國 立 交 通 大 學

電信工程學系

碩 士 論 文

基於一百奈米製程之具區塊層級最佳化的複合式溫度感應功率管理系統

A Comprehensive Thermal-Aware Power Management System with Block-Level Optimization in 100nm CMOS Technology

研究生：詹偉閔

指導教授：闕河鳴　博士

中　華　民　國　九　十　四　年　九　月

基於一百奈米製程之具區塊層級最佳化的複合式溫度感應功率管理系統

A Comprehensive Thermal-Aware Power Management System with Block-Level Optimization in 100nm CMOS Technology

研 究 生：詹偉閔　　Student: Wei-Min Chan

指導教授：闕河鳴　博士　　Advisor: Dr. Herming Chiueh

國 立 交 通 大 學

電 信 工 程 學 系 碩 士 班

碩 士 論 文

A Thesis

Submitted to Institute of Communication Engineering

College of Electrical Engineering and Computer Science

National Chiao Tung University

in Partial Fulfillment of the Requirements

for the Degree of

Master of Science

in

Communication Engineering

September 2005

Hsinchu, Taiwan

中 華 民 國 九 十 四 年 九 月

基於一百奈米製程之具區塊層級最佳化的複合式溫度感應功率管理系統

研究生：詹偉閔　　指導教授：闕河鳴　博士

國立交通大學
電信工程學系碩士班

摘要

由於電池壽命有限的緣故，現今的單晶片整合與手攜式系統強調低功率技術的應用。傳統的功率管理設計主要是針對動態功率損耗的降低，而最近的設計則開始將靜態功率損耗納入考量，因為在奈米製程下漏電流逐漸成為不可忽略的要素。最新的研究善用了單晶片設計思維的模組化特性，發展出區塊層級的控制技術來降低功率損耗，然而區塊之間的溫差對於單晶片系統設計所帶來的影響卻鮮少被討論。職是之故，本論文提出了複合式熱感應功率管理及其區塊層級的最佳化。所提出的設計應用了多種低功率技術來控制晶片內不同的功率損耗來源，同時顧及溫差所帶來的影響以維持各區塊間效能的一致性。模擬結果顯示，本設計對於電路的穩定性有顯著的提升，而對於漏電流也得到有效的鉗制。此模擬結果基於台積電一百奈米互補式金氧半導體製程。

A Comprehensive Thermal-Aware Power Management System with Block-Level Optimization in 100nm CMOS Technology

Student: Wei-Min Chan Advisor: Dr. Herming Chiueh

SoC Design Lab, Department of Communication Engineering,
College of Electrical and Computer Engineering, National Chiao Tung University
Hsin-Chu 30050, Taiwan

Abstract

Modern SoC integrations and mobile systems have emphasized low power techniques due to shortage of battery life. Conventional power management designs focused on the reduction of dynamic power consumption, recent designs begin to take leakage power into consideration since it becomes an important factor in nano-scale CMOS technology. Latest development has taken advantage of modularity in SoC design methodology to develop the block-level control technique for power reductions. However, thermal gradient over the system and its impacts to SoC designs are barely discussed. In this thesis, a block-level optimization of comprehensive thermal aware power management is presented. The proposed design applies several low power techniques to control different power sources and handles thermal impacts to provide performance coherence. As a result, optimal power reductions and performance coherence can be guaranteed within the whole system. The simulation results show a significant improvement in stability and leakage power reduction for most circuitries. These results are based on TSMC 100nm CMOS technology.

Acknowledgement

本篇碩士論文得以順利完成，首先要感謝的是我的指導教授闕河鳴博士。闕老師總在研究遇到瓶頸的時候給予寶貴的建議，使我得以突破與進步。除此之外，老師在平日培養學生獨立思考與分析問題的能力，更使我在研究領域建立正確的態度。

其次，感謝漳源、庭瑋、明崇、芳如，四位同窗在我研究上與生活上，給予諸多的幫助與指教，尤其懷念為了完成實驗而一同挑燈夜戰的時光。在此也謝謝學弟們的支持與幫忙，有了各位晶片系統設計實驗室成員的陪伴，才讓我有這麼一段快樂且充實的研究經驗。另外再感謝的是交大電子所博士班的華重憲學長，給予許多理論和實驗的寶貴建議。

最後，我要感謝父母的栽培，以及所有關心我的家人與朋友，唯有藉著大家的愛護和幫忙，才能造就今日的我。

我誠心感謝上述提攜或幫助過我的你們，謝謝大家並祝福大家。

詹偉閔

中華民國九十四年九月於新竹

Content

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Motivation

Power dissipation in modern VLSI designs has become the most critical issue in System-on-Chip era. The ever increasing on-chip integrations in recent decade have enabled a dramatically increase in system performance and scale. Unfortunately, accompanied with the performance improvement, a significant increase in power dissipation and heat density is introduced [1]. However, in modern VLSI circuitry of mobile systems, especially for handheld audio and video applications, low power considerations are becoming an important issue as battery life and geometry of mobile systems are limited [2, 3]. The demand for increased battery life will require designers to seek out new technologies and circuit techniques to maintain high performance and long operational lifetimes. Therefore, power and thermal issues have become the major limitation of such systems. Low power circuitry designs and architecture-level power reduction techniques have become more important in modern System-on-Chip implementation.

Generally, power consumption is categorized to three portions in VLSI circuitry: short circuit, leakage, and dynamic switching power. Traditionally, dynamic switching power is the dominant component of power consumption. However, as scaling trends continue in future generations, leakage power has become comparable to dynamic power, and potentially a dominant component of overall power consumptions [4-6]. Actually, leakage power is approaching 40% of the total power consumption in today's high-performance microprocessor design [6]. Therefore, unlike previous power reduction methods, modern power management techniques for SoC designs have to deal with static leakage power as equally important as dynamic switching power.

Several power reduction techniques can be adopted for modern power management designs, such as dynamic voltage scaling (DVS) [2, 7-11], clock gating [12, 13], power gating [1, 5, 6, 14-17], body bias [6, 18-23], and voltage islands [24-26]. DVS and clock gating can effectively reduce dynamic power consumption,

while power gating reduces static leakage power. Forward body bias can increase speed, while reverse body bias reduces power as well. Voltage islands can provide block-level control of power management, which makes control mechanism more flexible. These power reduction techniques are mainly circuit-level to architecture-level, meaning that given corresponding standard cell library, hardware-language coding remains unchanged and current EDA flows are compatible. It is possible to adopt these techniques to already established Software Intellectual Property (Soft-IP), which is in accordance with the concept of IP reuse in SoC designs.

Besides battery life, thermal impact is another major reason to utilize power management [8]. Power consumption of the chip contributes thermal impact, such as frequency sacrifice, leakage increase, and even circuit breakdown. Because of high costs of packaging and cooling requirements for reliability, power reduction is also addressed to face thermal impact. However, system temperature is not temporally and spatially constant, so adaptive power control doesn't imply adaptive thermal control. Thermal gradient over the system and its impact to SoC designs are barely discussed in current power management schemes.

Shown in Figure 1.1 to Figure 1.3 is the thermal impact to operating frequency, power consumption, and leakage power of a ring oscillator. We notice that there are about 13% frequency decrease, 18% total power consumption increase in the active mode, and 9X leakage power increase in the standby mode from 0 to 125°C. Functional blocks will not keep performance coherence at the same target frequency due to temperature difference. Leakage problems will become worse and worse as temperature rises. System will become asynchronous and power will not be controlled within specifications if we ignore these changes. Therefore, power management of SoC chip should take temperature variation into consideration.

In this thesis, we present a comprehensive thermal-aware power management (CTAPM) architecture and emphasis thermal optimization in block-level control. The proposed design not only deals with dynamic switching power and static leakage power, but also with temperature variation to dynamically adjust performance and power consumption of each block. Our goal is to make power and performance scalable as well as performance coherence between functional blocks in the meantime.

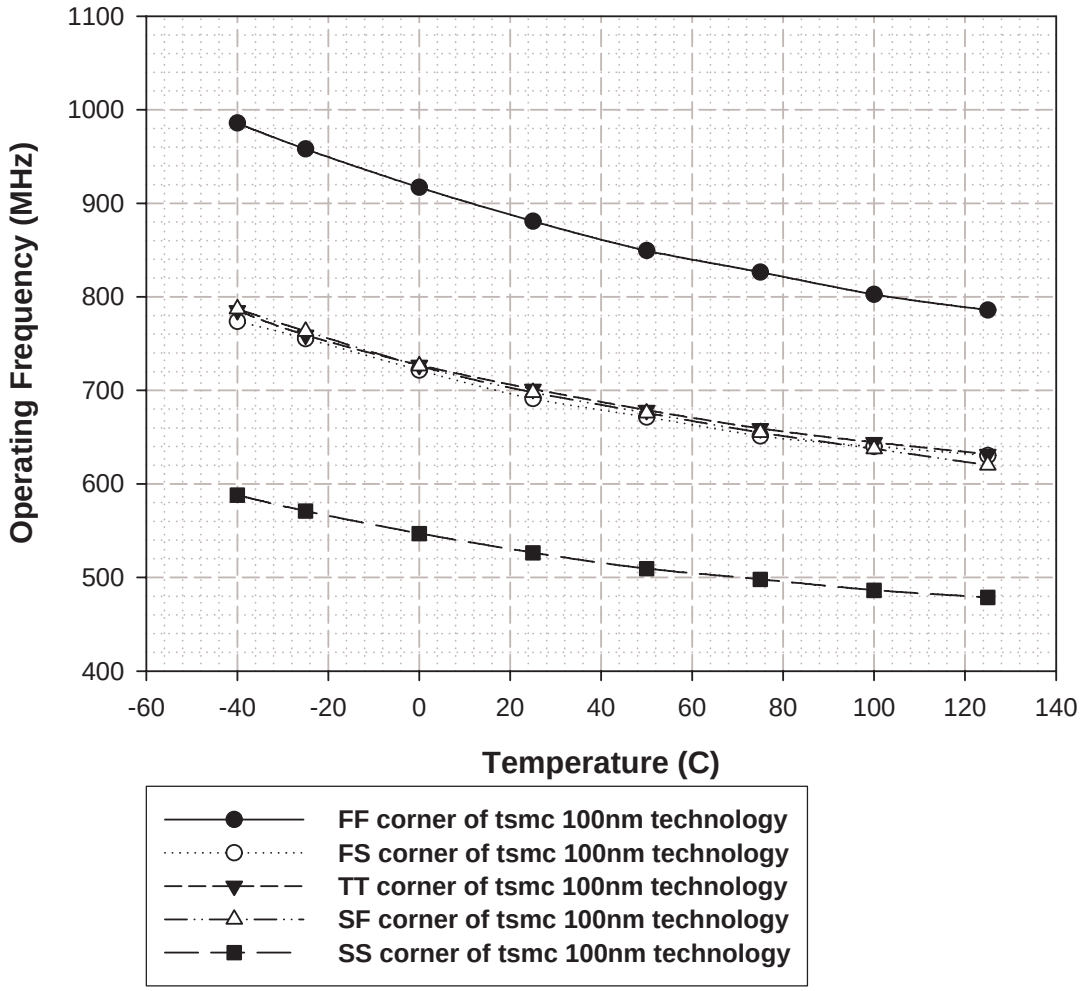

Figure 1.1 Thermal impact to operating frequency of a ring oscillator in TSMC 100nm CMOS technology. Supply voltage is 1.0V.

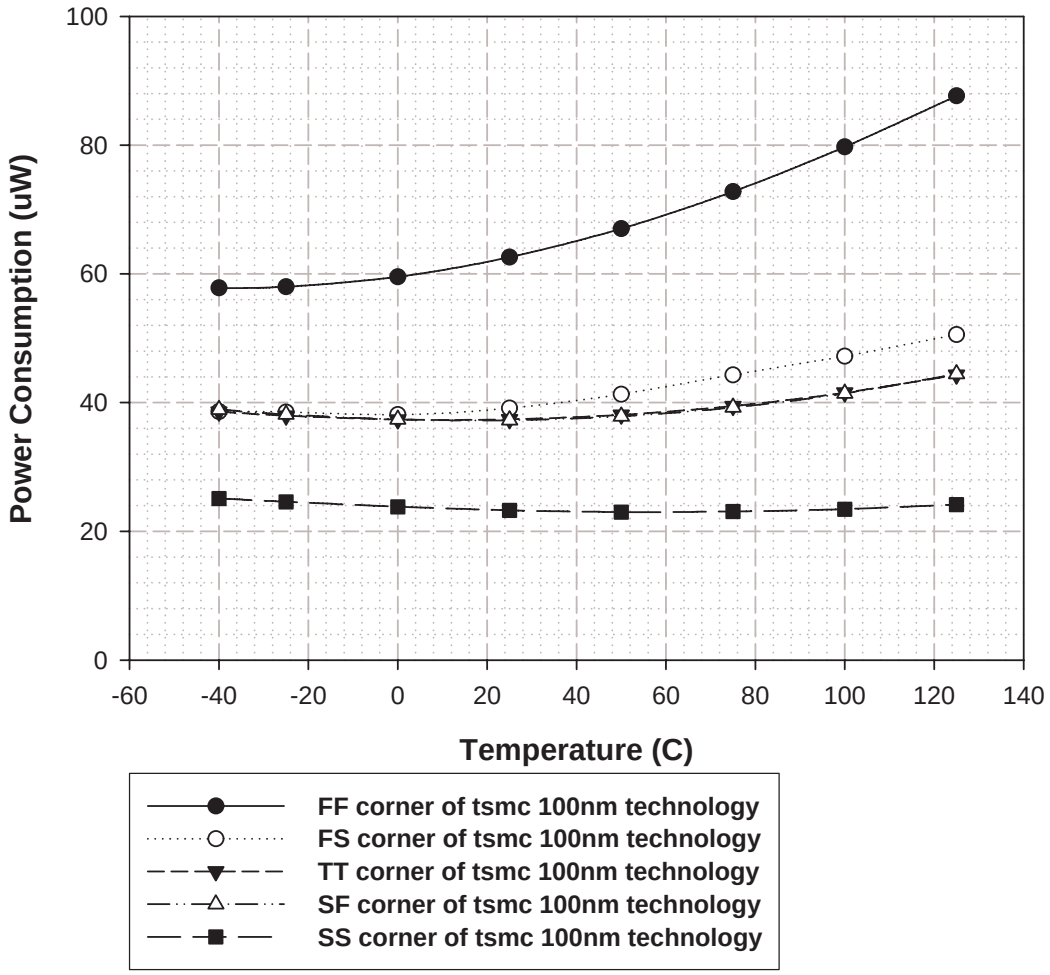

Figure 1.2 Thermal impact to power consumption of a ring oscillator in TSMC 100nm CMOS technology. Supply voltage is 1.0V.

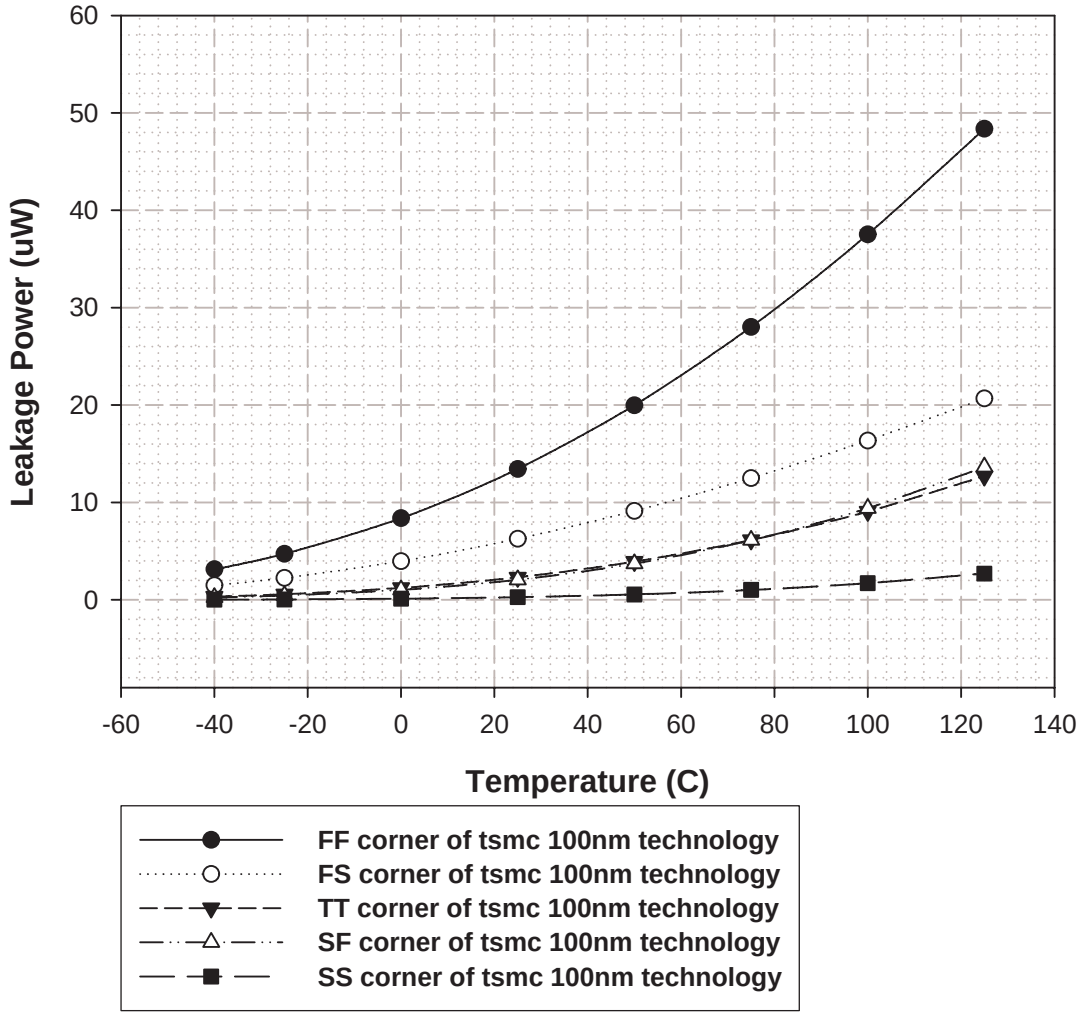

Figure 1.3 Thermal impact to leakage power of a ring oscillator in TSMC 100nm CMOS technology. Supply voltage is 1.0V.

1.2 Organization

In the beginning of Chapter 2, categories of power sources in digital ICs are discussed. Next comes the overview of low power techniques. After them, the Enhanced Intel SpeedStep® Technology, a state-of-the-art example of power management design, is presented. Finally, contributions and features of our comprehensive thermal-aware power management design end this chapter.

In Chapter 3, we propose the CTAPM unit architecture. Possible control flow chart, control state definition, and look-up table design are also introduced in this chapter. This CTAPM is designed to provide dynamic block-level power control and ensure performance coherence between functional blocks. Moreover, the test vehicle and analysis flow for power analysis are described in the end of this chapter.

In Chapter 4, the effect and efficiency of low power techniques are discussed. The experimental results of these different circuitry techniques are presented to see if they can provide needed functionalities for our CTAPM design. These results are based on TSMC 100nm CMOS technology,

In Chapter 5, conclusions and future works are addressed.

Chapter 2

Overview of Low Power Techniques

This chapter begins with the classification of power sources in VLSI circuitry. Several low power techniques are examined next to realize their functionalities. Afterwards, the Enhanced Intel SpeedStep® Technology is introduced as an example of modern power management design. Finally, contributions and features of our comprehensive thermal-aware power management (CTAPM) design are summarized in the end of this chapter.

2.1 Power Source Categories

In modern digital integrated circuits, power consumption is categorized to three components: dynamic, short circuit, and static leakage power.

2.1.1 Dynamic Power

Conventionally, dynamic power is the dominant component of total power consumption. Charging and discharging load capacitance contribute this power. Dynamic power is given by

$$P_{dynamic} = p_t \times f_{clk} \times C_{switched} \times V_{DD}^{\ 2} \tag{2.1}$$

where p_t is the switching activity, f_{clk} is the operating frequency, $C_{switched}$ is the total effective switched capacitance, and V_{DD} is the supply voltage. Obviously, scaling down V_{DD} is the most efficient way to reduce dynamic power, since dynamic power is proportional to the square of V_{DD}. However, V_{DD} cannot be scaled down unlimitedly because the operating frequency also degrades with lower supply voltage.

2.1.2 Short-Circuit Power

The second component of power consumption is called short-circuit power, which results from non-zero rise time and fall time of the input waveforms. As shown in Figure 2.1, non-zero input rise time and fall time make PMOS and NMOS turn on at the same period of time, setting up a DC path between V_{DD} and ground, therefore introducing the short-circuit current to consume power.

