were measured. The means and standard deviations of inverter leakages and delays are listed in Table 2, which shows that realistic variations in these parameters have little impact on leakage currents and delays. The same is true for the NAND gate.

To investigate the total cell area increase due to the optimization, a 10% channel length increase is assumed for a 2x NAND2 cell. Figure 10 shows the channel length increase and its impact on subthreshold leakage reduction (34%), total power reduction (8%), delay degradation (9%), and cell area

Table 1. Circuit simulations of noise margins for cells in a high-performance 90-nm technology (TT process corner, temperature = 125°C, VDD = 1.2 V, FO4 loads).

Cell Names	Original				After Optimization			
	V_{IUL}	V_{OUL}	V_{IUH}	V_{OUH}	V_{IUL}(%)	V_{OUL}(%)	V_{IUH}(%)	V_{OUH}(%)
2x Inverter	0.372	0.119	0.688	1.096	+0.5	−1.7	−0.6	+0.2
2x Aoi22	0.346	0.654	0.138	1.086	+2.6	−0.9	−3.0	+0.7
2x Nand2	0.428	0.736	0.135	1.071	+0.5	−0.5	−3.0	+0.5

change (0%). The new layout generated from Synopsys library layout tool CADABRA is shown side by side with the original cell layout in Figure 11. It is worthwhile noting that CADABRA has the ability to maintain the important features of the source layout, and it can also increase the number of vias for DFM purposes. The layout migration takes the original layout as a source layout and maintains the same placement and routing topology as far as possible except for increasing channel lengths. Recompaction is then applied to generate a new optimized layout for the low-leakage library. Since standard cell widths are in increments of layout columns/grids, quite often the layout optimized cell does not grow in size unless there is very little white space in the original layout. Results for more complex cells will be given in the next section.

TOOLS AND RESULTS

Based on the methodology described in the previous section, a leakage optimized library characterization flow was developed. This automatic flow consists of several commercial tools: 1) a SPICE simulator, 2) a standard cell library characterization tool, 3) a circuit optimizer, 4) a library layout generation and migration tool, and 5) a parasitic extraction tool. This flow has been applied to an established commercial library in a 90-nm technology. Results show that the flow works efficiently and achieves significant subthreshold leakage power reduction.

Optimization Flow

Three sets of data are required for the flow: 1) a subset of the Liberty timing library, 2) the GDSII cell layouts, and 3)

Table 2. Monte Carlo simulation of layout impact on an inverter: Leakage currents and delay in a 90-nm technology.

Leakage and Delay	Mean	σ
NMOS Leakage	243.1 nA	1.7 nA
PMOS Leakage	524.2 nA	454.9 pA
Output-Rise Delay	26.2 ps	47.3 fs
Output-Fall Delay	30.4 ps	46.3 fs

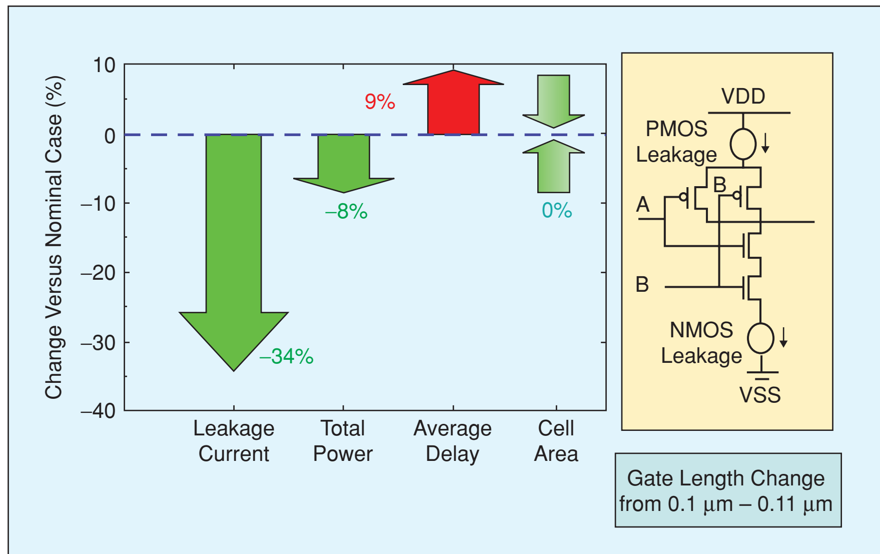

10. A NAND 2x2 cell. Device channel lengths are increased by 10% from 100 nm in a 90-nm technology. Simulation is performed under TT process corner, VDD = 1.0 v, and temperature of 125°C.

