Keeping Hot Chips Cool

Ruchir Puri
IBM T J Watson Research Center
Yorktown Hts, NY 10598
ruchir@us.ibm.com

Leon Stok
IBM
Somers, NY 10589
leonstok@us.ibm.com

Subhrajit Bhattacharya
IBM T J Watson Research Center
Yorktown Hts, NY 10598
sbhat@us.ibm.com

ABSTRACT
With 90nm CMOS in production and 65nm testing in progress, power has been pushed to the forefront of design metrics. This paper will outline practical techniques that are used to reduce both leakage as well as active power in a standard-cell library based high-performance design flow. We will discuss the design and cost issues for using different power saving techniques such as: power gating to reduce leakage, multiple and hybrid threshold libraries for leakage reduction and multiple supply voltage based design. In addition techniques to reduce clock tree power will be presented as power consumed in clocks accounts for a significant portion of total chip power. Practical aspects of implementing these techniques will also be discussed.

Categories and Subject Descriptors
B.7.2 [Hardware]: Integrated Circuits – Design Aids.

General Terms
Algorithms, Performance, Design.

Keywords
Low power, High-Performance, VLSI Design.

1. INTRODUCTION
Power dissipation is becoming the most challenging design constraint in sub-90 nanometer CMOS technologies. Exponential increase in sub-threshold and gate leakage power is slowing down CMOS scaling (Figure 1). Power consumption over 100W has become a common specification in multi-GHz high-performance designs. In addition, due to increasing variability in sub-90nm process, the variability in power dissipation has been a source of dropping yield, i.e., a significant portion of the yield is now power-limited in high-performance designs [5]. In order to continue the CMOS scaling, it is crucial to tame this exponential increase in power dissipation [2]. In this paper, we outline practical techniques that are used to reduce both leakage as well as active power in a standard-cell library based high-performance design flow.

DAC 2005, June 13–17, 2005, Anaheim, California, USA.

2. TECHNIQUES TO TAME POWER
Broadly speaking, power reduction techniques can be classified among two categories: Structural techniques, and Dynamic techniques [3]. Structural techniques are static in nature such that the power savings are built into the structure of the design (by construction) and no dynamic control at run-time is required. In contrast, dynamic techniques require active run-time control of a design parameter, e.g., frequency, voltage, or device threshold.

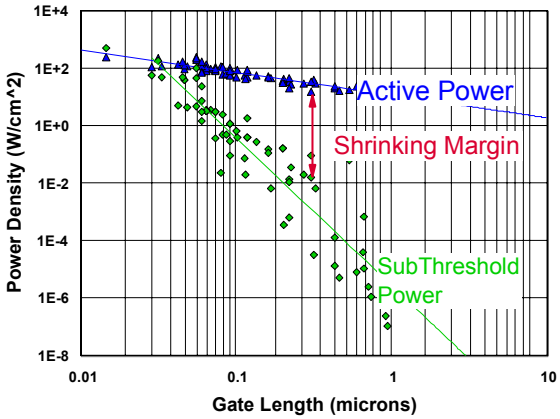

Figure 1. Variation in Power with CMOS sacling

Practical (both low-power and high-performance) designs utilize many well known structural techniques such as: *voltage Islands* (multiple-supply voltage design), *area reduction* of a design for minimizing active and leakage power, *multiple-threshold voltage* devices for reduction of sub-threshold leakage, *multiple-oxide thickness* for reduction of gate leakage, *latch clustering techniques* to minimize the loading on the clock network to minimize active power. Building parallelism at micro-architectural level is another example of structural power reduction technique. Designing multiple cores in the system trades-off scaling up of frequency with significant increase in power versus scaling out with multiple cores that yield equivalent performance accompanied with significant reduction in power consumption. In general, most of the techniques typically trade-off non-critical timing slack in the design for reduction in power (Figure 2). Thus, non-critical paths can be slowed down with lower supply voltage, higher-Vt, or thicker gate oxide devices in order to achieve significant reduction in power, which has a quadratic dependence on supply voltage.