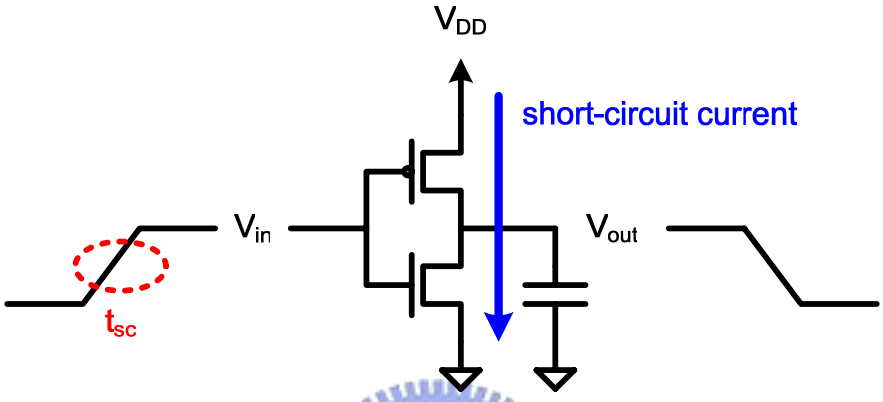

Figure 2.1 Short-circuit power is generated because of non-zero rise time and fall time of the input waveforms.

Short-circuit power can be expressed as

$$P_{short-circuit} = p_t \times t_{sc} \times f_{clk} \times I_{peak} \times V_{DD} = p_t \times f_{clk} \times C_{sc} \times V_{DD}^{\ 2} \qquad (2.2)$$

where t_{sc} is the time when the DC path exists, I_{peak} is the maximum of short-circuit current, and C_{sc} is the equivalent switched capacitance. Just like dynamic power, short-circuit power is also proportional to the switching activity. Besides, if V_{DD} can be scaled down below the sum of absolute values of NMOS and PMOS threshold voltage, which means $V_{thn} + |V_{thp}|$, then short-circuit power will be eliminated because NMOS and PMOS won't turn on at the same time.

2.1.3 Leakage Power

The third component of power consumption is leakage power, which grows larger and larger in nano-scale CMOS technologies. Leakage current can be divided into several portions, as illustrated in Figure 2.2 [4]. I_1 is the reverse-bias pn junction

leakage, I_2 is the subthreshold leakage, I_3 is the oxide tunneling current, I_4 is the gate current due to hot-carrier injection, I_5 is the gate-induced drain leakage (GIDL), and I_6 is the channel punchthrough current.

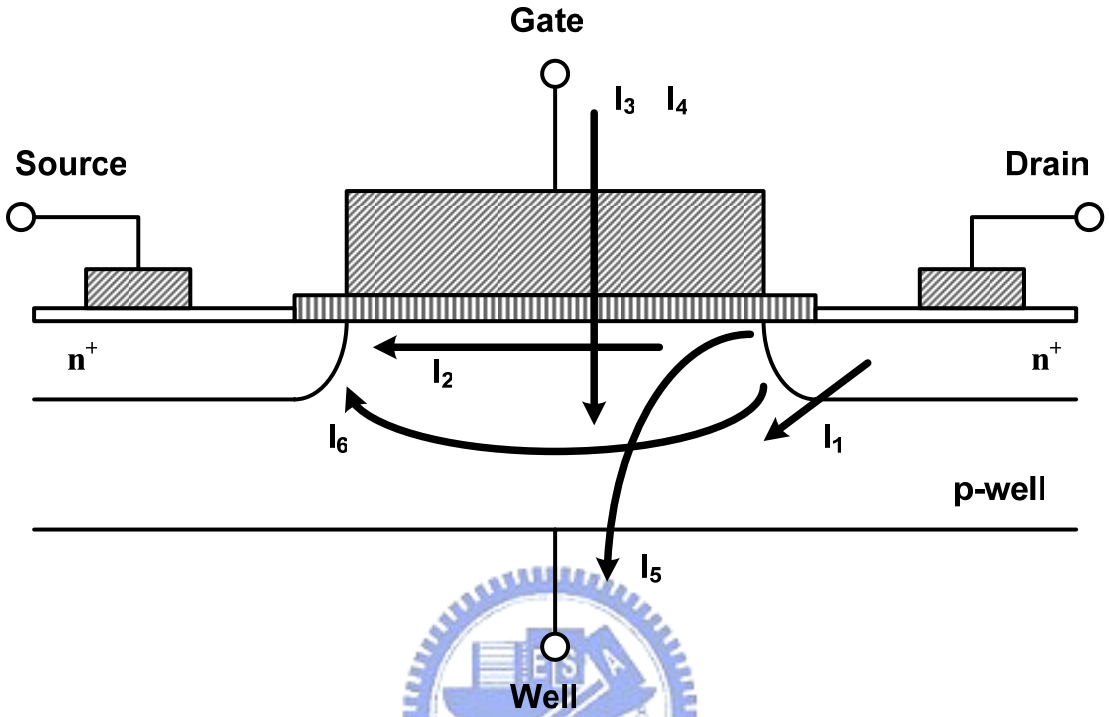

Figure 2.2 Summary of leakage current mechanisms of deep-submicrometer transistors.

Among these leakage sources, subthreshold leakage is dominant and can be modeled as [4]

$$I_{subth} = A \times e^{(V_G - V_S - V_{th0} - \gamma' \times V_S - \eta \times V_{DS}) / mV_T} \times (1 - e^{-V_{DS}/V_T}) \tag{2.3}$$

where

$$A = \mu_0 C_{ox} \frac{W}{L_{eff}} (V_T)^2 \times e^{1.8} \times e^{-\Delta V_{th}/\eta V_T} \tag{2.4}$$

V_{th0} is the zero bias threshold voltage, V_T is the thermal voltage, γ' is the linearized body effect coefficient, and η is the drain-induced barrier lowing (DIBL) coefficient. μ_0 is the zero bias mobility, C_{ox} is the gate oxide capacitance, m is the subthreshold swing coefficient of the transistor, and ΔV_{th} is a term introduced to account for transistor-to-transistor leakage variations. Equation (2.3) reveals that subthreshold leakage is larger when threshold voltage is lower due to their inverse exponential relationship. It is serious in nanometer devices since threshold voltage follows the scaling trends with the progress of CMOS technology.

Subthreshold leakage is the dominant component of leakage power consumption in deep-submicron and nanometer CMOS technologies. However, gate-tunneling leakage has the potential to exceed subthreshold leakage in future generations. The equation of gate tunneling leakage is given by [4]

$$J_{DT} = AE_{ox}^{\,2} \times e^{\{-B[1-(1-\frac{V_{ox}}{\phi_{ox}})^{3/2}]\}/E_{ox}} \tag{2.5}$$

where

$$A = \frac{q^3}{16\pi^2 \hbar \phi_{ox}},\ B = \frac{4\sqrt{2m^*}\,\phi_{ox}^{\,3/2}}{3\hbar q} \tag{2.6}$$

V_{ox} is the voltage drop across the gate oxide, ϕ_{ox} is the Si-SiO_2 interface barrier height for electrons, and E_{ox} is the field across the oxide. As shown in Figure 2.3 [4], gate-tunneling leakage can be further divided into five components, namely, parasitic leakage current through gate-to-source and gate-to-drain extension overlap region (I_{gso} and I_{gdo}), gate to inverted channel current (I_{gc}) which goes to the source (I_{gcs}) and the drain (I_{gcd}), and the gate to the substrate leakage current (I_{gb}).

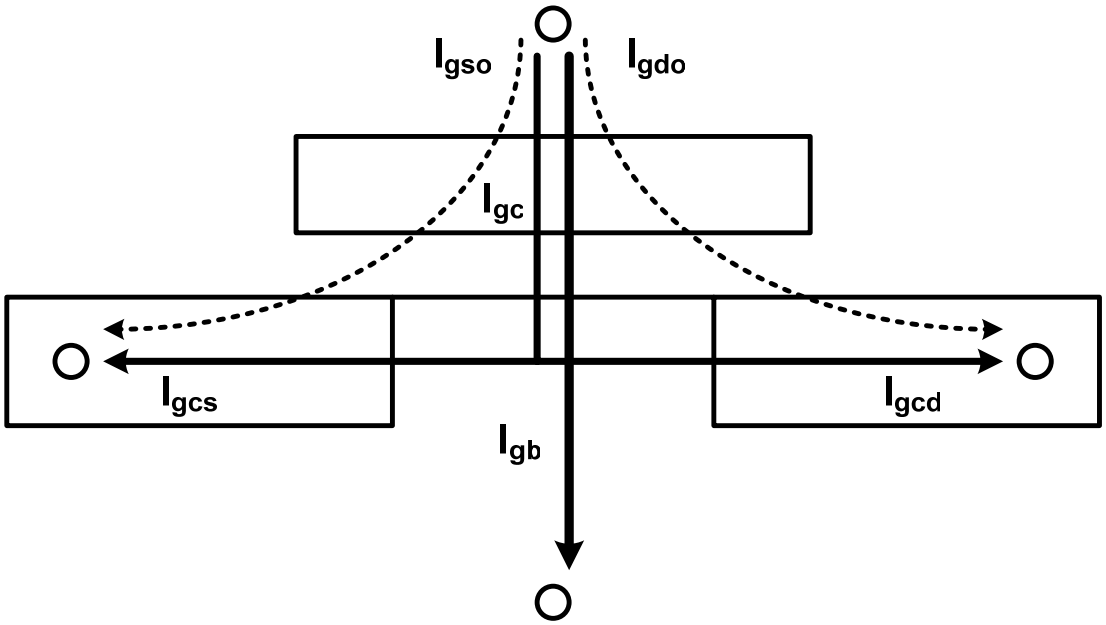

Figure 2.3 Components of gate-tunneling leakage current.

Unlike dynamic and short-circuit power which depend on the number of actively switching transistors, leakage power is a function of the total number (or area) of on-chip transistors, independent of their switching activity [1]. Although scaling V_{DD} down also helps to reduce subthreshold and gate-tunneling leakage, its limitations of frequency degrading and noise margin decrease still exist. However, power gating and reverse bias are two effective ways to suppress leakage power, which will be explained later in Section 2.2.

2.1.4 Summary

Power consumption of a VLSI chip can be categorized to three components: dynamic, short circuit, and leakage power. Traditionally, dynamic power is dominant, but leakage power has become comparable to dynamic power in nano-scale CMOS technologies. Subthreshold leakage is the dominant component of leakage power nowadays. However, gate-tunneling leakage has the potential to exceed subthreshold leakage in future generations.

Several low power techniques have been proposed to reduce and control power consumption dynamically, such as dynamic voltage scaling (DVS), clock gating, power gating, body bias, and voltage islands. These techniques will be introduced respectively in the next section.

2.2 Low Power Techniques

DVS, clock gating, power gating, body bias, and voltage islands are five major circuit-level techniques to reduce power consumption. They deal with different power sources and have individual control mechanisms.

2.2.1 Dynamic Voltage Scaling

Dynamic voltage scaling is the most efficient technique to reduce dynamic power. It also has benefits for leakage power saving. When the required performance of the target system is lower than the maximum performance, supply voltage can be dynamically reduced to the lowest possible extent that ensures proper operation. Significant power reduction is possible, since dynamic power of CMOS circuits is proportional to the square of the supply voltage.

Shown in Figure 2.4 is an example of the microprocessor system's desired throughput as a function of time [9]. The computational requirements can be considered to fall into one of three categories: compute-intensive, low-speed, and idle (standby). Low-speed and long-latency tasks only require a fraction of the full throughput to adequately run. Executing these tasks faster than needed has no discernible benefit. Therefore, the supply voltage has room to be scaled down in this situation to save dynamic power.

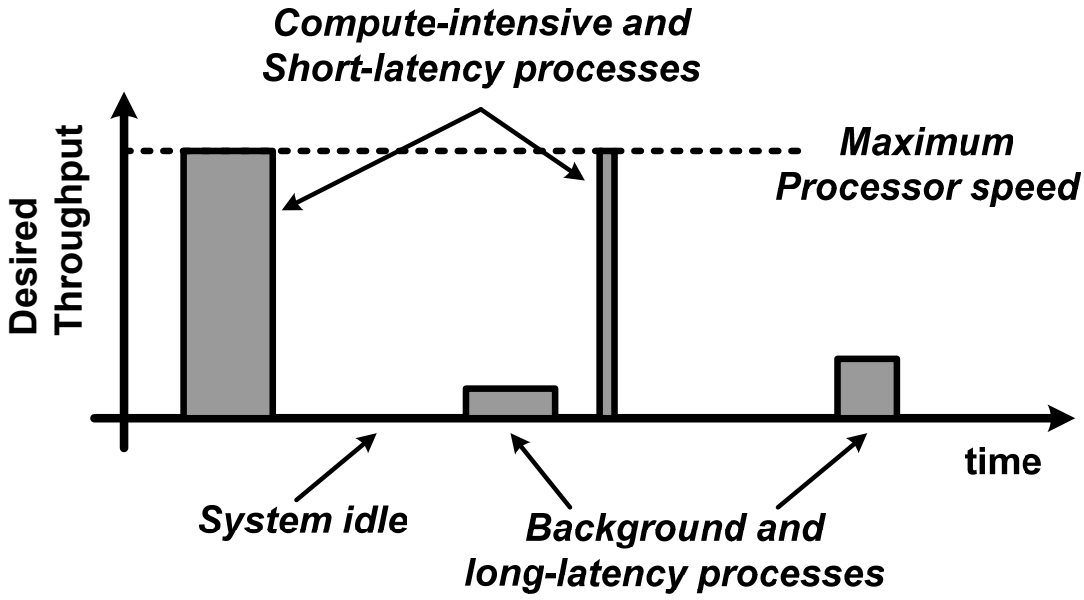

Figure 2.4 Processor usage model.

The conventional DVS system architecture is shown in Figure 2.5 [7]. The value of the supply voltage is determined by the computational loading and adjusted by a hard-wired frequency-voltage feedback loop. A ring oscillator, which is the replica of the critical path, is used to model the CMOS circuit delay for given supply voltage. However, this hardware-feedback approach does not provide efficient voltage control and has the following problems [7]:

1. Even in a same chip, critical path may be different along supply voltage, meaning that circuit delay of ring oscillator should have much margin to cover this variation.
2. Since fabrication process technology is different for each chip, circuit delay characteristics may differ a lot, meaning that all chips should be custom-designed to have same frequency-voltage relationship for efficient power reduction.
3. This approach cannot be applied for off-the-shelf processors, since ring oscillator cannot be inserted into ready-made chips.
4. In the multi-processor system, it is desirable to control supply voltages separately for each processor, which is impossible in this approach.

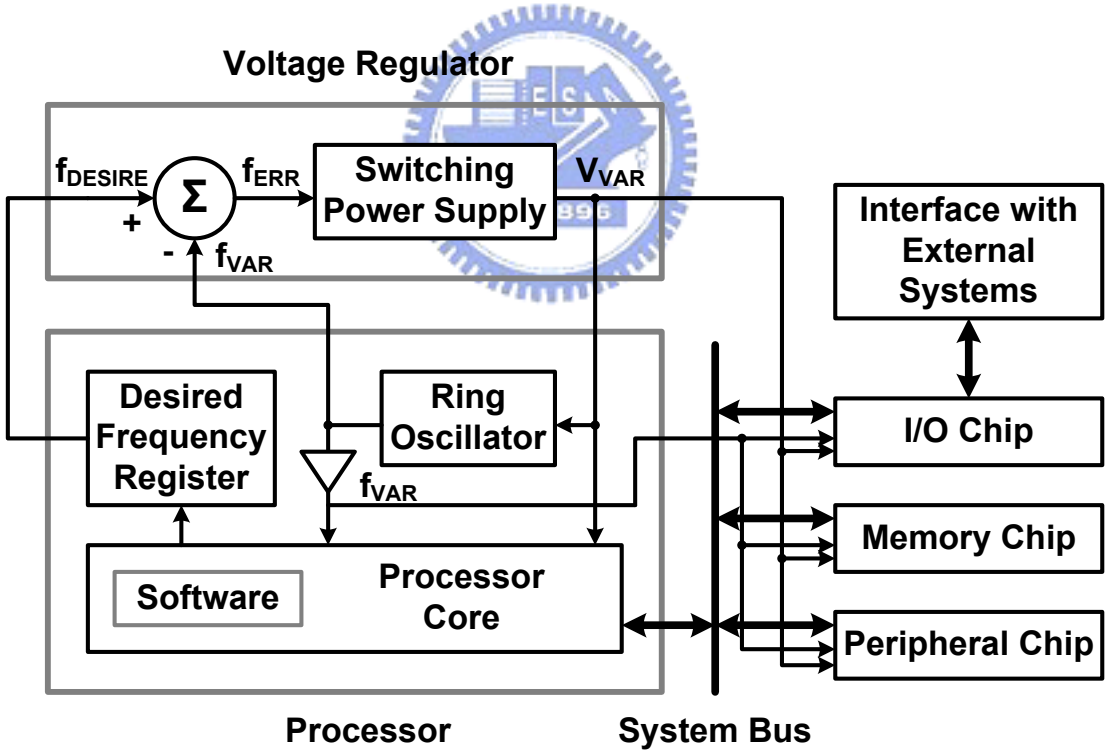

Figure 2.5 Conventional hardware-feedback DVS system architecture.

These problems can be solved by the software-feedback approach, as shown in Figure 2.6 [7]. Desired supply voltage is looked up from the device driver. The device driver has two lookup tables: one for frequency-voltage relationship of the chip, and

the other for transition delay to change clock frequency and supply voltage. These lookup tables are programmable and established by physical measuring results. Therefore, this software-feedback approach can prevent the conventional DVS system's problems.

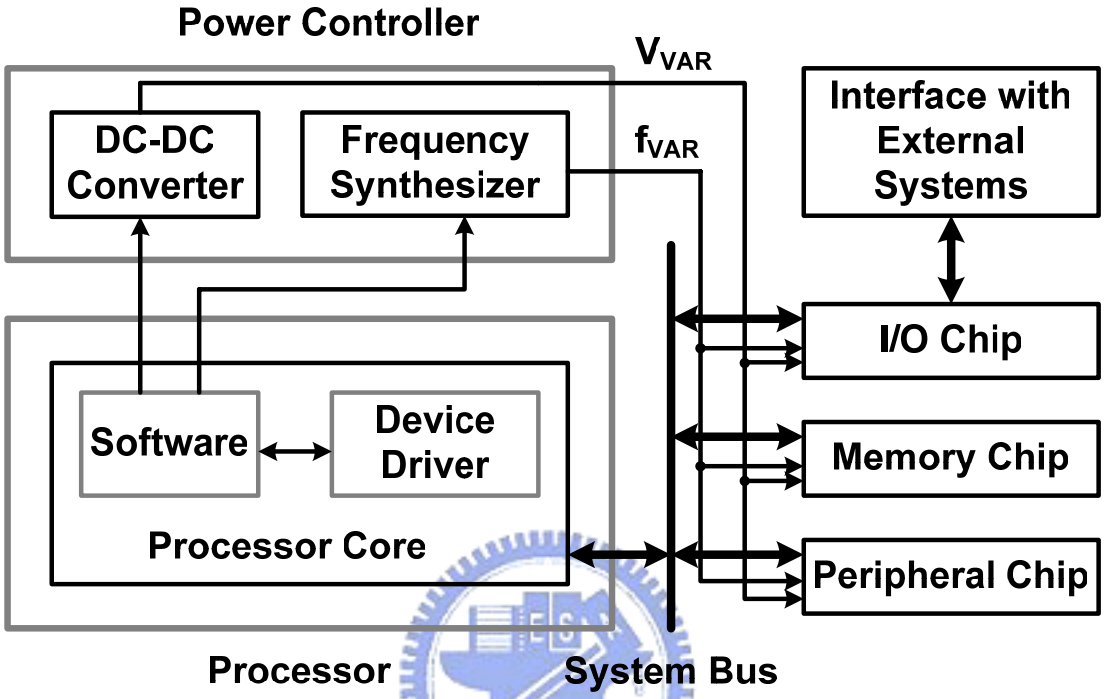

Figure 2.6 Improved software-feedback DVS system architecture.

The concept of software feedback is inspiring and essential in our comprehensive thermal-aware power management design. This concept can be applied to other low power techniques, so that these techniques can be controlled in the same way. Related issues will be discussed in Chapter 3.

2.2.2 Clock Gating

Clock gating is another efficient technique to reduce dynamic power. Clock switching power is a major component of total dynamic power because clock tree represents a very large load and clock signal switches all the time. If we can control and stop clock switching when the system is idle or data remain unchanged, thus unnecessary clock switching power can be saved.

Shown in Figure 2.7 is an illustration of the clock gating technique. The *clk_enable* signal is commonly set to high voltage level (logic 1) during normal operation. If the system is idle or data remain unchanged, then *clk_enable* will be set to low voltage level (logic 0) to stop local gated clock switching. An example of the operating timing diagram is shown in Figure 2.8. Finally, either AND gate or NOR gate can be used to generate gated clock, depending on which trigger type (positive-edge trigger or negative-edge trigger) of sequential elements is selected.