11. Migration of NAND 2x2 layout after subthreshold leakage current optimization. There is no cell area increase in the new layout while channel lengths are increased by 10%.

postlayout standard cell SPICE net lists. The first step passes the input files to a standard cell library characterization tool. The tool extracts information from the timing library such as cell name, input and output pins, timing arc sensitization, loading capacitances, and input slews. The characterization tool generates SPICE decks with appropriate arc sensitization and different combinations of input states for each standard cell. Capacitance load and input slew are selected based on the most commonly expected load and slope values. Measurements of leakage current, delay, and active area are generated, which act as constraints for the second step of the flow. The second step is circuit optimization. The circuit optimizer uses the SPICE input decks and performs circuit simulation and optimization. The optimizer changes the channel length and width for every transistor in the standard cell. The final step is layout and library generation and characterization. During optimization, constraints can be set, for example:

- total active area increase is less than 10%
- channel length increase is less than 10%
- channel width is within the range of +/− 20%
- delay change is within +/− 5%
- target average leakage reduction of 25%.

To make this flow more practical, additional settings are required:

- optimization step size for widths and lengths should be on minimum layout grid increment
- grouping multiple widths or lengths together based on circuit topology and layout.

Ideally, the channel lengths of every transistor would be an optimization variable. For large standard cells such as AOI, this may result in a large number of optimization variables. A more efficient method is to group several widths or lengths into macrovariables. The reduction in the number of free variables yields faster optimization run times and reduces layout complexity. The circuit optimizer performs simulations using a commercial HSPICE simulator. Based on the measurement of area, delay, and leakage

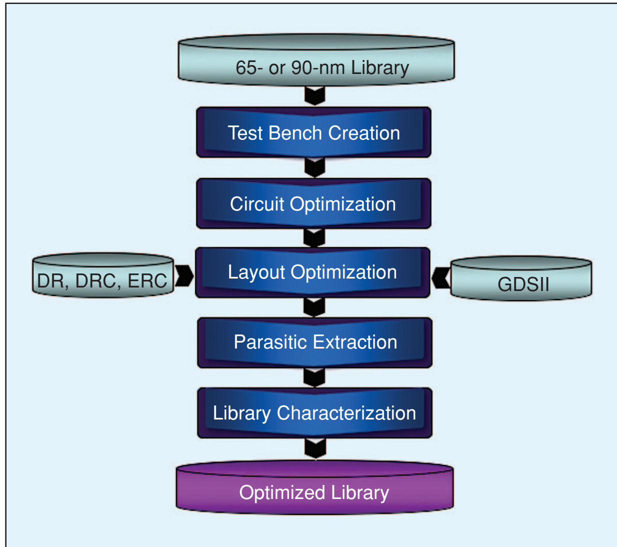

12. A complete flow of ASIC library optimization for leakage currents.

Table 3. Circuit optimization results based on a 90-nm CMOS high performance library.

Cell Names	Original		After Optimization	
	Active Area (fm^2)	I_{off} (nA)	ΔActive Area (%)	ΔI_{off} (%)
2x Inverter	140	337	10.29	−25.92
4x Buffer	390	1124	7.6	−27.6
1x Delay	265	457	6.04	−25.9
1x Mux2	427	650	4.33	−25.39
2x Xnor2	567	1267	6.24	−32.89
2x Aoi22	560	473	12.06	−26.68
2x Nand2	280	394	9.96	−26.45
6x Nand2	832	1770	7.61	−25.42
Average	432.63	809	8.02	−27.03

Table 4. Circuit optimization and layout migration results for a 90-nm technology CMOS library.

