Sizing and Characterization of Leakage-Control Cells for Layout-Aware Distributed Power-Gating

Pietro Babighian ‡ Luca Benini * Enrico Macii ‡

‡ Politecnico di Torino
Torino, ITALY 10129

* Università di Bologna
Bologna, ITALY 40136

Abstract

This paper proposes a methodology for sleep transistor sizing for usage in a novel, single-threshold leakage cut-off approach, where power gating cells are distributed row-by-row in a fully placed circuit. Sizing equations are obtained by performing SPICE simulations for a 130nm technology. Furthermore, the layout of a test case is considered and power and delay values are extracted in order to demonstrate the practical impact of our solution.

1 Introduction

A number of leakage-reduction techniques move from the observation that sub-threshold current in a stack of OFF transistors is greatly reduced with respect the single transistor case. Quantitative analyses reported in the literature [1] show that leakage current can be decreased by one order of magnitude by simply stacking two transistors.

To reduce the performance impact associated with transistor stacking, a common technique is to connect a number of CMOS gates to a *virtual ground* node, which is then connected to the ground node through a large *sleep transistor*, whose gate is driven by a sleep-control signal.

We propose a novel sleep transistor insertion approach which is aware of the physical placement of the cells in a logic circuit and introduces a minimal amount of perturbation in the layout of the circuit. This paper focuses mainly on the development of analytical models for sizing the sleep transistor as a function of its load (AC and DC), given a specified acceptable switching speed penalty.

2 Row-Based Distributed Power-Gating

We propose to implement sleep transistors as special-purpose, *leakage control* cells and insert them in a fully placed design. Given a row-based placement, as produced by a commercial physical design tool, leakage-control cells are added at both sides of each row. A leakage cut-off circuit can be implemented by connecting a NMOS sleep transistor in series to the pull-down networks of the logic cells to be gated. In accordance to the number of logic gates to be controlled and the interconnection load capacitance C_L, the sleep MOSFET must be sized properly to bound acceptable performance loss. As a consequence, sleep transistors can be quite large and their gate capacitance C_G could be too high for the driving strength of the sleep control signal. Thus, insertion of properly sized buffering may be required.

3 Sleep Transistor Sizing

In general, the NMOS sleep transistor collects discharge currents from many cells on the same row, hence sizing depends on its aggregate current load. In the following, we conservatively assume that all gates connected to a sleep transistor can switch at the same time. Ensuring a limited performance degradation when the circuit is active is not sufficient. In fact, the time required to exit the sleep state must also be tightly controlled. As mentioned previously, when the sleep transistor is OFF, the virtual ground voltage V_g raises significantly above zero, due to the stacking effect. In general, we can expect that $V_g = V_{DD} - V_{TH}$ after a long sleep time. Upon exiting the sleep state, the sleep transistor is turned on but the gated logic is not active immediately. It takes some time to wake the circuit up since the voltage at the virtual ground must be first discharged to zero. This is required for the pull-down networks of the gated logic to operate correctly.

Limiting the slow-down in active state and reducing the activation time are the objectives of sleep transistor sizing. We have performed extensive characterization of both phenomena to obtain the transistor sizing curves.

We focus first on the activation time constraint. We have set a maximum activation time constraint (i.e., 5% of a fanout-4 inverter delay), and we have performed SPICE simulations in order to characterize the dependency of the sleep transistor width W from the virtual ground capacitance C_L, which depends on the number and type of the gated cells. During switch-on of the sleep transistor, we have performed transient analysis of its *drain* voltage and we have studied the impact of the virtual ground capacitance.

In our approach, leakage-control cells are inserted at the sides of the same row where the cells to be gated are placed. This strategy allows us to easily model the virtual ground capacitance C_L as a linear function of the row's length. For instance, if we assume that the row contains N identical cells, C_L can be expressed as:

$$C_L = K_1 + K_2 N \quad (1)$$

where K_1 and K_2 are fitting coefficients obtained from layout extraction of a set of sample circuits. In order to verify the validity of the linear model, we obtained K_1 and K_2 by least square fitting on a set of sample designs, where capacitance extraction was performed using the Cadence

1530-1591/04 $20.00 (c) 2004 IEEE

Silicon Ensemble placement and routing tool. On the basis of the results we have obtained, Equation (1) can be re-written as:

$$C_L = 0.70 + 1.43N \quad (2)$$

Statistical significance of the linear regression model was confirmed with a high level of confidence. In general, given a row-based placement, the load capacitance of the virtual ground is easily extracted and appropriate sizing of the sleep transistor is performed. Equation (2) can be used to obtain an approximation of C_L, before placement and routing. By referring to the characterization curve of Figure 1, the size of the sleep transistor can be determined for the given performance bound.