Dynamic power reduction techniques such as: clock gating to reduce active power [10], power gating, i.e., switching off power to inactive components to reduce leakage power, dynamically varying clock frequency and power supply voltage regulation

depending on the workloads, as well as dynamically varying device thresholds with well-biasing techniques [9] have been used in several commercial designs.

In the following sections, we briefly discuss the implementation issues with some of the power reduction techniques mentioned above, such as multiple supply voltage islands, multiple-threshold voltage, latch clustering technique, and power gating.

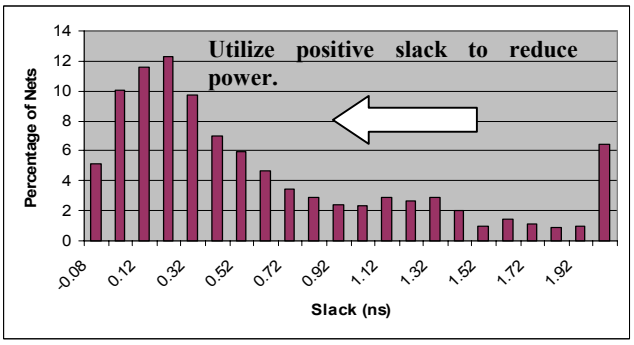

Figure 2. Trading off non-critical slack for power

2.1 Multiple-supply voltage islands

The use of multiple supply voltages presents some unique physical and electrical challenges [2][4] Level shifters need to be introduced between the various voltage regions. The physical layout needs to be designed to ensure the efficient delivery of the correct voltage to various voltage regions. More flexibility can be gained by using appropriate level shifters. For example, in an SoC that integrates a processor core with memory and control logic, the performance critical processor core requires the highest voltage to maximize its performance [7]. However, the on chip memory and control logic may not require the highest voltage operation and can be operated at a reduced voltage to save significant active power without compromising system performance. In addition, voltage flexibility at the unit level allows pre-designed standard components from other applications to be reused in a new SoC application. In contrast to the macro-based voltage island methodology, a flexible physical design approach which allows flexible voltage islands and enables a fine grained implementation of the dual-supply voltage assignment in a placement driven synthesis framework could have advantages in high-performance designs. However, it requires a significantly more complex design methodology for implementation. A flexible voltage island processor core is shown in Figure 3, where we can assign different voltages at both macro and cell levels, and it has more freedom in terms of layout style by allowing multiple voltage islands within the same row. First, timing closure is run with the entire circuit timed at VddH. For deep submicron circuits, interconnect delay dominates the gate delay. Thus we need rough placement information to identify critical versus non-critical cells. Once global placement is determined and timing is more or less closed, we can perform the flexible voltage island generation, by assigning non-critical cells to a lower supply voltage.

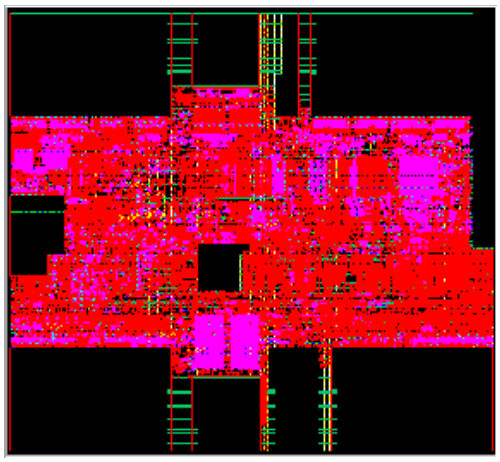

Figure 3. Processor core with flexible voltage island.

2.2 Multiple-threshold Voltage Technique

Numerous publications have addressed the issues related to multiple threshold voltage optimization techniques. In order to reduce sub-threshold leakage, timing non-critical gates are assigned to higher-device thresholds so that they consume less power at the expense of some performance on these non-critical paths. Significant reduction in leakage power can be obtained by utilizing a multiple-Vt library which is a standard offering in sub-90nm technologies[2] [4][6] .