Cell Name	Leakage Reduction (%)	Total Power Reduction (%)	Delay Increase (%)	New Cell Width (Columns)	Cell Width Change (1 Column = 0.28 mm)	Migration Run Time (h:min:sec)
2x Inverter	36	9	9	3	0	0:00:12
4x Buffer	37	11	13	6	0	0:00:28
1x Delay	33	14	11	8	−1	0:00:26
1x Mux2	33	17	13	10	1	0:00:55
2x Xnor2	35	17	11	11	0	0:00:50
2x Aoi22	33	8	11	7	0	0:00:43
2x Nand2	34	8	9	4	0	0:00:16
6x Nand2	35	12	9	11	1	0:00:39
4x D Flip-Flop	35	12	15	29	1	0:12:43
4x Scan Enable D Flip-Flop	34	15	15	55	3	0:57:17
AVERAGES	**35%**	**12%**	**12%**			**00:07:33**

current, the circuit optimizer modifies the input variables and then resimulates. This loop repeats until the optimization targets are met. If the targets cannot be met, the optimization terminates after a prespecified run time is reached. Once new transistor sizes are obtained, the new library generation/migration is completed by layout generation and library characterization (see Figure 12).

Circuit Optimization and Layout Migration Results

This flow was tested on a production ASIC library in a 90-nm technology at 1.2 V supply and 125°C. Eight representative combinatorial standard cells were chosen for testing. The selected cells cover most unique types of standard cells. As shown in Table 3, an average of 27% subthreshold leakage reduction is achieved at an average 8% active area increase. In the table, "Active Area" is the active channel area, and I_{off} is the cell leakage current. The delay degradation is less than 1%. These techniques are also applicable to sequential cells. If more delay degradation is permitted, one can choose to increase channel lengths only and maintain the same transistor widths. This can significantly speed up the optimization process. Table 4 shows the results of 10% channel length increase for all transistors in eight combinatorial cells and two sequential cells in a different library. The average leakage reduction is ~35% at the average cost of 12% delay degradation. The average layout migration time is about 7 min on a 3 GHz Intel Xeon Machine.

CONCLUSIONS

Leakage current is of great concern for designs in nanometer technologies. In 90- and 65-nm technologies, subthreshold leakage current dominates total leakage current. For a typical ASIC circuit running at several hundred megahertz frequencies, the subthreshold leakage power can be as high as 60% of total power. An important method for minimizing power in ASIC libraries is reducing leakage current. In this article, a complete automated leakage optimization flow that changes channel lengths and widths with cell delay and active area constraints was discussed. Optimization results show that there is ~30% leakage current reduction with a few percent active area and delay increase. There is increase in dynamic power, but the net total power reduction is significant. A uniform increase of 10% in gate length results in ~35% leakage reduction at the cost of ~12% delay degradation. The total cell area changes are minimal in both cases. The optimization flow begins with SPICE net lists from an existing library, optimizes leakage currents subject to performance metrics and active area increase constraints, and finishes with new layout generation and characterization. Investigations indicate that the leakage optimization has little impact on cell noise margin and layout parasitic modifications do not affect optimization results. The efficient automatic layout-to-layout cell leakage optimization flow is most suitable for leakage minimization and library migration for 90- and 65-nm ASIC libraries. Future work includes applying the flow to situations where layout-dependent DFM and process variation objective functions are also optimized.

REFERENCES

[1] "2005 international technology roadmap for semiconductor" [Online]. Available: http://public.itrs.net/

[2] K. Roy, S. Mukhopadhyay, and H. Mahmoodi-Meimand, "Leakage current mechanisms and leakage reduction techniques in deep-submicrometer CMOS circuits," *Proc. IEEE*, vol. 91, no. 2, pp. 305–327, Feb. 2003.

[3] A. Chandrakasan, I. Yang, C. Vieri, and D. Antoniads, "Design considerations and tools for low-voltage digital system design," in *Proc. Design Automation Conf.*, June 1996, pp. 728–733.