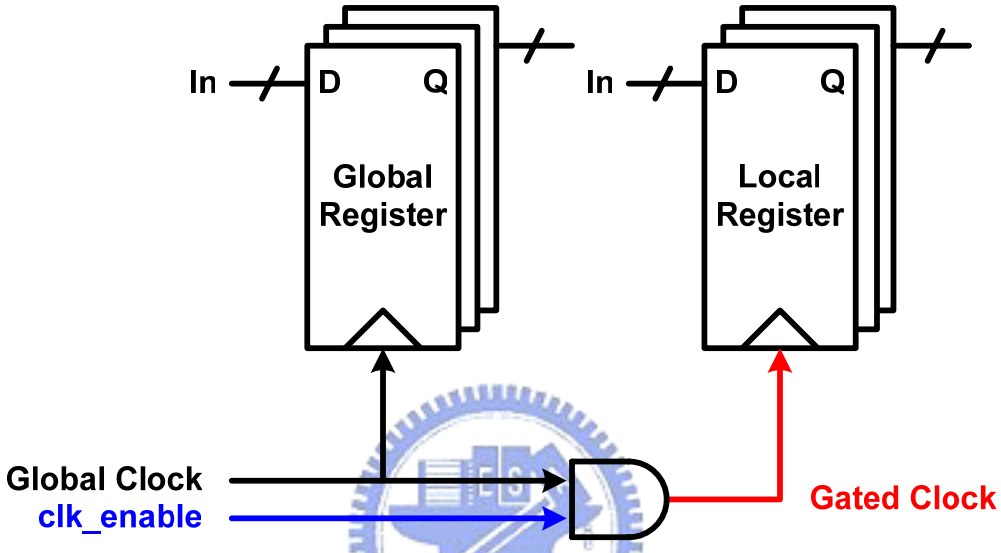

Figure 2.7 Illustration of clock gating.

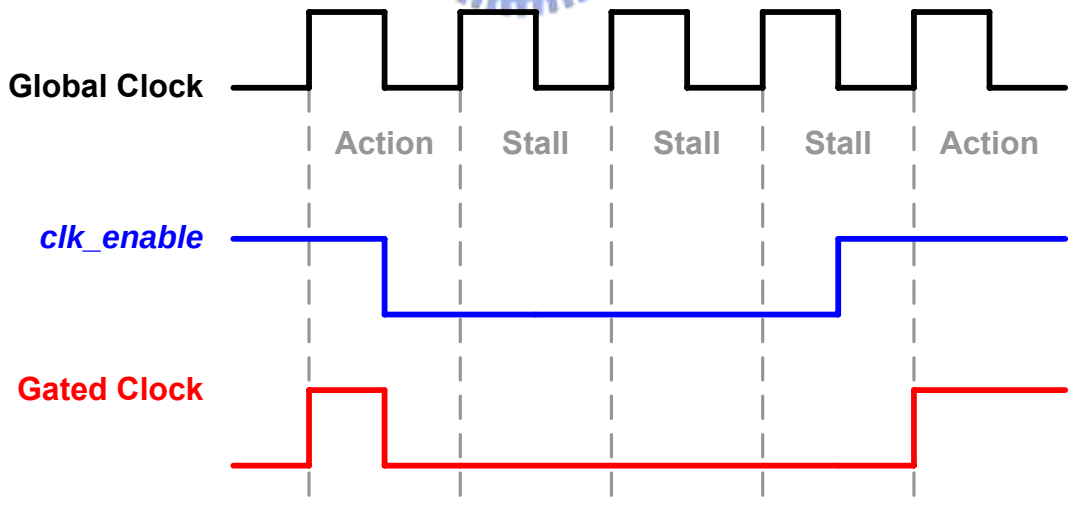

Figure 2.8 Operating timing diagram of clock gating, which is the positive-edge triggered case.

2.2.3 Power Gating

Power gating is the most efficient way to reduce leakage power. Figure 2.9 is an illustration of this technique [5]. In the active mode, the *ctrl* signal remains high, and the power gating devices act as minor power-rail resistance which does not affect the correctness of the function. In the standby mode, however, *ctrl* is set to low. The power gating devices act as open switches that cut off the direct connection between power supply and internal circuits, so as to reduce leakage current [15].

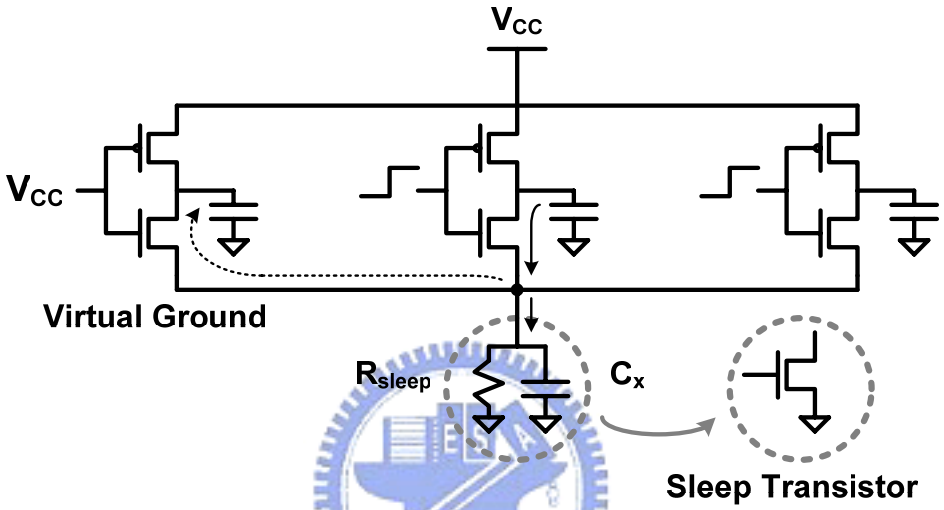

Figure 2.9 Illustration of power gating.

Actually, sleep transistors, as power gating devices, can be PMOS (header-type) or NMOS (footer-type), as shown in Figure 2.10 [15]. Because PMOS sleep transistors require larger silicon area to be capable of sourcing the maximum instantaneous current in the active mode, NMOS is commonly chosen to be power gating devices.

Besides, if sleep transistors are turned off more strongly, further leakage reduction can be achieved. Figure 2.11 shows the block diagram of two enhanced cut-off switch schemes having multi-threshold-voltage CMOS (MTCMOS) and variable-threshold-voltage (VTCMOS) [17]. However, using higher threshold voltage transistors as sleep transistors requires extra mask cost for foundry and larger silicon area to be capable of sinking the maximum instantaneous current in the active mode. Therefore, using single threshold voltage transistors as sleep transistors only or utilizing reverse body bias technique to adjust threshold voltage of sleep transistors are more recommendable.

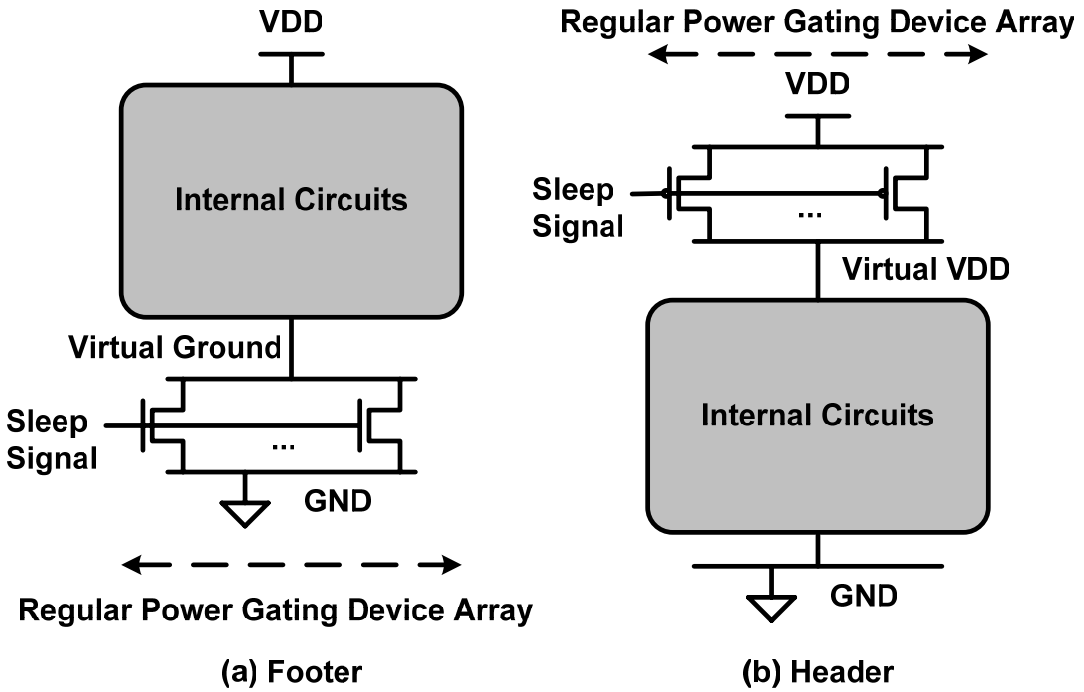

Figure 2.10 (a) NMOS footer-type sleep transistors for power gating. (b) PMOS header-type sleep transistors for power gating.

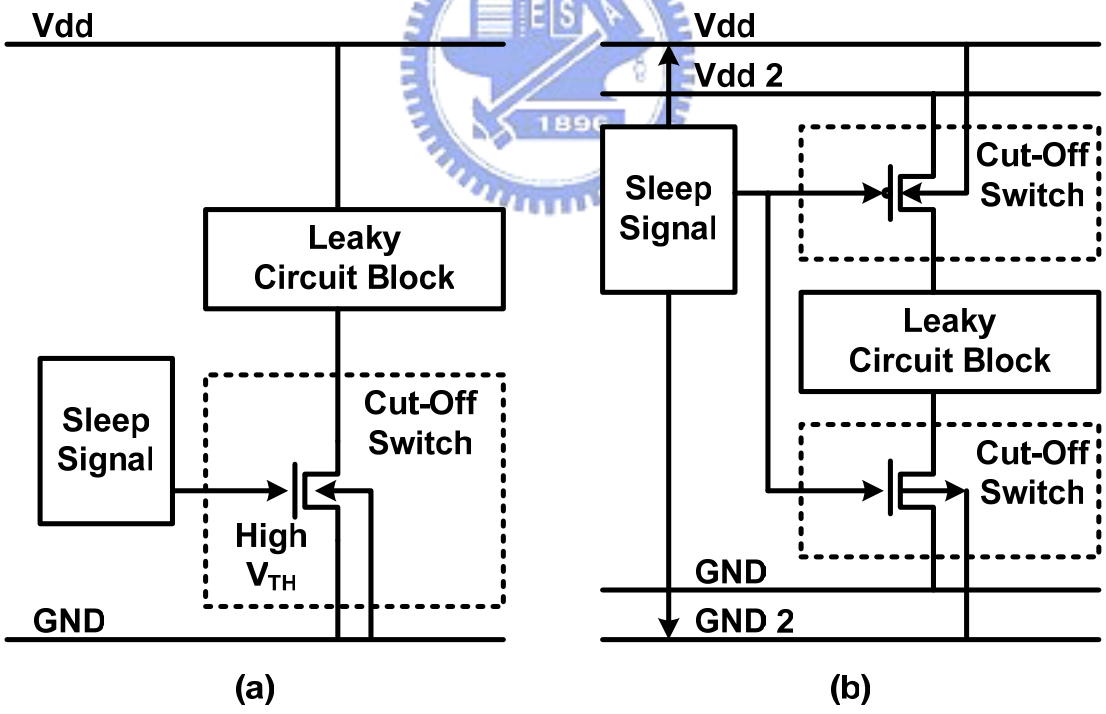

Figure 2.11 Block diagram of two enhanced cut-off schemes. (a) MTCMOS (b) VTCMOS (also known as reverse body bias).

2.2.4 Body Bias

Body bias can be applied for either dynamic or leakage power reduction. In the active mode, forward body bias (FBB) is applied to reduce supply voltage to save dynamic power at the same operating frequency. In the standby mode, reverse body bias (RBB) is applied to raise threshold voltage to save leakage power.

Figure 2.12 shows the V_{th}-hopping scheme that the substrate bias of transistors can be dynamically changed [23]. V_{BSP} and V_{BSN}, the body voltage of PMOS and NMOS, are isolated from V_{DD} and ground, connected to two possible voltage levels to adjust the threshold voltage of the target processor. In the active mode, the *Vt_low_enable* signal is asserted, inducing FBB to lower the threshold voltage, so as to reduce supply voltage to save dynamic power at the same operating frequency. In the standby mode, however, *Vt_high_enable* is asserted, inducing RBB to raise the threshold voltage of the target processor to save leakage power. The *CONT* signal represents the operation mode and decides the values of *Vt_low_enable* and *Vt_high_enable*. It can also be used to control the discrete frequency levels of the target processor.

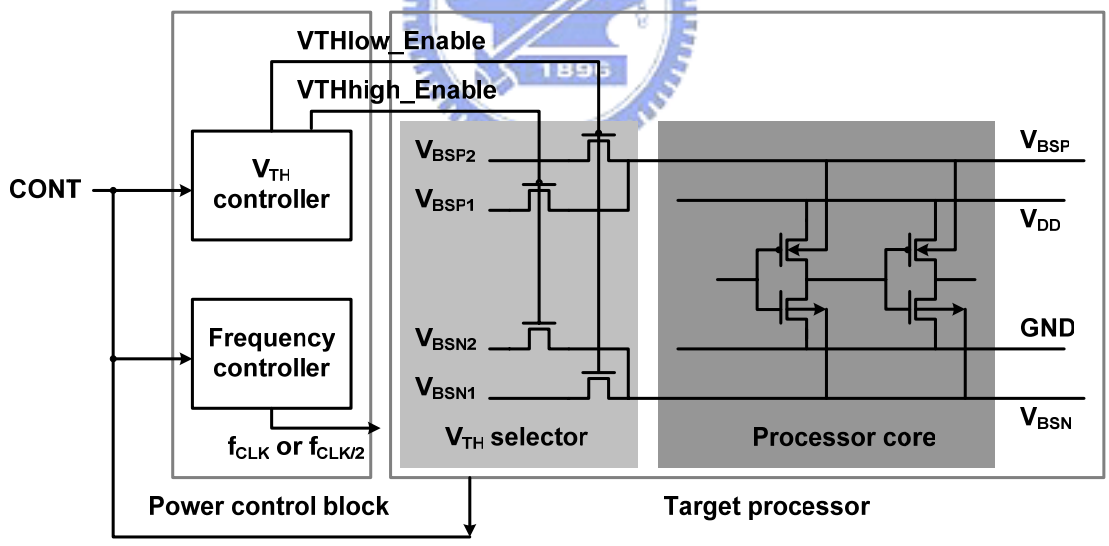

Figure 2.12 Schematic diagram of V_{th}-hopping.

As mentioned before, body bias can be combined with power gating technique to provide further leakage power savings. Shown in Figure 2.13 is a design example which applies these two techniques [6]. On the one hand, FBB can compensate part of frequency degradation due to power gating insertion in the active mode. On the other

hand, RBB can turn off sleep transistors more strongly to further reduce leakage power in the standby mode, which is also known as the VTCMOS cut-off scheme of power gating.

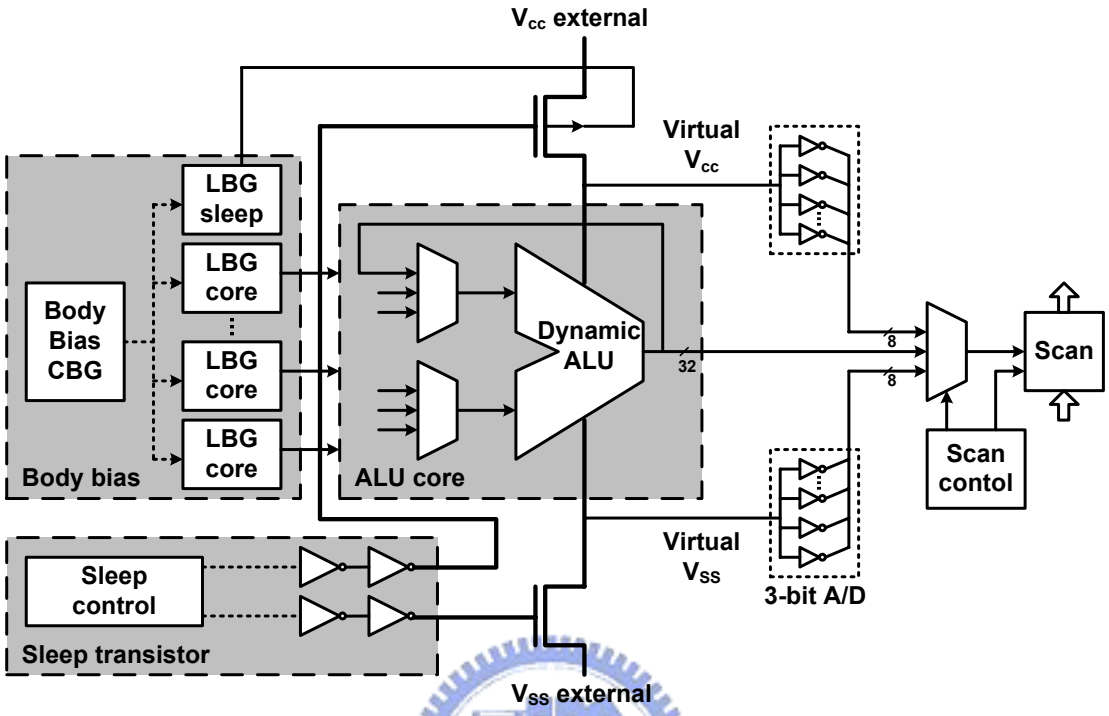

Figure 2.13 Design example utilizing both body bias and power gating.

2.2.5 Voltage Islands

Voltage Islands provide block-level control of the system. In the active mode, local supply voltage can scale down to reduce dynamic power while maintaining the same performance for islands at non-critical paths. In the standby mode, leakage savings are possible when the island is not being used, regardless of whether or not the rest of others are operating. Block-level control makes it easier to spatially optimize the system.

The concept of voltage islands was proposed by IBM in 2002 and originated from the use of dual supply voltages (Dual-V_{DD}) [24, 27]. To begin with, gates on critical paths are assigned to operate at higher V_{DD} and non-critical portions of the circuit operate at lower V_{DD}, reducing the total power consumption without degrading performance. For example, clustered voltage scaling (CVS) is a classic design style with Dual-V_{DD} [24], as shown in Figure 2.14(a).

Afterwards, as shown in Figure 2.14 (b), the concept of voltage islands restores the concept of individual voltage optimization of functional blocks to SoC design [25]. Individual functional blocks of the SoC design can have power characteristics unique from the rest of the design, and can be optimized accordingly. For example, in a system integrating a processor core with on-chip memory and control logic, the processor core may require the highest voltage for performance requirements. However, the on-chip memory and control logic can be operated at low voltage to save dynamic power without compromising system performance. In addition, this macro-based design methodology allows pre-designed standard components to be reused.

Shown in Figure 2.14 (c) is so-called generic voltage islands [26], a new design style with Dual-V_{DD}. In the case of designs that are highly performance critical as well as severely power constrained, it is useful to have a finer grained control over supply voltages or even within a macro of the chip. However, it needs advanced synthesis and placement EDA tools to deal with relative timing issues, and optimization flow is not supported by available commercial EDA tools.

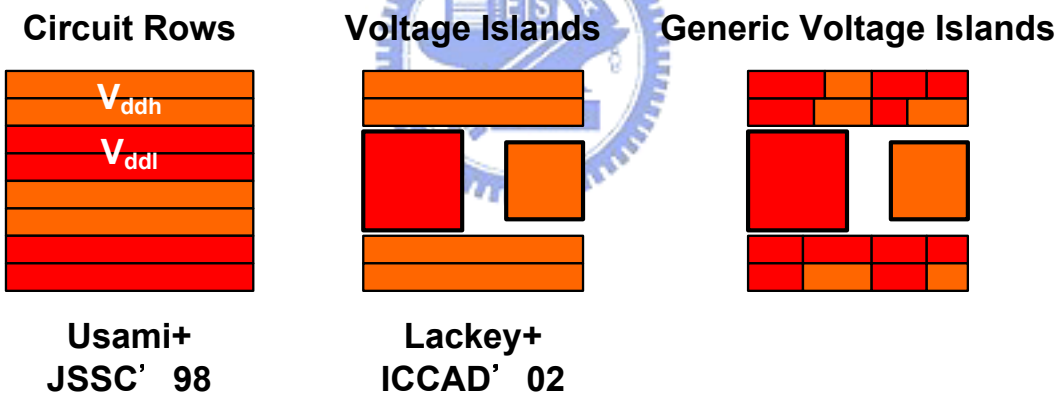

Figure 2.14 Design styles with dual supply voltages. (a) CVS (b) voltage islands (c) generic voltage islands.

Voltage islands itself is a static solution for low power criteria, but it can be combined with power gating to perform dynamic block-level control of leakage power. It can also be combined with DVS to dynamically and spatially reduce active power, since the power supply of each voltage island is separated and isolated.

2.2.6 Summary

DVS provides overall power reduction by dynamically scaling down the supply voltage. Clock gating saves clock switching power when the system is idle or data remain unchanged by stopping local gated clock switching. These two techniques mainly focus on dynamic power saving. On the contrary, power gating efficiently suppresses leakage power in the standby mode by cutting off the direct connection between power supply and internal circuits.

Body bias has two ways to use: FBB and RBB. Forward body bias can reduce dynamic power at the same target frequency or increase speed with the same supply voltage, whereas reverse body bias reduces leakage power by raising threshold voltages of NMOS and PMOS transistors. Moreover, voltage islands provide block-level optimization and spatial flexibility for control mechanism. Table 2.1 is the summary.

Table 2.1 Summary of circuit-level low power techniques.