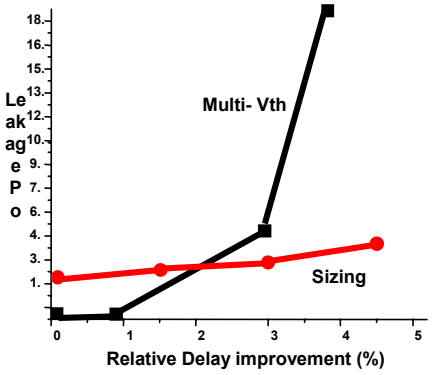

Figure 4. Delay-leakage power trade-off with multi-Vt and gate sizing.

As shown in Figure 4, resizing gates is a much better way to achieve desired performance while minimizing power in leakier CMOS technologies. Therefore, a preferred methodology for multiple threshold voltage optimization in high-performance designs is to perform logic and physical optimizations with nominal threshold voltage devices and towards the end of the timing closure flow, selectively substitute low-Vt cells on the timing critical paths to boost performance, and high-Vt cells on the non-critical paths to reduce leakage. While substituting high-threshold voltage cells, we must also take the Vt mistrack of manufacturing process into account and keep some timing margin in order to avoid high-threshold voltage cells on the critical paths in real hardware. Gate sizing must be concurrently performed

along with multiple threshold voltage optimization in order to achieve a good power-performance trade-off.

2.3 Latch Clustering for Power reduction

It is well known that clocking can contribute up to 50% of the total active power in high-performance multi-GHz designs. In well designed balanced clock trees, most of the power is consumed in the lower level clock buffers, i.e., portion of the tree closer to the latches. Clocking power can be significantly reduced by constraining the latches closer to the local clock buffers which are generating the local clock signal for latches. As shown in Figure 5, this implies clustering the latches around the local clock buffer (LCB). Intuitively, it may appear that a very flexible placement of latches closer to the data cells may be beneficial for performance. However, implementation in real high-performance designs have shown that constraining the latches could have positive effect on timing as well due to clock skew being minimized, as all the latches end up being clustered around local clock buffers (as shown in the Figure 6). Savings in power are obtained as a result of the reduction in wire load being driven by the clock buffer. As shown in Figure 7, Out of twelve local clock buffers in one of the cross-sections of a high-performance design, capacitive load on C1/C2 (two phase clock) signal can be reduced by over 50% in many cases, relative to unconstrained latch placement. This reduces the total capacitance by up to 40% in many cases, which directly translates into power saving for the local clock buffer.

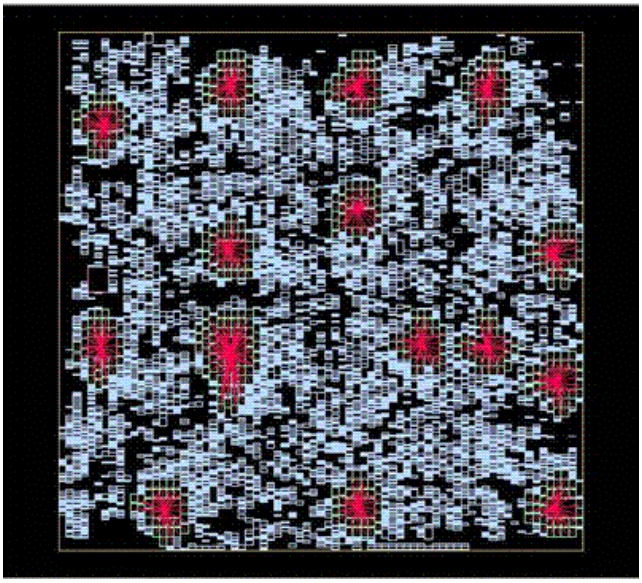

Figure 5. Latch Clustering around LCBs in a high performance Macro (design partition)