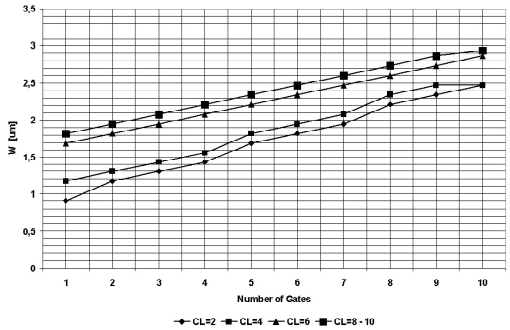

Figure 1: Sizing of the sleep transistor in OFF-state mode.

A similar sizing curve was obtained for the active-state gate switching speed degradation constraint. In this case, the characterization experiments were run by keeping the sleep transistor in the ON-state, and by forcing a falling edge on the output of the gated cells.
A 130nm CMOS technology was considered, with a V_{DD} of $1.3V$. All reported values of C_L are normalized to the gate capacitance of the smallest inverter implementation in our technology (i.e., $0.6fF$).

4 Validation: Case Study

The viability and effectiveness of the proposed leakage control methodology has been assessed on a small benchmark circuit, consisting of a total of 55 standard cells, 16 of which are synchronous, load-enable flip-flops and the remaining 39 are combinational gates. The standard cell library we used for our experiments is the 130nm HCMOS9 by STMicroelectronics.
The layout of the circuit, generated using Cadence Silicon Ensemble consists of a total of six rows. Sleep transistor insertion was performed on the third row of the layout, which contains a total of twelve combinational cells. Such cells were partitioned into two clusters. Two leakage control cells were then designed using the Cadence Virtuoso tool, following the rules of the standard cells in the library; each cell contained a sleep transistor and the appropriate buffering circuitry. The two leakage control cells were placed into the layout at each side of the considered row and connected to the above mentioned clusters separately. Figure 2 provides a detailed view of row three.
Post-layout simulation results are collected in Table 1. Leakage power savings range from 65% to 87%, depending on the cell (average savings over all the cells is 75%). The delay overhead is very limited, and never above 4.5%, thus compliant with the initial constraint of a maximum delay increase of 5%.

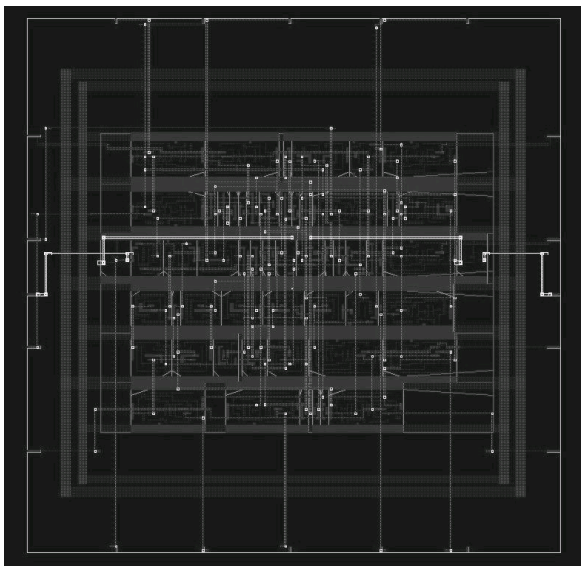

Figure 2: Details of the layout of row three.

Cell	*No Leak Control*		*Leak Control*		*Δ Power* [%]	*Δ Delay* [%]
	P_{LK} [nW]	*Delay* [ps]	P_{LK} [nW]	*Delay* [ps]		
G1	7.761	132.3	2.717	137.0	65.0	-3.5
G2	7.881	132.2	2.371	135.0	69.9	-2.1
G3	11.278	120.8	1.466	126.1	87.0	-4.3
G4	1.951	161.3	0.547	166.3	71.9	-3.1
G5	1.988	158.3	0.467	165.3	76.5	-4.4
G6	2.161	161.0	0.737	168.3	65.8	-4.5
G7	4.967	130.3	0.745	135.1	85.0	-3.6
G8	5.704	185.0	1.254	189.0	78.0	-2.1
G9	1.136	146.7	0.353	152.5	68.9	-3.9
G10	1.968	240.0	0.446	248.0	77.3	-3.3
G11	4.967	182.6	0.745	188.2	85.1	-3.0
G12	3.054	369.6	0.916	385.7	70.0	-4.3

Table 1: Leakage power savings and delay degradation.

5 Conclusions

In this paper, we have presented a novel methodology for sub-threshold leakage power reduction based on the idea of inserting sleep transistors into standard-cell circuits. The paper has focused on the development of sizing models of the sleep transistor which have been validated on a real layout, leading to leakage power savings around 75%, on average, for all the cells in the circuit.

References

[1] D. Lee, W. Kwong, D. Blaauw, D. Sylvester, "Simultaneous Subthreshold and Gate-Oxide Tunneling Leakage Current Analysis in Nanometer CMOS Design," *ISQED-03* pp. 287-292, San Jose, CA, March 2003.