Technology	Time domain	Space domain	Dynamic power saving	Leakage power saving
Dynamic Voltage Scaling	Yes		Yes	
Clock Gating	Yes		Yes	
Power Gating	Yes			Yes
Forward Body Bias	Yes		Yes	
Reverse Body Bias	Yes			Yes
Voltage Islands		Yes	Yes	Yes

2.3 Enhanced Intel SpeedStep® Technology

In this section, the Enhanced Intel SpeedStep® Technology [28], a state-of-the-art example of power management design and widely used in the Intel processor family, is introduced. The old version of this technology enables real-time dynamic switching of the voltage and frequency between two performance modes based on processor demand. However, the Enhanced Intel SpeedStep® Technology enables the processor to switch between multiple frequency and voltage points instead of two. Switching between states is software controlled.

Following are the key features of the Enhanced Intel SpeedStep® technology [28]:

- Multiple voltage/frequency operating points provide optimal performance at the lowest power.
- Voltage/Frequency selection is software controlled by writing to processor model specific registers (MSRs) thus eliminating chipset dependency.
  - If the target frequency is higher than the current frequency, Vcc is ramped up by placing a new value on the VID pins and the PLL then locks to the new frequency.
  - If the target frequency is lower than the current frequency, the PLL locks to the new frequency and the Vcc is changed through the VID pin mechanism.
  - Software transitions are accepted at any time. If a previous transition is in progress, the new transition is deferred until its completion.
- Low transition latency and large number of transitions possible per second.
  - Processor core (including L2 cache) is unavailable for up to 10μs during the frequency transition.
- Improved Intel® Thermal Monitor mode.
  - When the on-die thermal sensor indicates that the die temperature is too high, the processor can automatically perform a transition to a lower frequency/voltage specified in a software programmable MSR.
  - The processor waits for a fixed time period. If the die temperature is down to acceptable levels, an up transition to the previous frequency/voltage point occurs.
  - An interrupt is generated for the up and down Intel Thermal Monitor transitions enabling better system level thermal management.

The Enhanced Intel SpeedStep® Technology used a DVS based mechanism, changing both the voltage as well as the frequency at run time, to adjust the power consumption of the system to the thermal conditions and to maximize the battery life [8]. An example of a DVS cycle is described in Figure 2.15 [8]. Initially, the processor is operating at high V_{DD} and high frequency generating high power. When the on-die thermal sensor indicates that the die temperature is too high, or the workload of the processor is lower than its capability, then the processor halts the execution and locks the phase-locked loop (PLL) at a new frequency. Once finished, the V_{DD} starts changing to a new lower value. The full power saving is achieved once the voltage transition is finished. A frequency transition up is done in reverse order to ensure normal work of the processor. Finally, the target frequency and voltage points are programmable by BIOS or operating system. Table 2.2 is an example of the Enhanced Intel SpeedStep® Technology operating point for Intel Pentium® M Processor [28].

Besides the Enhanced Intel SpeedStep® Technology, the Intel Pentium® M processor supports the AutoHALT Power-Down, Stop Grant, HALT/Grant Snoop, Sleep, DeepSleep, and Deeper Sleep states for optimal power management. Figure 2.16 is its low power state machine and Table 2.3 explains definitions of these states in detail [28].

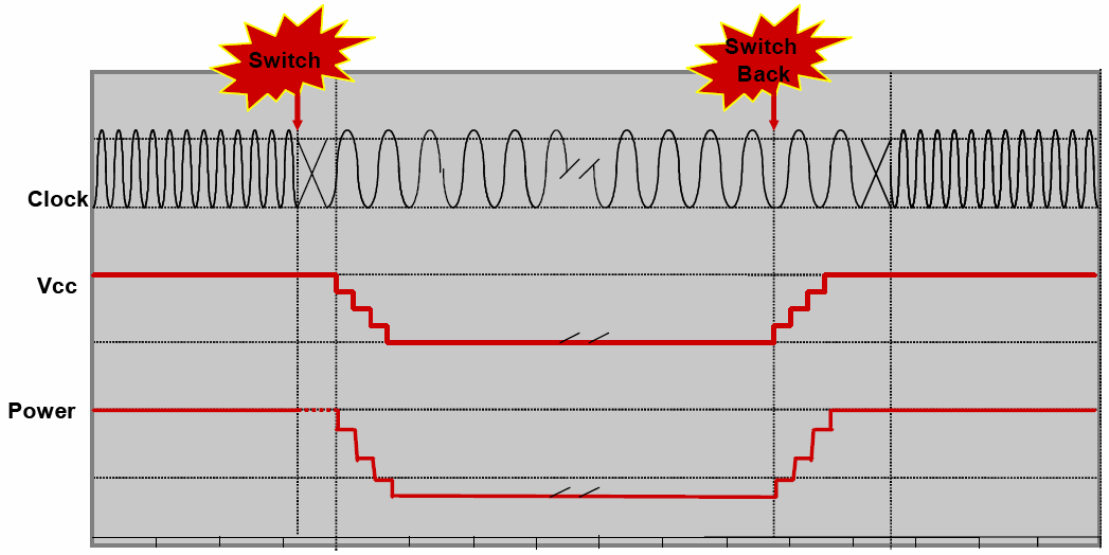

Figure 2.15 Dynamic voltage scaling cycle.

Table 2.2 Voltage specifications of Intel Pentium® M Processor. The Intel Pentium® M processor uses six voltage identification pins, VID[5:0], to support automatic selection of power supply voltages.

Processor	Frequency (MHz)	V_{DD} (V)	VID
Intel Pentium® M Processor 765	2.1	1.340	010111
	1.8	1.276	011011
	1.6	1.228	011110
	1.4	1.180	100001
	1.2	1.132	100100
	1.0	1.084	100111
	0.8	1.036	101010
	0.6	0.988	101101
Intel Pentium® M Processor 755	2.1	1.340	010111
	1.8	1.292	011010
	1.6	1.244	011101
	1.4	1.196	100000
	1.2	1.148	100011
	1.0	1.100	100110
	0.8	1.052	101001
	0.6	0.988	101101

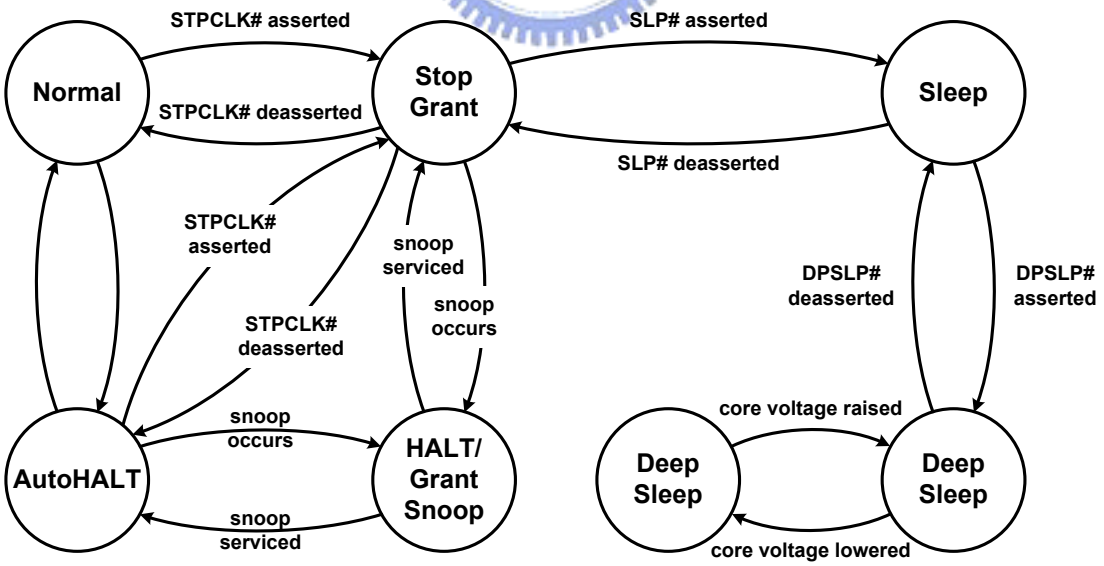

Figure 2.16 Low power state machine of Intel Pentium® M processor.

Table 2.3 State and control signal definitions of Intel Pentium® M processor.

Normal State	This is the normal operating state for the processor.
AutoHALT Power-Down State	AutoHALT Power-Down is a low-power state entered when the processor executes the HALT instruction. While in AutoHALT Power-Down state, the processor will process bus snoops and interrupts.
Stop-Grant State	When the STPCLK# pin is asserted, the Stop-Grant state of the processor is entered 20 bus clocks after the response phase of the processor-issued Stop Grant Acknowledge special bus cycle. While in Stop-Grant state, the processor will process snoops on the FSB and it will latch interrupts delivered on the FSB.
HALT/Grant Snoop State	The processor responds to snoop or interrupt transactions on the FSB while in Stop-Grant state or in AutoHALT Power-Down state. During a snoop or interrupt transaction, the processor enters the HALT/Grant Snoop state.
Sleep State	A low power state in which the processor maintains its context, maintains the PLL, and has stopped all internal clocks. The Sleep state can be entered only from Stop-Grant state. In the Sleep state, the processor is incapable of responding to snoop transactions or latching interrupt signals. No transitions or assertions of signals (with the exception of SLP#, DPSLP# or RESET#) are allowed on the FSB while the processor is in Sleep state.
Deep Sleep State	Deep Sleep state is a very low power state the processor can enter while maintaining context. BCLK may be stopped during the Deep Sleep state for additional platform level power savings. When the processor is in Deep Sleep state, it will not respond to interrupts or snoop transactions.
Deeper Sleep State	This state is functionally identical to the Deep Sleep state but at a lower core voltage.
STPCLK#	*STPCLK# (Stop Clock), when asserted, causes the processor to enter a low power Stop-Grant state. The processor issues a Stop-Grant Acknowledge transaction, and stops providing internal clock signals to all processor core units except the FSB and APIC units. The processor continues to snoop bus transactions and service interrupts while in Stop-Grant state. When STPCLK# is deasserted, the processor restarts its internal clock to all units and*

	resumes execution. The assertion of STPCLK# has no effect on the bus clock; STPCLK# is an asynchronous input.
SLP#	*SLP# (Sleep), when asserted in Stop-Grant state, causes the processor to enter the Sleep state. During Sleep state, the processor stops providing internal clock signals to all units, leaving only the Phase-Locked Loop (PLL) still operating. Processors in this state will not recognize snoops or interrupts. The processor will recognize only assertion of the RESET# signal, deassertion of SLP#, and removal of the BCLK input while in Sleep state. If SLP# is deasserted, the processor exits Sleep state and returns to Stop-Grant state, restarting its internal clock signals to the bus and processor core units. If DPSLP# is asserted while in the Sleep state, the processor will exit the Sleep state and transition to the Deep Sleep state.*
DPSLP#	*DPSLP# when asserted on the platform causes the processor to transition from the Sleep State to the Deep Sleep state. In order to return to the Sleep state, DPSLP# must be deasserted.*

In summary, the Intel Pentium® M processor applies DVS in the active mode, which means utilizing the Enhanced Intel SpeedStep® Technology in Normal state; in the standby mode, namely in Stop Grant, Sleep, DeepSleep, and Deeper Sleep states, clock gating is used to reduce clock switching power. Although the Enhanced Intel SpeedStep® Technology has taken die temperature limit into consideration, however, it doesn't deal with the problem of performance coherence between functional blocks due to temperature gradient. Besides, leakage power is barely emphasized in the Intel Pentium® M processor design. In our CTAPM system, performance coherence and leakage power suppression are especially addressed.

2.4 Summary

Power consumption is categorized to dynamic, short circuit, and static leakage power in modern VLSI circuitry. Several low power techniques have been used to deal with different power sources. DVS, clock gating and forward body bias can reduce dynamic power in different ways, while power gating and reverse body bias are aimed at leakage power suppression. The Enhanced Intel SpeedStep® Technology is a state-of-the-art example of power management design. However, just like other developments, the block-level control for multiple power sources and the awareness of thermal impact are not achieved yet. Therefore, we propose a new power management design in this thesis.

Contributions and features of our CTAPM design are listed in the following:

1. The word *"Comprehensive"* means that this design adopts several low power techniques mentioned above to dynamically deal with both dynamic power and static power consumption. Some researches have used more than one technique in their designs [3, 6, 22, 29] but only one power source was discussed or both of them were controlled by static solutions for the most cases.
2. The phrase *"Thermal-Aware"* means that the problem of performance coherence between functional blocks due to temperature gradient is ensured in our power management design. This problem was not revealed since previous works mainly focused on the full-chip control. Digital circuit designers are also hard to be aware of it during hardware-language coding. However, the block-level control provides more flexibility and efficiency. Therefore, this problem has to be solved to maintain system stability.
3. Because the software-feedback approach is chosen, this CTAPM design is programmable to handle both die-to-die process variation and within-die process variation. The yield can hence be improved.
4. Not only power analysis, design trade-offs between different techniques and technologies are also discussed in this thesis. The analysis flow is still the same for different technologies, while the selection of low-power techniques will be case by case.

In order to implement these design concepts, the CTAPM unit architecture and related issues are introduced in Chapter 3. The proposed full-version architecture design can be modified or simplified while maintaining the same functionality. More details and the experimental results are described in Chapter 4.

Chapter 3

Comprehensive Thermal-Aware Power Management

The comprehensive thermal-aware power management (CTAPM) system is proposed in this chapter. In power phase, the CTAPM takes care of both dynamic power and leakage power; in thermal phase, the asynchronous problem between functional blocks due to temperature gradient is handled. Related issues like the architecture map, control state machine, and look-up table design are also introduced in this chapter. In addition, the test vehicle and analysis flow are described in Section 3.2 and 3.3.

3.1 System Design of CTAPM

As implied by the name, the function of CTAPM is to control power consumption of functional blocks of the chip, and take thermal issues into account at the same time. As shown in Figure 3.1, the CTAPM receives temperature and leakage information from the thermal sensor and job monitor of each target circuit. It also accepts system orders from the embedded operation system (OS) or specific scheduling software, such as system state and desired performance. Besides, there should be an application interface for users to input commands externally. Afterward, the CTAPM decides and sends control parameters to peripheral circuits to produce desired supply voltage and clock frequency, .etc. Finally, the target circuits will act in the desired performance and power envelopes.

The overall architecture map, control state machine and look-up table design are major parts in system design of CTAPM. The following of this section will describe these three topics in sequence.

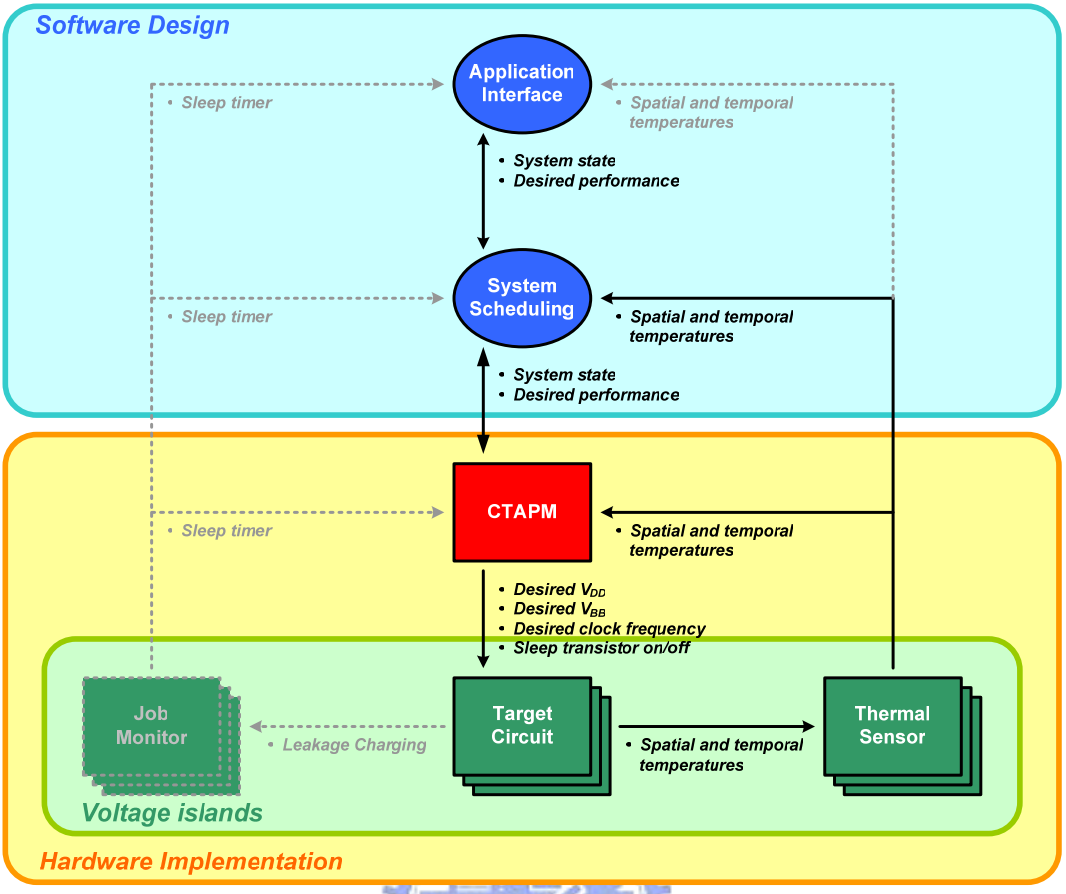

Figure 3.1 Illustration for the functionality of CTAPM. The dotted parts are not covered in this thesis.

3.1.1 Architecture Map

The proposed CTAPM unit architecture is shown in Figure 3.2. Techniques including dynamic voltage scaling (DVS), clock gating, power gating, and body bias are combined to control system power and adapt to thermal impact. DVS and clock gating can effectively reduce dynamic power consumption, while power gating reduces static leakage power. Forward body bias can increase speed, while reverse body bias reduces power as well. Individual details of these low power techniques have been described in Chapter 2. Besides, these techniques can be applied at the full-chip level, or they can be applied on a finer block-level granularity. As the result, the concept of voltage islands is also included in power management design.

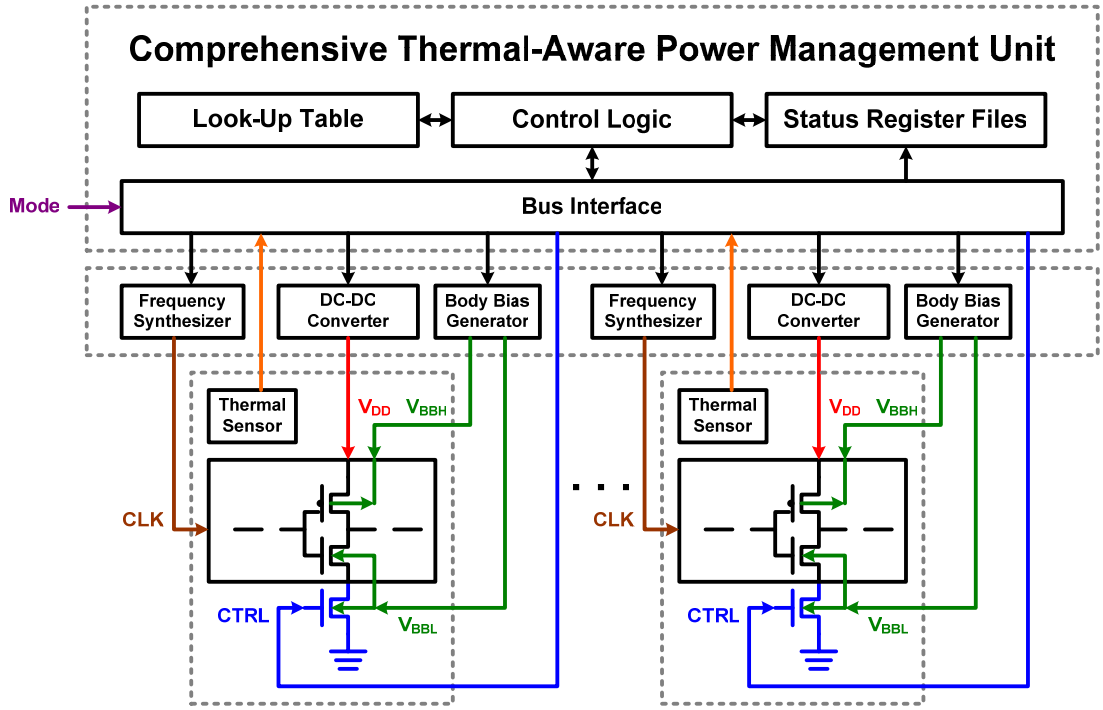

Figure 3.2 Architecture of proposed comprehensive thermal-aware power management unit. Peripheral circuits, such as DC-DC converter and temperature sensor, are required.

The CTAPM unit contains control logic circuits, look-up tables, status register files, and a bus interface. Look-up table design is the most important design parameter of the CTAPM unit, which will be discussed later. The control logic chooses optimized values of parameters based on look-up tables and system status. Status register files record system information from software and hardware. Finally, the bus interface is used for communication.