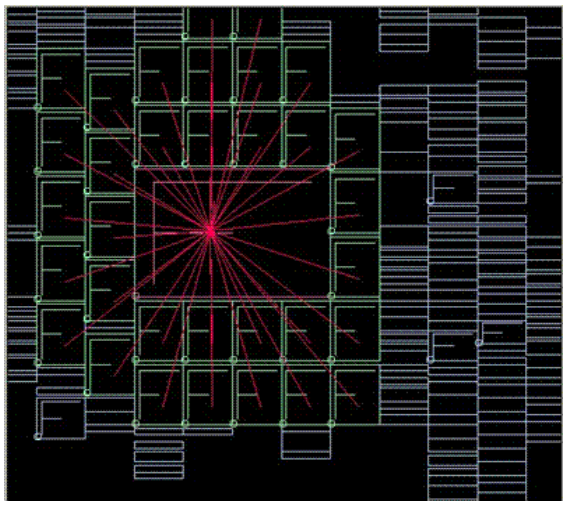

Figure 6. Cluster of Latches around a single LCB

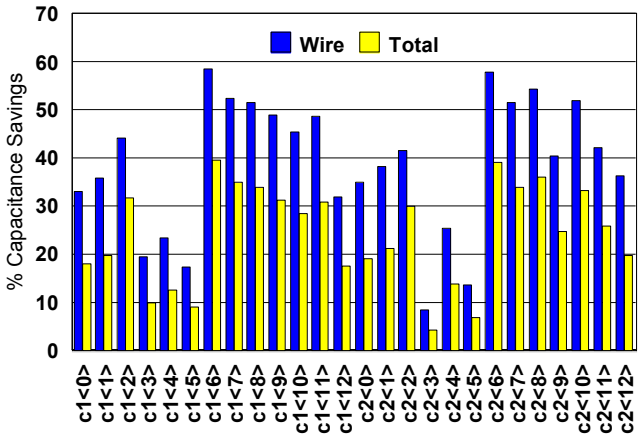

Figure 7. Capacitance reduction on C1/C2 clock signals from LCB to latches due to latch clustering .

2.4 Power Gating to reduce Leakage Power

Exponential increase in leakage power has been one of the most challenging issues in sub-90nm CMOS technologies. Power gating is one of the most effective techniques to reduce both sub-threshold leakage and gate leakage as it cuts off the path to the supply [3][8]. Conceptually, it is a straightforward technique, however, the implementation can be quite tricky in high-performance designs where the performance trade-off is constrained to less than 2% of the frequency loss due to power gate (Footer/Header switch insertion). Figure 8 shows a simple schematic of a logic block that has been power gated by a header switch (PFET) or a footer switch (NFET). Obviously, footer switches are preferred due to the better drive capability of NFETs. Operationally, if the logic block is not active, the SLEEP signal can turn off the NFET (footer switch) and the virtual ground (drain of NFET) will float towards Vdd (supply voltage), thereby reducing the leakage by orders of magnitude.

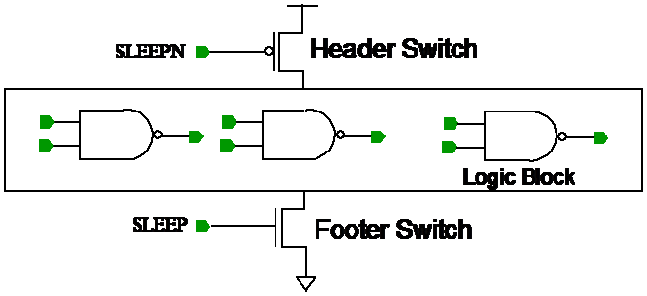

Figure 8. Power gating using Header/Footer switches.

Obviously, introducing a series transistor (footer/header) in the logic path results in a performance penalty. This performance penalty can be mitigated by making the size of the footer/header larger so as to reduce the series resistance. However, the leakage benefit reduces with increasing size of power gate. Practically, in low-power applications, over 2000X leakage saving can be obtained at the expense of 8-10% reduction in performance. However, in high-performance designs, this is a relatively large performance penalty. So, larger power gate sizes are chosen (approximately 6-8% of logic area) to achieve less than 2% performance penalty with over 20X leakage reduction.