[4] R. Brodersen, M. Horowitz, D. Markovic, B. Nikolic, and V. Stojanovic, "Methods for true power minimization," in *Proc. IEEE Int. Conf. Computer-Aided Design*, San Jose, CA, 2002, pp. 35–42.

[5] S. Mukhopadhyay, A. Raychowdhury, and K. Roy, "Accurate estimation of total leakage in nanometer-scale bulk CMOS circuits based on device geometry and doping profile," *IEEE Trans. Computer-Aided Design*, vol. 24, no. 3, pp. 363–381, Mar. 2005.

[6] S. Borkar, "Circuit techniques for sub-threshold leakage avoidance, control, and tolerance," in *Proc. IEEE Electron Devices Meeting*, San Francisco, Dec. 2004, pp. 421–424.

[7] K. Zhang, U. Bhattacharya, Z. Chen, F. Hamzaoglu, D. Murray, N. Vallepalli, Y. Wang, B. Zhang, and M. Bohr, "SRAM design on 65-nm CMOS technology with dynamic sleep transistor for leakage reduction," *IEEE J. Solid-State Circuits*, vol. 40, no. 4, pp. 895–901, Apr. 2005.

[8] M. Mui, K. Banerjee, and A. Mehrotra, "Supply and power optimization in leakage-dominant technologies," *IEEE Trans. Computer- Aided Design*, vol. 24, no. 9, pp. 1362–1371, Sept. 2005.

[9] N. Jayakumar, S. Dhar, and S. Khatri, "A self-adjusting scheme to determine the optimum RBB by monitoring leakage currents," in *Proc. Design Automation Conf.*, Anaheim, CA, June 2005, pp. 43–46.

[10] F. Gao and J. Hayes, "Total power reduction in CMOS circuits via gate sizing and multiple threshold voltages," in *Proc. Design Automation Conf.*, Anaheim, CA, June 2005, pp. 31–36.

[11] D. Lee, D. Blaauw, and D. Sylvester, "Static leakage reduction through simultaneous Vt/Vox and state assignment," *IEEE Trans. Computer-Aided Design*, vol. 24, no. 7, pp. 1014–1029, July 2005.

[12] L. Yuan and G. Qu, "Enhanced leakage reduction technique by gate replacement," *Proc. Design Automation Conf.*, Anaheim, CA, June 2005, pp. 47–50.

[13] S. Lichtensteiger, L. Wissel, J. Engel, and P. Sulva, "Modeling leakage in ASIC libraries," *IEEE Custom Integrated Circuits Conf.*, San Jose, CA, Sept. 2005, pp. 609–612.

[14] N. Sirisantana, L. Wei, and K. Roy, "High-performance low-power CMOS circuits using multiple channel length and multiple oxide thickness," in *Proc. IEEE Int. Conf. Computer Design: VLSI Computers Processors*, 2000, pp. 227–232.

[15] Y. Taur and T.H. Ning, *Fundamentals of Modern VLSI Devices*. Cambridge, U.K.: Cambridge Univ. Press, 1998, ch. 2, pp. 94–95.

[16] B.J. Sheu, D.L. Scharfetter, P.K. Ko, and M.C. Jeng, "BSIM: Berkeley short-channel IGFET model for MOS transistors," *IEEE J. Solid-State Circuits*, vol. 22, no. 4, pp. 558–566, Aug. 1987.

[17] B. Chatterjee, M. Sachdev, S. Hsu, R. Krishnamurthy, and S. Borkar, "Effectiveness and scaling trends of leakage control techniques for sub-100 nm CMOS technologies," in *Proc. Int. Symp. Low-Power Electronics*, Seoul, Korea, 2003, pp. 122–127.

Xiaoning Qi, Sam C. Lo, Alex Gyure, Yansheng Luo, Mahmoud Shahram, Kishore Singhal, and *Don B. MacMillen* are with Synopsys, Inc, Mountain View, California. E-mail: Mahmoud.Shahram@synopsys.com.

CD ■