Peripheral circuits are needed to implement these low power techniques. DC-DC converter is used to dynamically provide the desired supply voltage of the target circuit. Frequency synthesizer is applied to change the operating frequency between different performance levels. If the functional block is in the standby mode, the gated local clock will be paused to save clock switching power. Body bias generator provides the desired body voltages of core transistors. If forward body bias (FBB) is expected in the active mode, then the body voltage of PMOS (V_{BBH}) will be below V_{DD} and the body voltage of NMOS (V_{BBL}) will be above ground. Otherwise, if reverse body bias (RBB) is preferred in the standby mode, then V_{BBH} will be above V_{DD} and V_{BBL} will be below ground.

In addition, the thermal sensor in each voltage island is needed to report the operating temperature to CTAPM. Even at the same performance level, the supply voltage and body bias could vary according to temperature variation to make sure the consistence of the speed.

3.1.2 Control State Machine

Figure 3.3 is the low power state machine of CTAPM. Compared with that of Intel Pentium® M processor [28], Figure 3.3 is simpler because some functional states handling exceptions are omitted, but stronger in power phase due to leakage control mechanism, which is important especially for future technology generations and SoC designs. Besides, since CTAPM provides block-level control, each functional island should have its own state rather than unity.

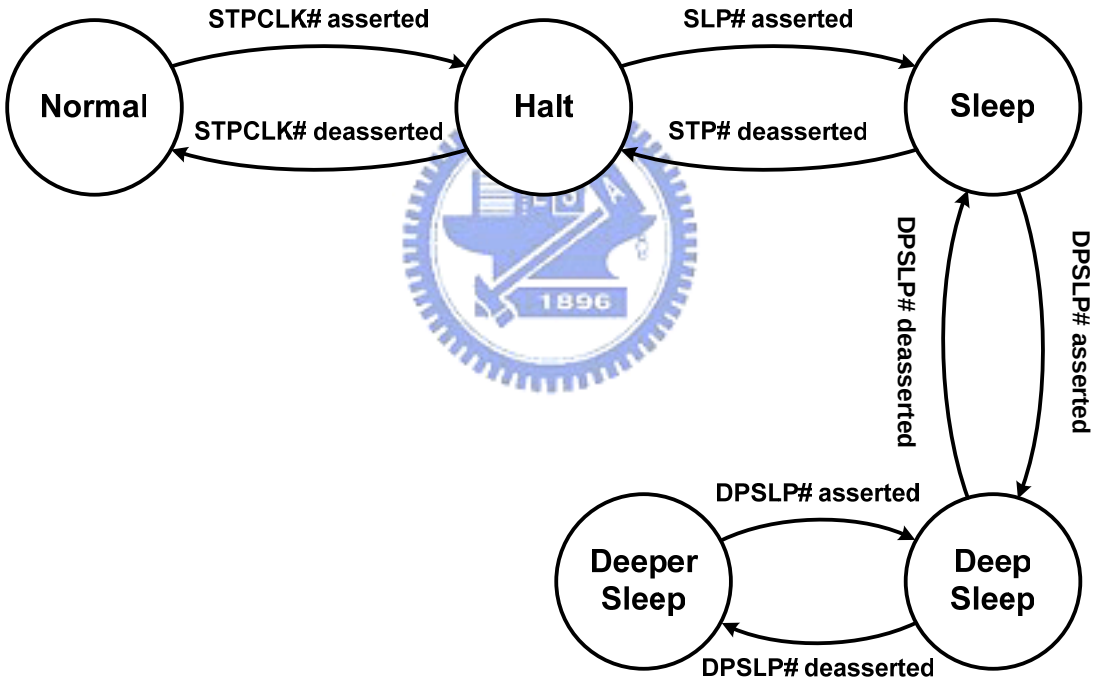

Figure 3.3 Low power state machine of CTAPM.

From Figure 3.3, Normal State is identical to the active mode in previous used terminology, while other four states can be seen as different standby modes in the mean time. Normal State is the common operating state for the target functional block. When the island is about to execute few HALT instructions (this condition can be predicted or scheduled early in compile time), it will enter HALT State to pause the gated local clock to save clock switching power. If the functional block is idle for a

longer time, then it will enter Sleep State, Deep Sleep State, and Deeper Sleep State in order. Corresponding low power solutions, like power gating and reverse body bias, will be applied sequentially, too. Timing specifications need to be investigated and decided by measurement or simulations. However, it is not covered in this thesis. Table 3.1 is the summary.

Table 3.1 State definitions and functionalities of CTAPM.

Normal State	Definition	Normal operating state for the target functional island
	Functionality	*DVS to make performance and power scalable* *FBB to ensure performance coherence between blocks*
HALT State	Definition	Entered when few consecutive HALT instructions executed
	Functionality	*Gated clock off to reduce local clock switching power*
Sleep State	Definition	Entered after the island idle for a while
	Functionality	*Sleep transistor off to reduce leakage power*
Deep Sleep State	Definition	Entered after the island in Sleep State for a while
	Functionality	*RBB to reduce leakage power even more*
Deeper Sleep State	Definition	Lowest power state the system can enter
	Functionality	*Lowest V_{DD} and highest RBB to reduce leakage power most*

3.1.3 Look-Up Table Design

As shown in Figure 3.2, the look-up table is among the most important design parameter of the CTAPM unit. It provides frequency-voltage information for software feedback of the system, either in operation mode transitions or toward temperature variation. In this thesis, we focus on the look-up table design for the block-controlling parameters, especially for adjust frequency-voltage relationship under temperature variation. Besides power control mechanism, we utilize DVS and FBB on the temperature axis to meet performance requirements. We hope that power consumption can be controlled by the CTAPM system but also deal with thermal impact to performance coherence in the meantime.

Figure 3.4 illustrates the construction of look-up tables. Each functional block should have its own look-up table in the CTAPM unit. Normal State is divided into five performance levels, followed by HALT State, Sleep State, Deep Sleep State, and Deeper Sleep State. The temperature axis is divided into six steps. *VDDID* and *VBBID* represent the optimized supply and body bias voltages, while *CLKSP* and *CKTSP* stand for enabling clock gating and power gating or not.

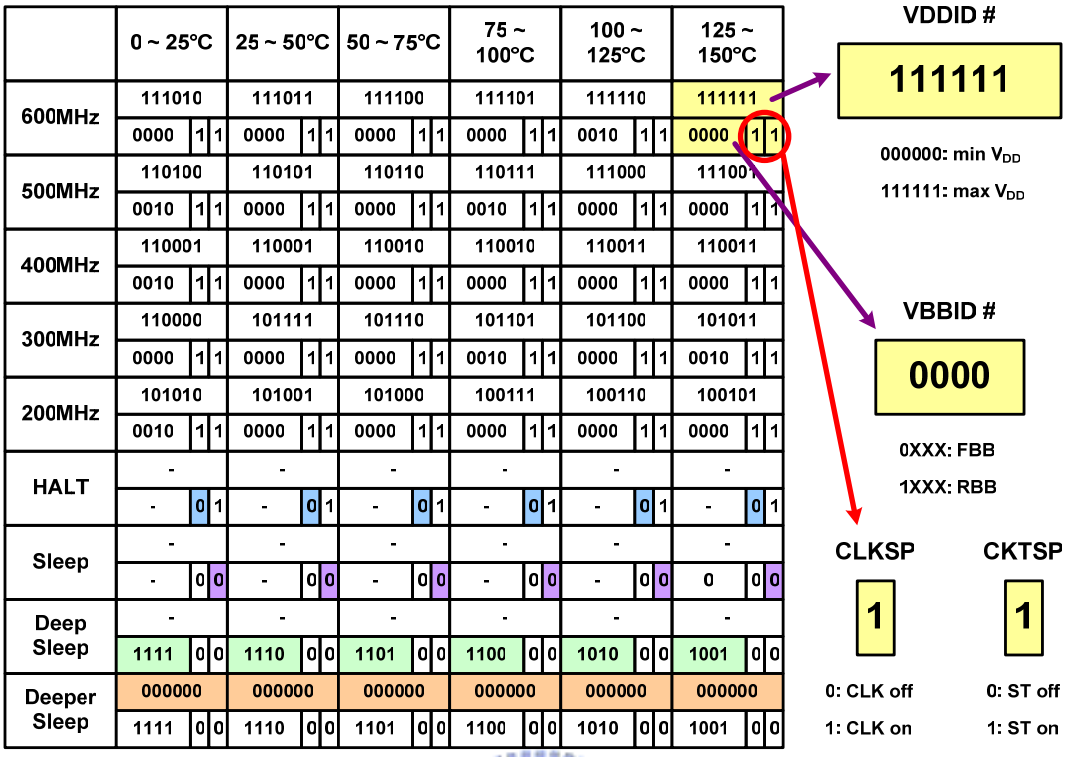

Figure 3.4 Diagram of the look-up table design for CTAPM.

Actually, Figure 3.4 primarily uses DVS to face temperature variance. If FBB is applied dynamically in the small scale of temperature to adjust frequency loss, there will be an extra look-up table in each temperature step for FBB to decide the value, which is like Figure 3.5.

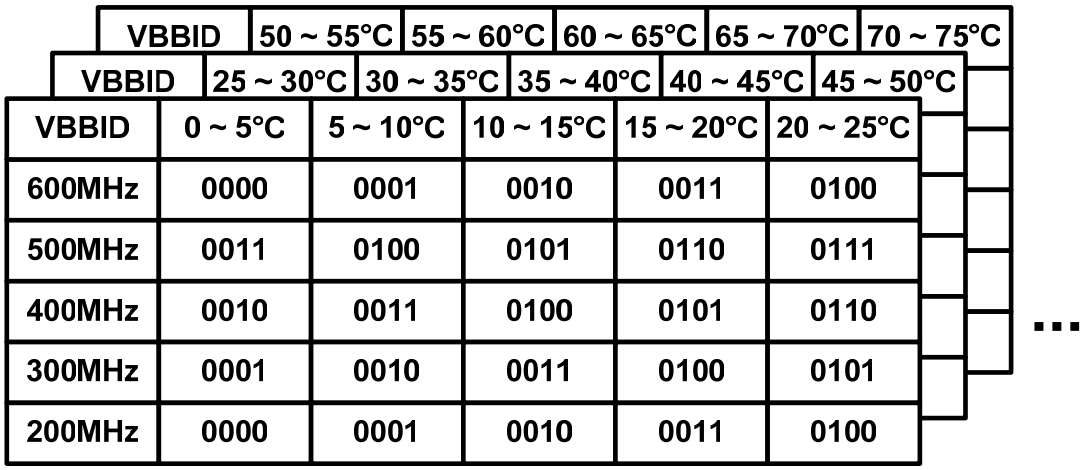

Figure 3.5 Illustration of the FBB look-up tables.

3.2 Test Vehicle

To verify the concepts of CTAPM, a 101-stage ring oscillator is chosen as the test vehicle, as shown in Figure 3.6. The 101-stage ring oscillator is chosen for following reasons: First, the ring oscillator is composed by inverters, which consumes the largest static leakage power because of the least stacking effect. Therefore, this test vehicle matches the trend of leakage growth, and our experimental results are applicable for future generations. Second, the ring oscillator is a common component of digital circuits as the timing emulator, especially in hardware feedback systems. Moreover, the scale of the 101-stage ring oscillator is comparable to a 32-bit adder, meaning that the experimental results are practicable.

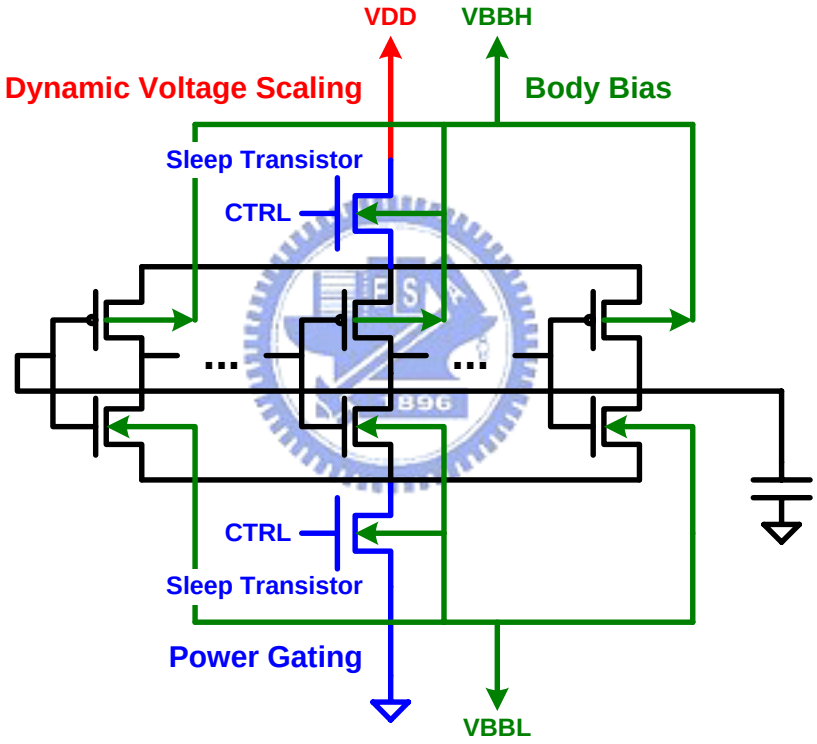

Figure 3.6 A 101-stage ring oscillator is used as the test vehicle. Several power reduction concepts are applied to this circuit.

Figure 3.7 is the layout illustration of our test vehicle. Differing from the conventional layout style, source and body terminals of transistors are separated for DVS and body bias. Besides, sleep transistors are inserted in to the original circuit for power gating. However, if body bias is applied to NMOS transistors, then the triple-well technology (see Figure 3.8) is needed for block-level control to isolate the p-well bias of each voltage island.

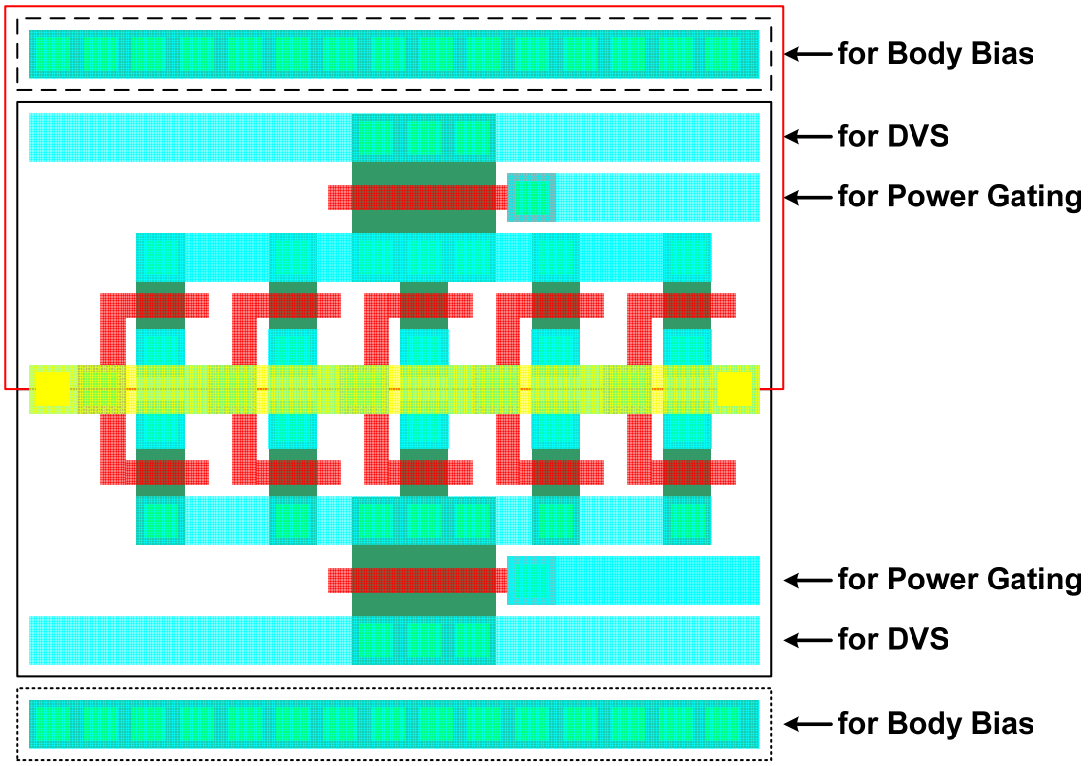

Figure 3.7 Layout illustration of the 101-stage ring oscillator. Source and body terminals of transistors are separated for DVS and body bias.

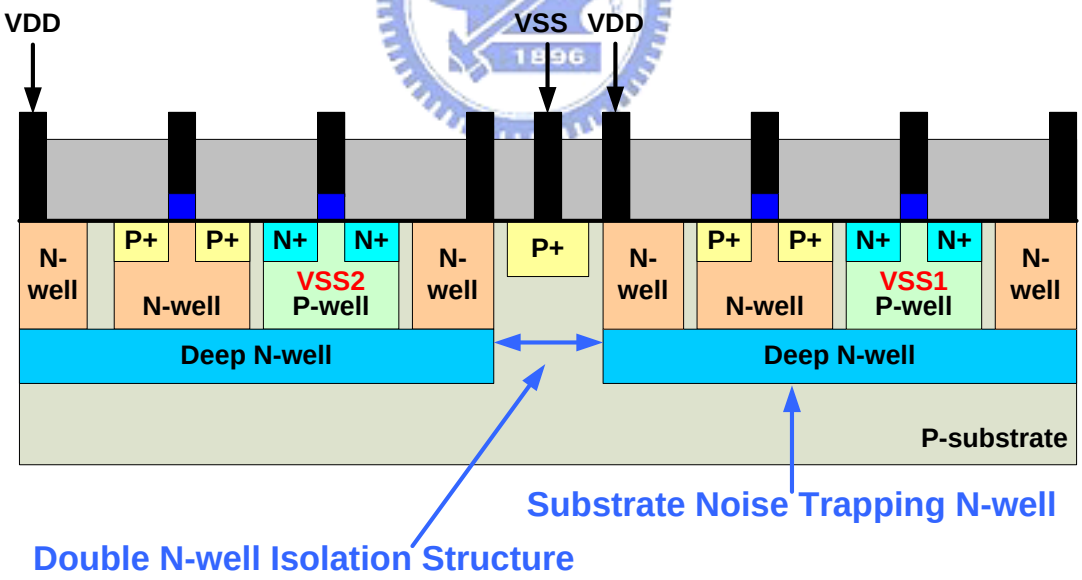

Figure 3.8 Triple-well technology. Deep n-wells are implemented to prevent crosstalk between p-wells of NMOS transistors.

3.3 Analysis Flow

Power gating insertion is done before DVS simulation and body bias simulation, since DVS and body bias only change the supply voltage and body voltages of CMOS, which do not impact the structure of the core logic circuit, but power gating technique does. Sleep transistors are added to the original circuit for power gating insertion, which has a great impact on system performance, noise margin, dynamic power, and static power consumption. As a result, when designing a system integrating several low power techniques such as DVS, body bias, and power gating, the first work is to decide the type and size of sleep transistors. Figure 3.9 is the analyzing flow for experiments.

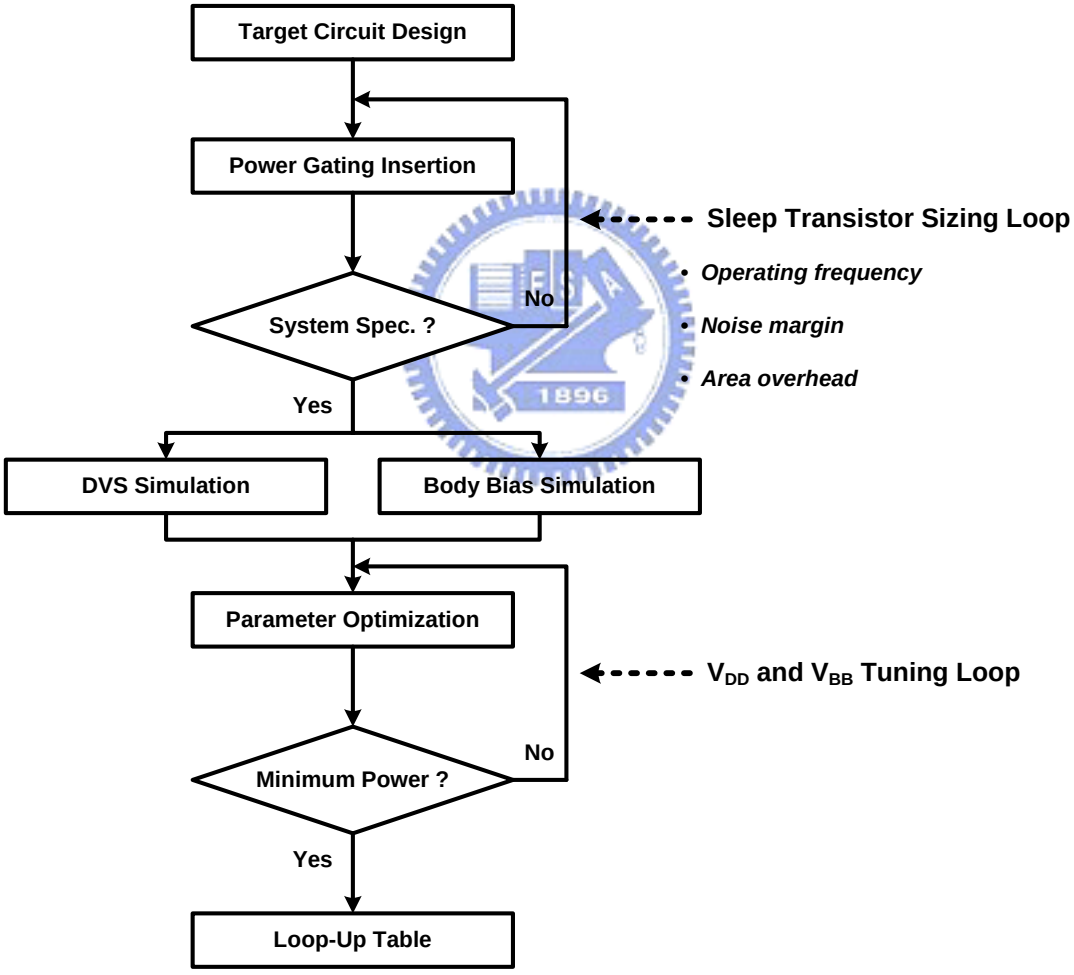

Figure 3.9 Analysis flow for experiments.