In general, power gating can be physically implemented in the designs using: Macro-based coarse grained power gating and intra-macro fine power gating (similar to multiple-supply voltages). In macro-based implementation, the footer (or header) switches surround the boundary of the macro, as shown in Figure 9. From a physical design point, this implementation is easier due to the fact that it does not disturb the internal layout of the macro. However, it has a potential drawback in terms of larger IR drop on the virtual ground supply. For IP blocks, this is the preferred implementation technique for power gating. Fine grained power gating, as shown in Figure 10, where the footer switches are implemented within the logic in a regular layout are more desirable in a high-performance design where the voltage drop across the power gate as well as IR and EM (electro-migration) requirements are more stringent.

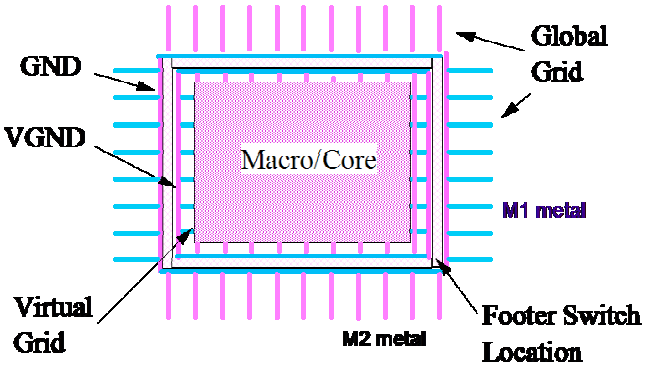

Figure 9. Coarse-grained power gating with Macro/Core

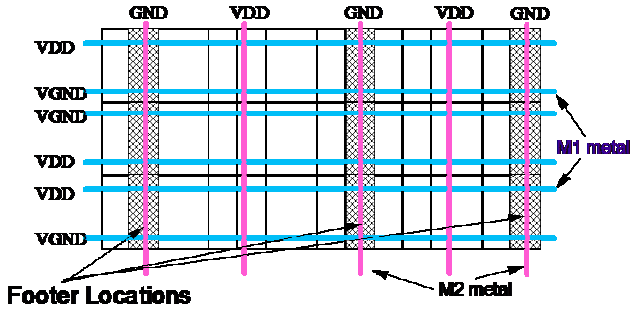

Figure 10. Fine-grained power gating within Macro

3. CONCLUSIONS

Power dissipation has become the most challenging design constraint in sub-90 nanometer CMOS technologies. In order to continue the CMOS scaling, it is crucial to tame the exponential increase in power dissipation. In this paper, we outlined some practical techniques that are used to reduce both leakage as well as active power in a standard-cell library based high-performance design flow. Issues related to implementing these techniques in high-performance designs were outlined as well.

REFERENCES

[1] Trevillyan, L. et al., *An integrated design environment for technology closure of deep-submicron IC Designs*, IEEE Design & Test, Feb. 04.

[2] Puri, R. et al, *Pushing ASIC performance in a power envelope*, DAC 2003.

[3] Kosonocky, S. et al, *Low Power Circuits and Technlogy for wireless digital systems*, IBM Journal of R&D, Vol. 47, No. 2/3, 2003.

[4] Puri, R., et al., *Minimizing power with Flexible Voltage Islands*, ISCAS 2005.

[5] Nassif, S., *Delay Variability: Sources, Impact and Trends*, ISSCC, 2000.

[6] Warnock, J., "*Circuit Design with Leaky Transistor*", ISSCC Workshop on Design in power constrainted era, 2003.

[7] Lackey, D. et al, "*Managing power & performance for SoC Designs using Voltage Islands*", ICCAD 2002.

[8] Kim, S. et al., *Understanding and minimizing ground bounce during mode transition of power gating structures*., ISLPED 2003.

[9] Keshavarzi, A et al, *Effectiveness of reverse body biasing for scaled technologies*, ISLPED 2001.

[10] Roy, K., *Clcok Gating for Microprocessor power reduction*, High Performance Computer Architecture Symposium, 2003.