After the target circuit design, simulations about power gating insertion will proceed. The type of sleep transistors has to be chosen first, then the sleep transistor sizing loop will be activated. Trade-offs between the operating frequency, noise margin, and area overhead drive this sizing loop. Once these three specifications are acceptable, the size of sleep transistors will have been decided.

DVS and body bias simulations come after power gating insertion. First, basic characteristic simulations will be done to realize the effect and efficiency of these techniques. Afterward, V_{DD} and V_{BB} tuning loop should be activated to optimize design parameters. Each performance level and temperature interval should be corresponding to a pair of optimized V_{DD} and V_{BB} to minimize power consumption. At last, the loop-up table of CTAPM can be built. Detailed simulation results can be seen in Chapter 4.

Chapter 4

Power Analysis

The experimental results of different circuitry techniques utilized in proposed comprehensive thermal-aware power management (CTAPM) system are presented in this chapter. We first begin with power gating analysis in Section 4.1. Section 4.2 discusses characteristics of dynamic voltage scaling (DVS) and construction of the V_{DD} look-up table. Section 4.3 and 4.4 examine the effect and efficiency of reverse body bias (RBB) and forward body bias (FBB). However, the simulation of clock gating is skipped since it is a mature technique. These experimental results are based on TSMC 100nm CMOS technology.

4.1 Power Gating Analysis

In this section, first we observe device leakage characteristics before we use these transistors. Then, we choose NMOS or PMOS as the sleep transistor for our test vehicle under some considerations. Section 4.1.3 shows sleep transistor sizing by trade-offs between operating frequency, dynamic power, and virtual ground rising. The minimum scale of voltage islands is also defined. In the end of this section is the summary.

4.1.1 Device Leakage Characteristics

Before proceeding power gating analysis, characteristics of device leakage should be known first. Figure 4.1 is the device leakage current simulation results based on TSMC 100nm CMOS technology (minimum transistor length is used) at room temperature. We can see that the PMOS device leakage is as a rule proportional to the transistor width; namely, the wider the transistor is the larger leakage current it produces. However, the NMOS device leakage is not that normal. Leakage current grows surprisingly large when the device width is approaching minimum width. There is no appropriate physical explanation yet for this unpredicted phenomenon, but it reveals two important messages: First, since digital circuits are generally composed of minimum size transistors for area constraints, larger than expected leakage current is

produced inevitably to worsen power consumption. Moreover, to earn more effect and efficiency of power gating technique, NMOS approaching minimum size should be avoided to be sleep transistors.

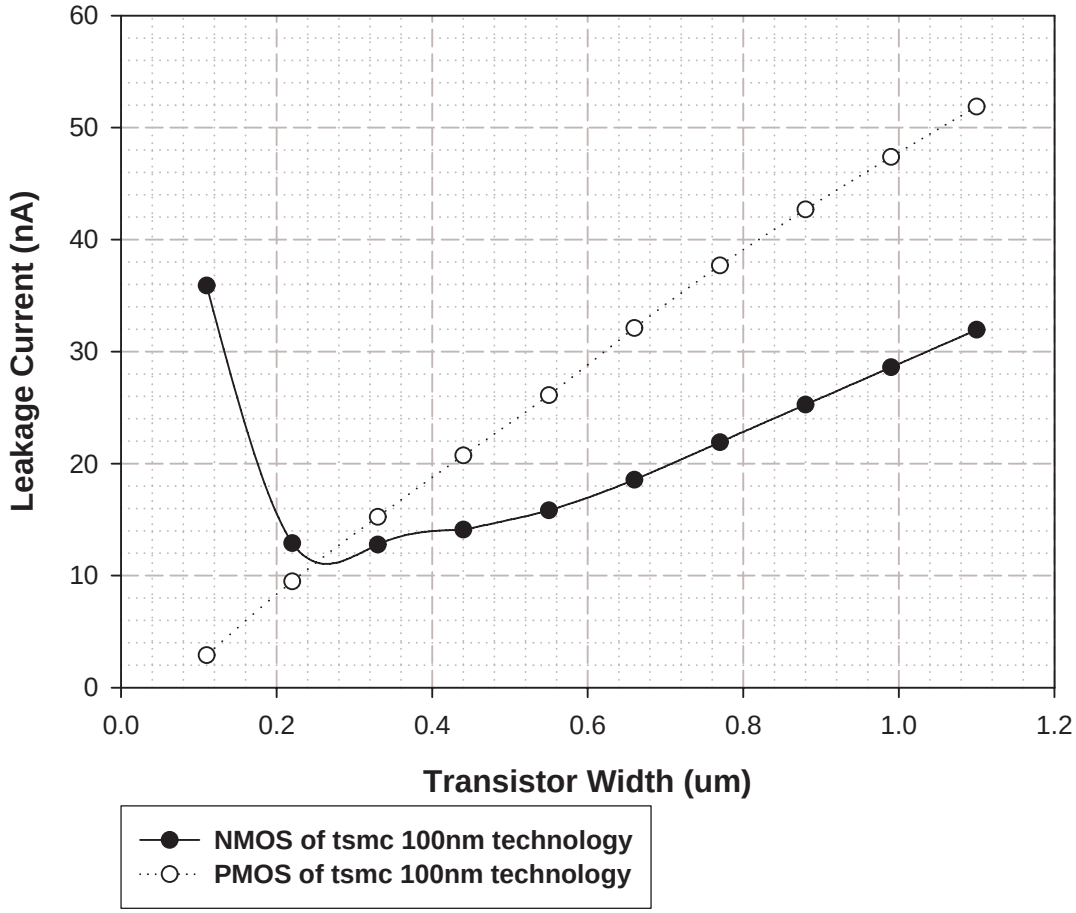

Figure 4.1 Device leakage simulation of TSMC 100nm CMOS technology. Unless otherwise noted, all data used for the figures and tables in this chapter are simulated at room temperature.

4.1.2 Sleep Transistor Type Selection

As mentioned in Chapter 2, sleep transistors for power gating can be PMOS (header-type) or NMOS (footer-type); only one polarity sleep device is actually required to reduce leakage. Because the mobility of electric holes is smaller than electrons, header-type PMOS sleep transistors require more silicon area than NMOS to source the maximum instantaneous current in the active mode. Therefore, NMOS is commonly chosen to be sleep transistors.

Figure 4.2 shows the relationships between operating frequency and sleep transistor width of the 101-stage ring oscillator (described in Chapter 3) under TSMC 100nm CMOS technology. It shows that PMOS sleep transistor leads to more frequency sacrifice than NMOS sleep transistor for the same transistor size. In fact, for example, PMOS sleep transistor needs to be 14X minimum size to compete with 3X minimum size of NMOS sleep transistor for the same frequency sacrifice.

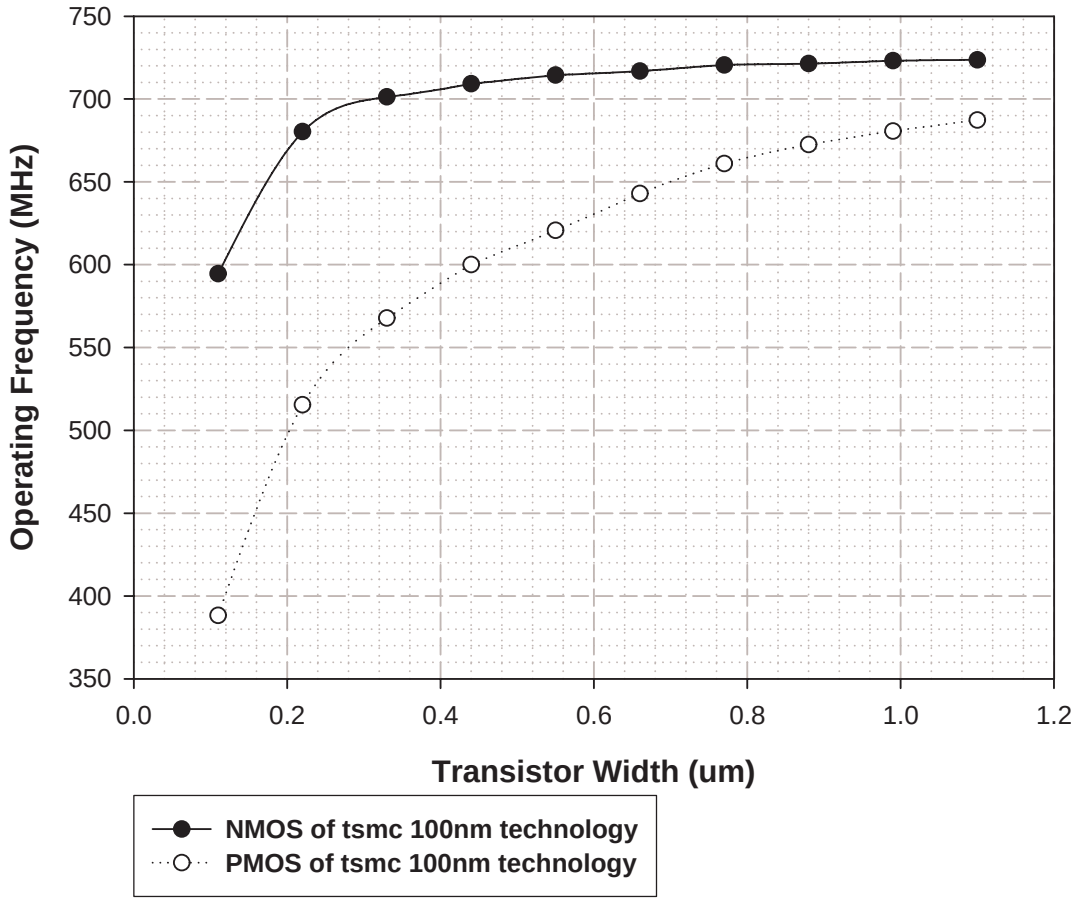

Figure 4.2 Speed comparison of different kinds of the sleep transistor of the ring oscillator under TSMC 100nm CMOS technology. NMOS sleep transistor is more competitive because of less area overhead for the same frequency sacrifice.

Besides area overhead, virtual supply noise and power consumption should be taken into consideration when selecting sleep transistor type. However, there is nearly no difference in choosing which type of sleep transistor about these two issues. Shown in Table 4.1 are frequency/power/noise trade-offs of different sleep transistor

types. Due to equivalent resistance but larger parasitic capacitance, using PMOS as the sleep transistor results in slightly less virtual supply noise and more power consumption than NMOS.

Since these differences are quite small, selecting sleep transistor type should mainly depend on area overhead. Nevertheless, manufacturing issues need to be kept in mind when combining power gating with reverse body bias technique to do the block-level control. More details will be seen in Section 4.5.

Table 4.1 Comparison of different sleep transistor types.

	Case 1		Case 2	
Sleep Transistor Type	9X PMOS	2X NMOS	14X PMOS	3X NMOS
Operating Frequency (MHz)	680.61	680.25	701.75	701.26
Virtual Supply Noise (mV)	52.05	56.46	31.49	35.85
Power Consumption (uW)	35.56	35.32	38.76	37.40

4.1.3 Sleep Transistor Sizing

Sleep transistors sizing is the most important and hardest part of power gating insertion due to frequency loss, area overhead, and other trade-offs. If sized too large, then valuable silicon area would be wasted and leakage power saving would be reduced. On the other hand, if sized too small, then the circuit would be too slow because of increased resistance to ground. A rule of thumb about sleep transistor sizing is to be 1/10 of total PMOS or NMOS area of the target circuit, depending on which kind of sleep transistor chosen. However, this rule is only a rough estimation, not really an optimization.

As shown in Figure 4.3, 3X minimum size of NMOS sleep transistor is chosen for our test vehicle, which brings about 4% frequency loss, 9% dynamic power saving, and less than 50mV virtual ground raising. This size is only 3% of total NMOS area of the ring oscillator, which is more aggressive than the sizing rule mentioned above. Although different target circuits have their own optimized sleep transistor sizes, the same analysis flow can be applied. By sketch characteristic curves like Figure 4.3 and needed specifications, optimized sleep transistor sizes of different target circuits can be clearly decided.

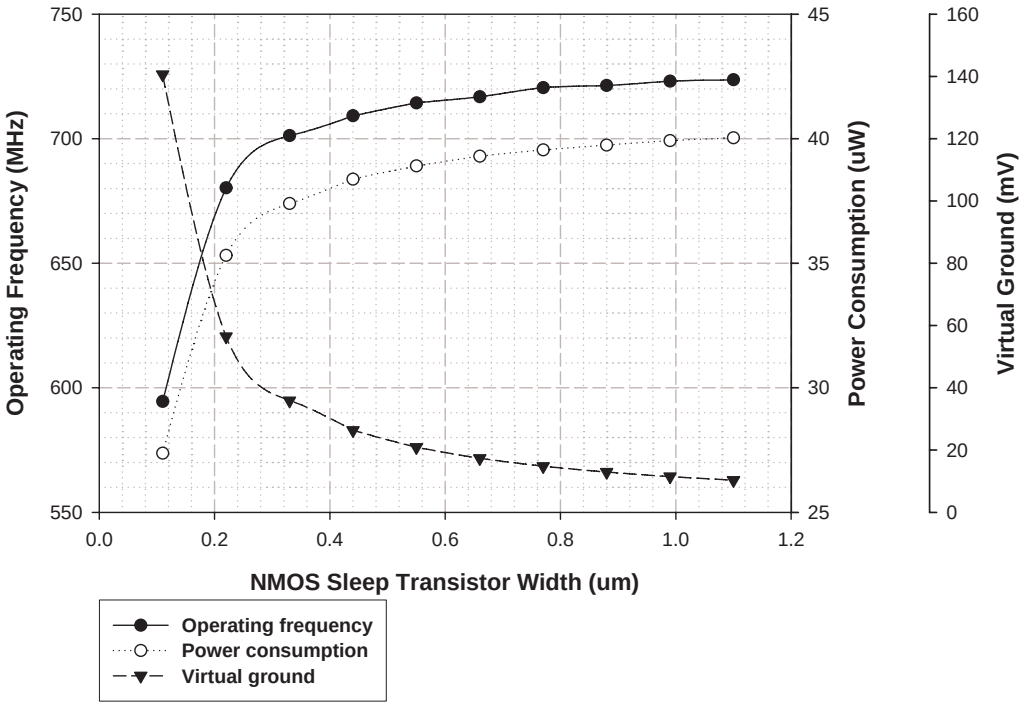

Figure 4.3 Relationship between frequency/power/virtual ground and NMOS sleep transistor width of the ring oscillator under TSMC 100nm CMOS technology. Without power gating, operating frequency and power consumption are 732.69MHz and 41.13uW, respectively.

The use of power gating is to reduce leakage power. As shown in Figure 4.4, leakage power can be suppressed from the order of microwatt to nanowatt, which means more than 90% leakage power can be reduced applying this technique. For example, since 3X minimum size of NMOS sleep transistor is chosen for our test vehicle, then leakage power can be reduced from 2.3uW to 12.66nW at room temperature. However, as noticed in Section 4.1, NMOS approaching minimum size should be avoided to be sleep transistors. In other words, the scale of the target circuit should be large enough to prevent optimized size of sleep transistor falls in this region. Otherwise, leakage power saving will be reduced, or extra area wasted will be produced to enlarge the sleep transistor. Because the scale of the ring oscillator is comparable to a 32-bit adder and its optimized sleep transistor size is just outside this unusual leaky region, it can be concluded that power gating is preferable in system and sub-system design, but not suitable in cell library design. As a result, the minimum scale of voltage islands is defined.

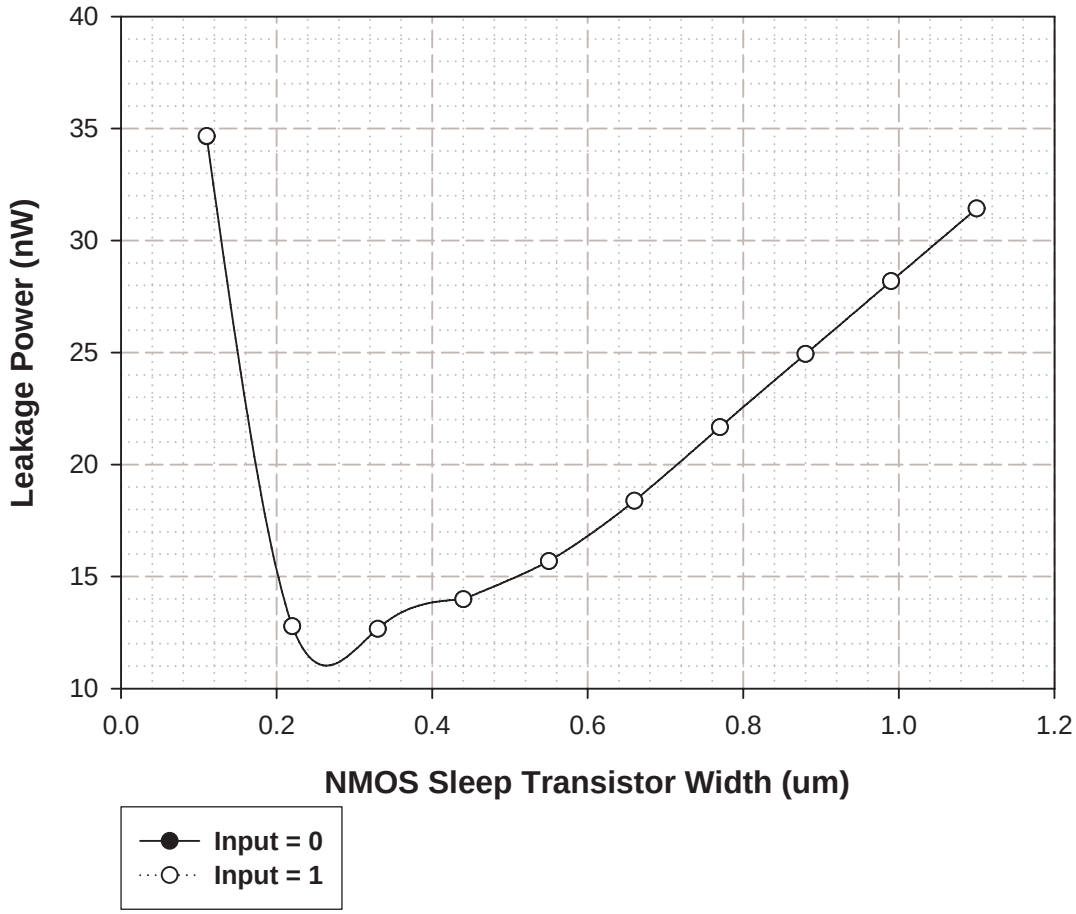

Figure 4.4 Relationship between leakage power and NMOS sleep transistor width of the ring oscillator under TSMC 100nm CMOS technology. Without power gating, leakage power is about 2.3uW.

4.1.4 Summary

Sleep transistors for power gating can be implemented by PMOS or NMOS. Because of the area constraint and better leakage reduction efficiency, we choose NMOS to be sleep transistors. Besides, Sleep transistor sizing is the most important and hardest design decision due to frequency loss, area overhead, and other trade-offs. 3X minimum size of NMOS sleep transistor is chosen for our test vehicle, which brings about 4% frequency loss, 9% dynamic power saving, and small virtual ground rising. More than 90% leakage saving can be achieved applying power gating technique. Details about co-working with reverse body bias will be discussed later in Section 4.3. Finally, the scale of voltage islands should be beyond a sub-system, such as a 32-bit adder, .etc, to prevent the use of leaky sleep transistors.

4.2 DVS Analysis

The goal of DVS analysis in this section is to build V_{DD} look-up tables for CTAPM. We first realize that the circuit operating at different supply voltages has different responses toward temperature gradient. Then the V_{DD} look-up table can be divided to three zones to complete individually.

4.2.1 Characteristic Simulation

Before building V_{DD} look-up tables for CTAPM, characteristics of DVS technique should be known first. Figure 4.5 shows relationships between operating frequency and supply voltage in five different technology corners. Figure 4.6 is the corresponding power statistics. It is verified in Figure 4.5 that the operating frequency is linearly proportional to the supply voltage. It is also described in Chapter 2 that the dynamic power is proportional to V_{DD}^2. As a result, if the operating frequency can be varied with the supply voltage, then we can say that the dynamic power is proportional to V_{DD}^3. Since the dynamic power is the dominant component of total power consumption in the active mode, the total power consumption is close this relationship, too. Therefore, Figure 4.6 is shown in that way.

However, Figure 4.6 reveals that the marginal benefit of reducing the supply voltage too low is quite little. When the dynamic power is reduced to be competed with leakage power, it is not worth scaling the supply voltage down further. Besides, if supply voltages of two adjacent functional blocks differ too large, needed level shifters will be too hard to design. Moreover, the efficiency of the on-chip DC-DC converter is worse in low voltage level transformation [30]. Last but not least, noise margins will be destroyed due to low voltage level of logic 1. For these reasons, even though the required operating frequency is low enough, scaling the supply voltage below 0.5V should be avoided.

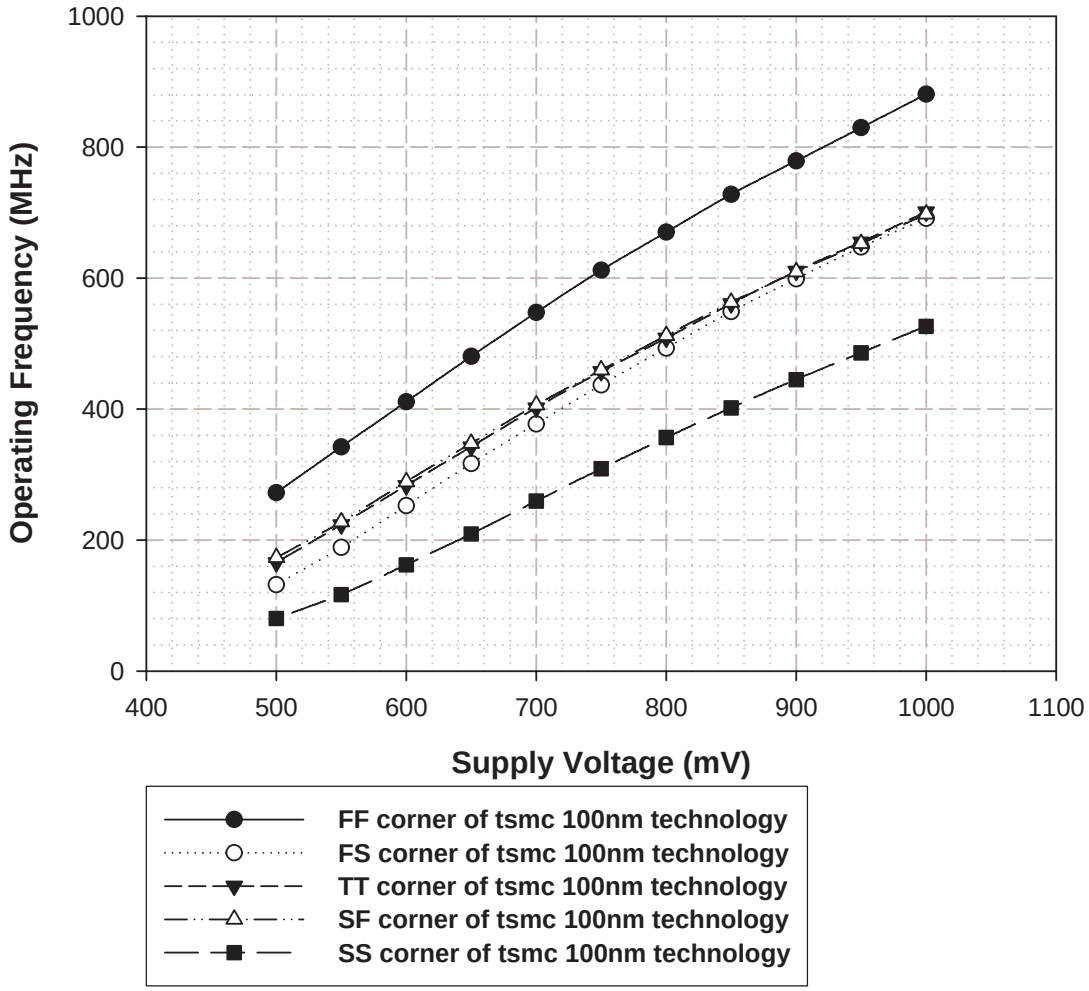

Figure 4.5 Relationships between operating frequency and supply voltage in five different technology corners.

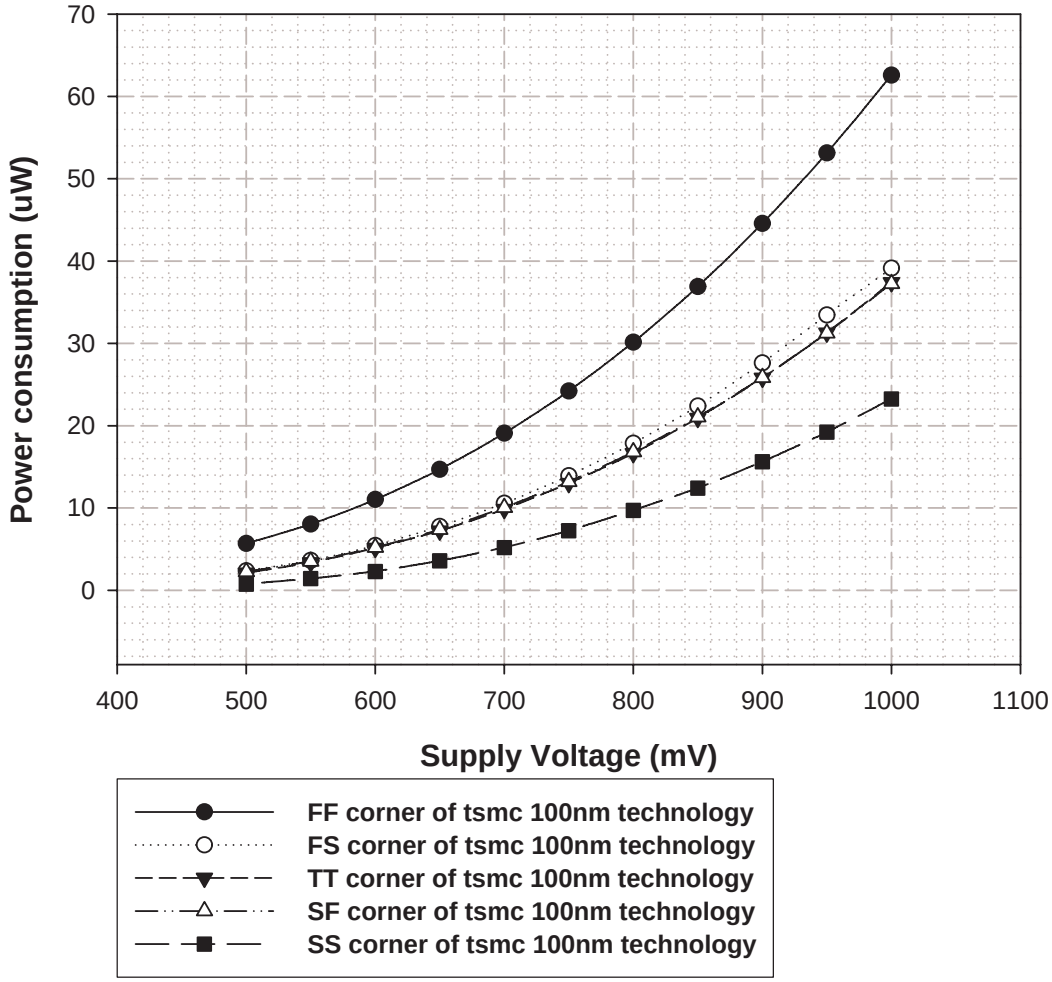

Figure 4.6 Relationships between power consumption and supply voltage in five different technology corners.

Figure 4.7 shows relationships between operating frequency and supply voltage in eight different temperature levels. Figure 4.8 is the corresponding power statistics. We notice that supply voltage range can be divided into high, middle, and low V_{DD} zones. When operating in the high V_{DD} zone, frequency decreases as temperature increases, which causes the problem of performance coherence between functional blocks. In the middle V_{DD} zone which is round 700mV in our simulation, frequency decreases slightly as temperature increases. However, frequency increases as temperature increases in the low V_{DD} zone, resulting in no performance coherence problem at all.

The reason of these phenomena is as follows: As temperature rises, the mobility decreases due to more phonon scattering, and the threshold voltage is reduced due to Fermi level. In high and middle V_{DD} zones, the mobility is dominant, so that the drain current and operating frequency are both reduced. However, the threshold voltage is the dominant factor in the low V_{DD} zone, and thus increases the drain current and operating frequency. While this concept is well-known to analog designers, it has so far not been emphasized in digital circuit design.

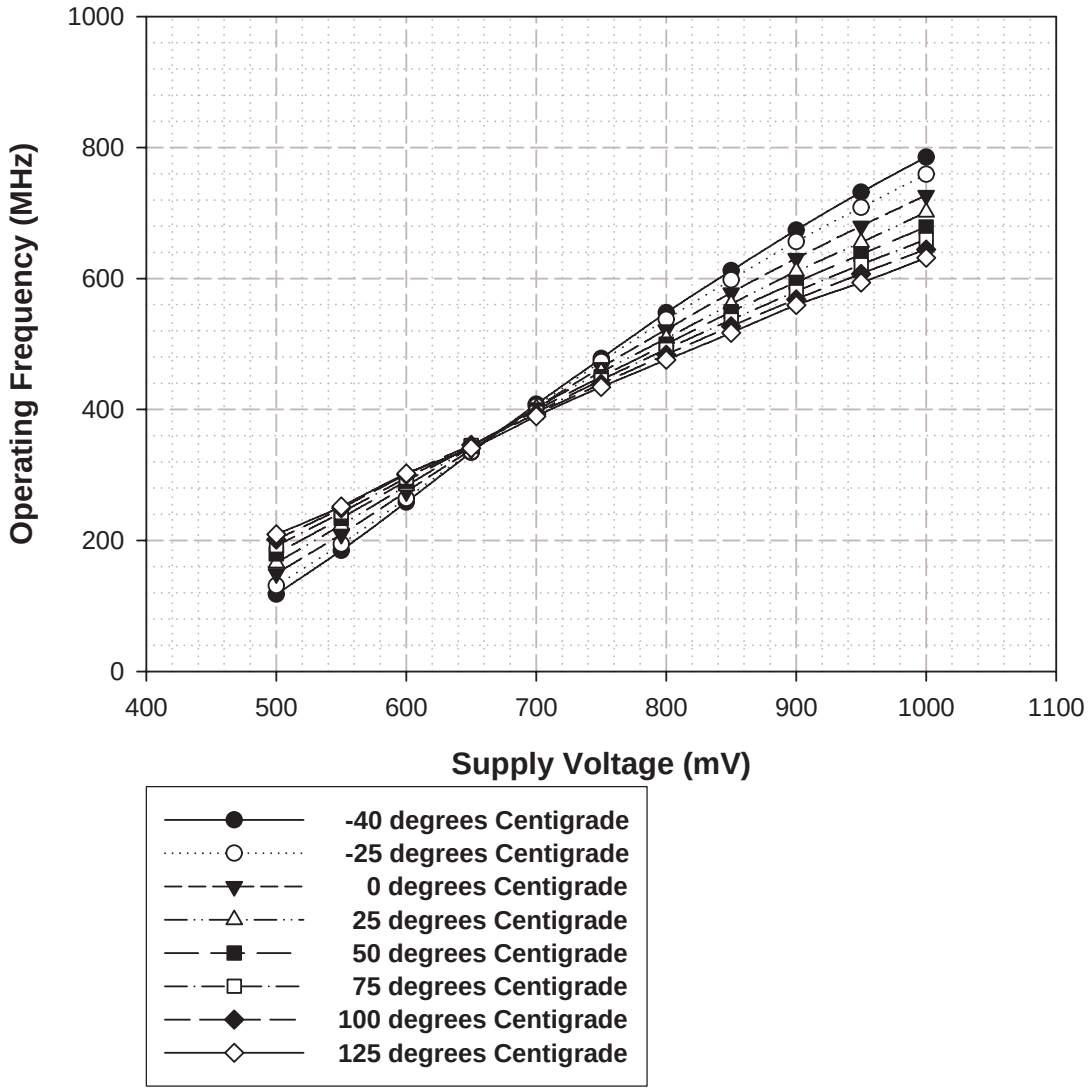

Figure 4.7 Relationships between operating frequency and supply voltage in eight different temperature levels.

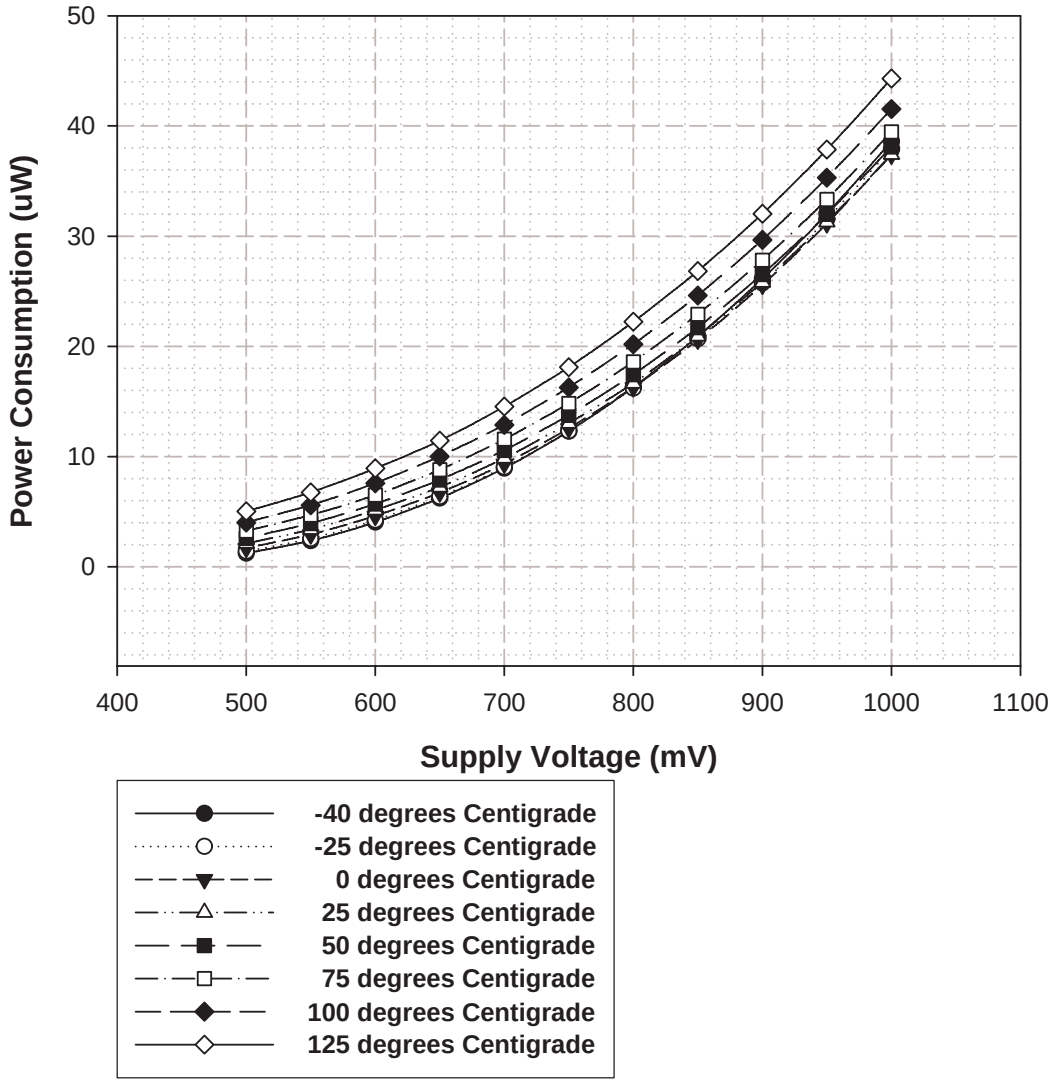

Figure 4.8 Relationships between power consumption and supply voltage in eight different temperature levels.

4.2.2 V_{DD} Look-Up Table Construction

Since frequency increases as temperature increases in the low V_{DD} zone, there is no performance coherence problem in this supply region. Therefore, we raise a margin of supply voltage in high V_{DD} and middle V_{DD} zones to keep performance coherence, but keep minimum required supply voltage in the low V_{DD} zone. Table 4.2 is simulation results of the V_{DD} look-up table. Without utilizing DVS, the supply voltage must be the highest one, 980mV in the table, to ensure normal operation. On the contrary, it only needs to be 500mV for 200MHz operation at 150°C if we utilize

DVS on both frequency and temperature dimensions. It can lead to 74% total power reduction and maintain performance coherence at the same time. By appropriate coding, the logical V_{DD} look-up table can be built like Table 4.3.

Table 4.2 Numerical V_{DD} look-up table for CTAPM. 10mV is the minimum step.

V_{DD} (mV)	-40~-25°C	-25~0°C	0~25°C	25~50°C	50~75°C	75~100°C	100~125°C	125~150°C
600 MHz	860	870	890	910	930	940	960	980
500 MHz	780	780	790	800	810	820	830	840
400 MHz	700	700	700	710	710	710	720	720
300 MHz	630	630	620	620	610	610	600	600
200 MHz	570	560	550	530	520	510	500	500

Table 4.3 Logical V_{DD} look-up table for CTAPM. Each VID number corresponds to an independent V_{DD} level, respectively.

VID	-40~-25°C	-25~0°C	0~25°C	25~50°C	50~75°C	75~100°C	100~125°C	125~150°C
600 MHz	100100	100101	100111	101001	101011	101100	101110	110000
500 MHz	011100	011100	011101	011110	011111	100000	100001	100010
400 MHz	010100	010100	010100	010101	010101	010101	010110	010110
300 MHz	001101	001101	001100	001100	001011	001011	001010	001010
200 MHz	000111	000110	000101	000011	000010	000001	000000	000000

In our initial plan, DVS is used in the large steps of temperature to meet performance requirements, and FBB will be applied in the small steps of temperature to adjust frequency changes due to temperature variation. There should be no need to raise a margin of the supply voltage in high V_{DD} and middle V_{DD} zones because of performance coherence needs. However, the effect of FBB is surprisingly too small to compensate frequency decrease based on used TSMC 100nm technology, which will be described later in Section 4.4. Consequently, DVS is used alone to solve the problem in this situation, raising a margin of the supply voltage to ensure performance coherence.

4.3 RBB Analysis

In this section, we present experimental results about reverse body bias technique. Owing to previous works about this technique mainly focusing on the effect of leakage suppression, we put our emphasis on the apply methods instead. As shown in Figure 4.9, although converging quickly, applying reverse body bias on PMOS is the least effective since the leakage current of OFF-state PMOS devices is much smaller than that of NMOS devices. Applying reverse body bias on both PMOS and NMOS leads to the largest leakage saving. However, applying reverse body bias on NMOS only is efficient due to lower power overhead of the charge-pump circuit, which is needed to provide body bias voltage.

Figure 4.10 is the case co-working with power gating technique. Because the magnitude of leakage current of the circuit is decided by the OFF-state NMOS sleep transistor in this situation, applying reverse body bias on this sleep transistor only is as effective as on all PMOS and NMOS devices. In addition, applying reverse body bias on NMOS sleep transistor only introduces the least loading and power overhead of charge-pump circuits.

What is worth to mention is that the magnitude of reverse body bias voltage has its limit; in fact, applying too large reverse body bias will bring about obvious second order effects (SCE), such as gate induced drain lowering (GIDL). Previous research indicates (see Figure 4.11) that beyond a certain optimal RBB voltage, which is about 0.3X V_{DD} for 0.13um technology, the transistor OFF-state current starts to increase due to increased GIDL leakage [19]. We follow the reference to choose 0.3V as optical reverse body bias voltage to avoid the GIDL effect because the lack of BSIM4 model.

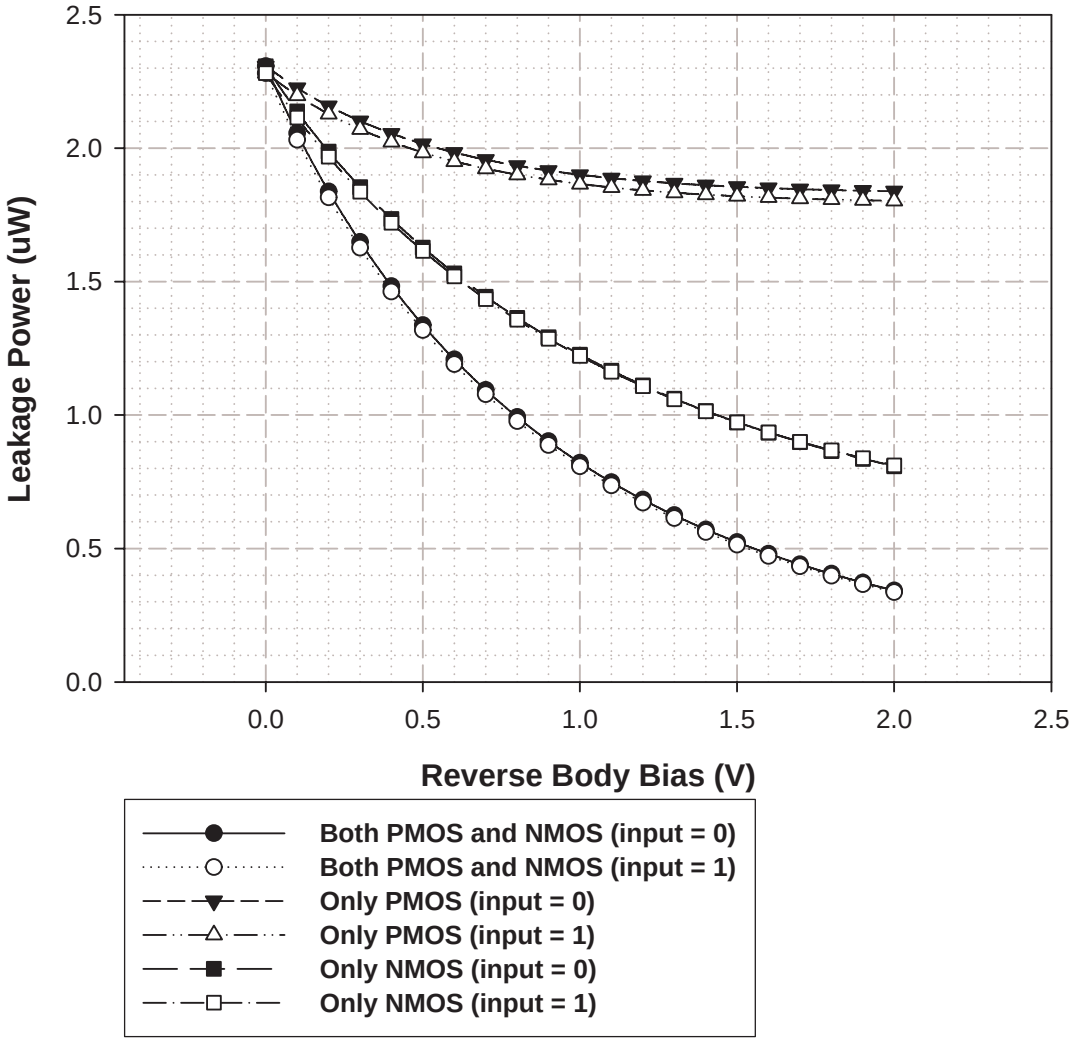

Figure 4.9 Relationships between leakage power and reverse body bias in different applying methods.

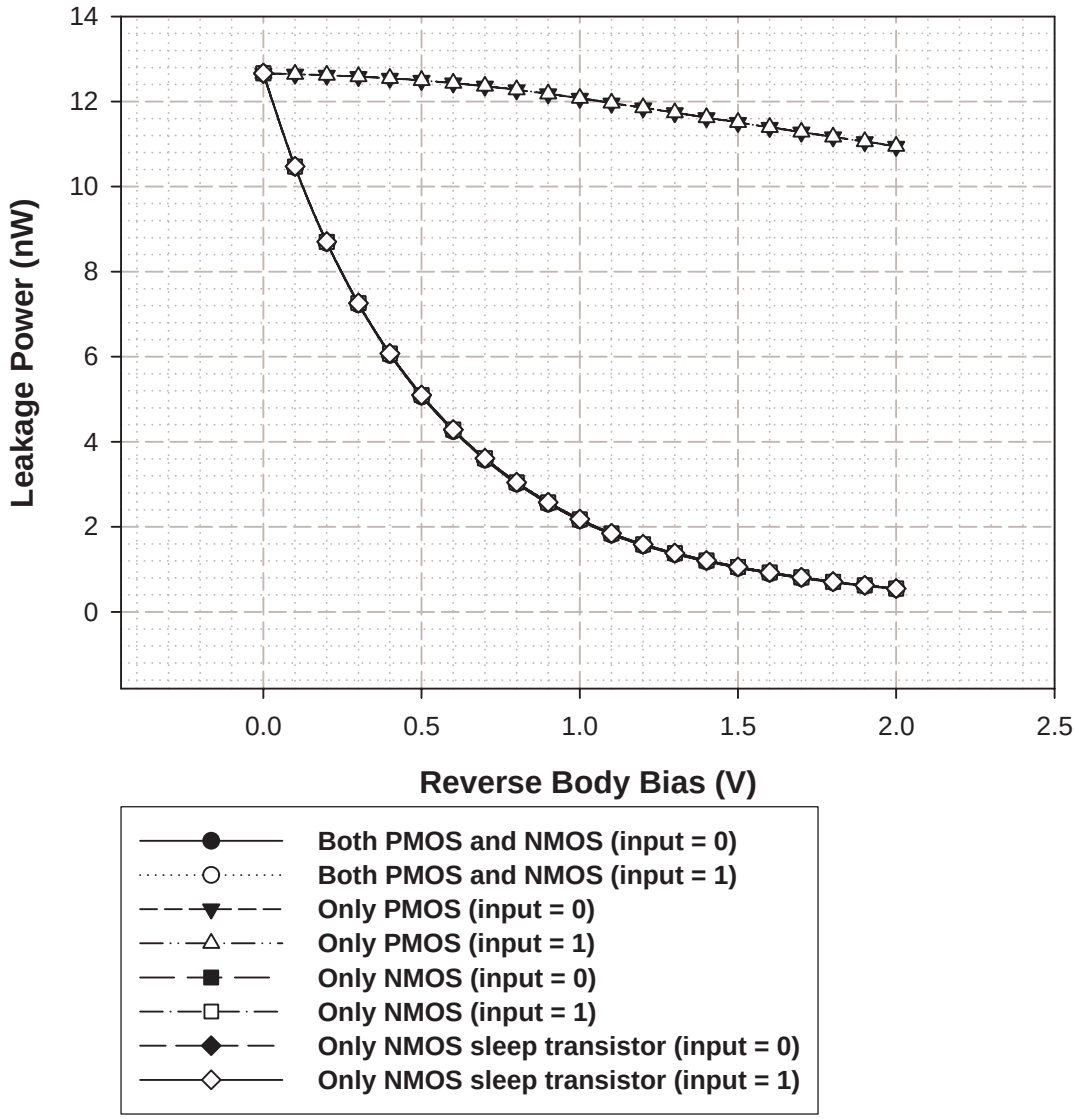

Figure 4.10 Relationships between leakage power and reverse body bias co-working with power gating in different applying methods.

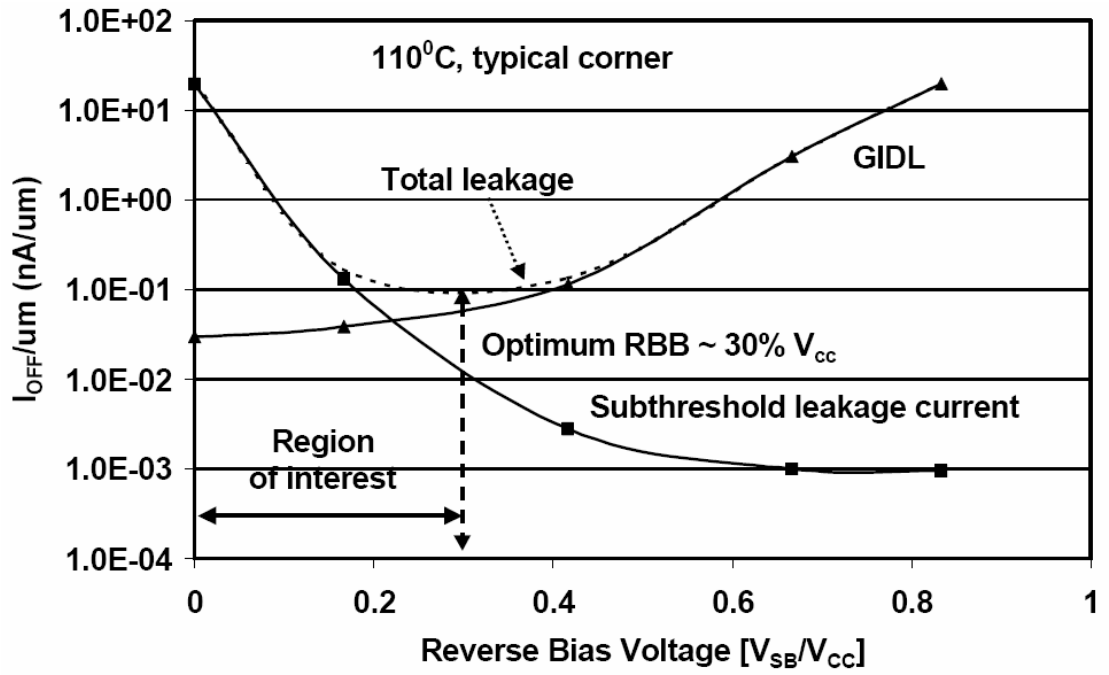

Figure 4.11 Optimum RBB voltage for 130nm technology

Utilizing power gating technique will reduce leakage power from 2.3uW to 12.66nW, and combining power gating with 0.3V reverse body bias can further reduce leakage power to 7.2nW. These simulation results show that power gating and reverse body bias are powerful to reduce leakage power.

4.4 FBB Analysis

Forward body bias is originally arranged to co-work with DVS to solve the problem of performance coherence in our CTAPM design. However, this technique doesn't show enough competence for the job based on used TSMC 100nm technology. Related simulation results will be expressed to explain the statement in this Section.

Reviewed from Section 4.2, supply voltage range can be divided into high, middle, and low V_{DD} zones. When operating in high and middle V_{DD} zones, frequency decreases as temperature increases, which causes the problem of performance coherence between functional blocks. However, frequency increases as temperature increases in the low V_{DD} zone, resulting in no performance coherence problem at all. Figure 4.12 and Figure 4.13 clearly illustrate these phenomena once more. In this place, we define that the 600MHz speed mode is in the high V_{DD} zone, and the 500MHz and 400MHz speed modes are in the middle V_{DD} zone, whereas the last two low speed modes are in the low V_{DD} zone.

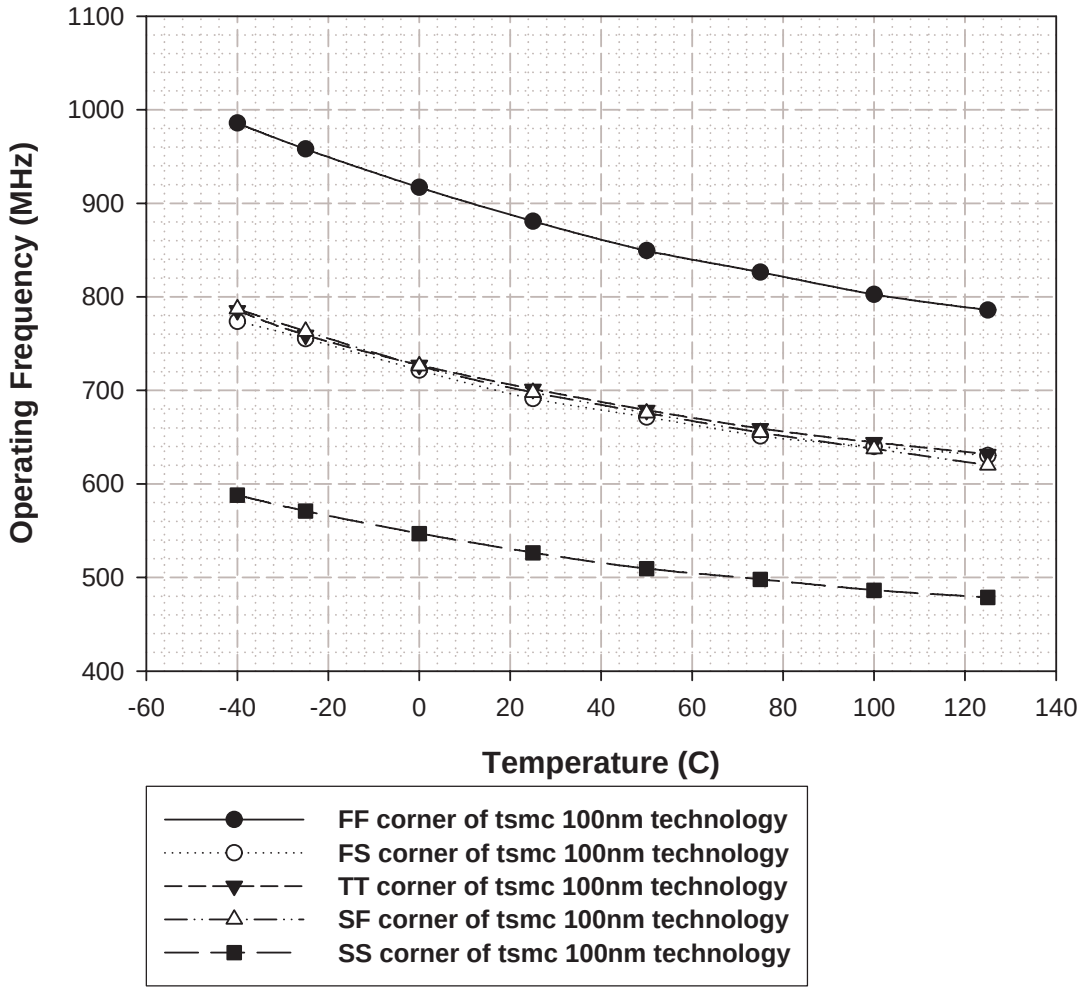

Figure 4.12 Relationship between operating frequency and temperature in five technology corners for the high V_{DD} zone. Supply voltage is 1.0V. Frequency decreases as temperature increases.

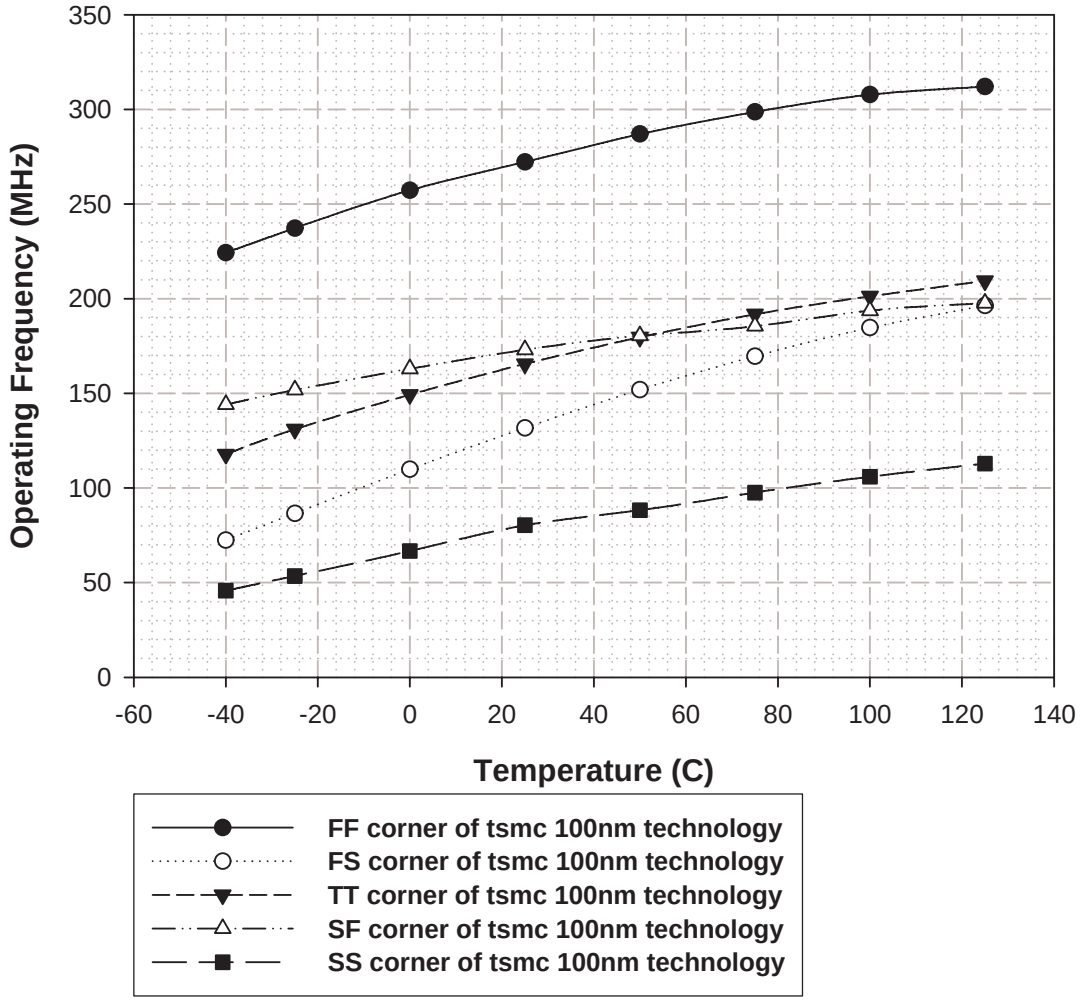

Figure 4.13 Relationship between operating frequency and temperature in five technology corners for the low V_{DD} zone. Supply voltage is 500mV. Frequency increases as temperature increases, so that there is no problem of performance coherence.

Following is the evidences showing that FBB is not that effective to increase speed. As shown in Figure 4.14, applying FBB on PMOS is more powerful than NOS because of stronger body effect for the same FBB voltage level. Even so, it still can't catch the frequency loss due to temperature raise. The speed of the ring oscillator decreases about 20MHz from 25°C to 50°C when the supply voltage is 1.0V, but applying FBB can only provide about 15MHz performance gain. Besides, although applying heavier FBB on PMOS can further increase the speed, the power overhead will be amazingly huge. For example, applying 650mV FBB on PMOS can compensate the frequency loss, but its power consumption at room temperature rises

from 37.40uW to 373.14uW, which is absolutely not applicable. Hence, we choose to raise 10 ~ 20mV of the supply voltage to deal with frequency sacrifice due to temperature rise in Section 4.2, since its power overhead is quite little.

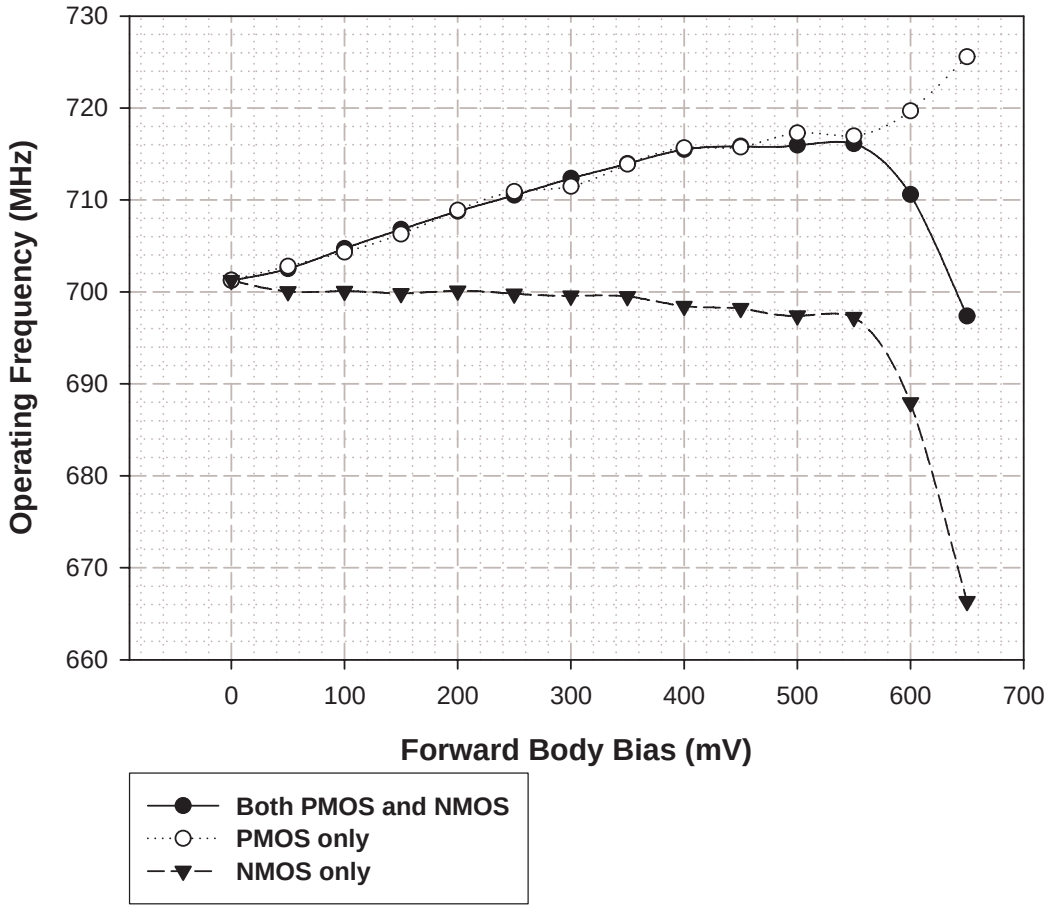

Figure 4.14 Relationship between operating frequency and forward body bias at 25°C in the high V_{DD} zone. Supply voltage is 1.0V.

Figure 4.15 shows similar results as Figure 4.14. We notice that the applicable FBB range is narrower at high temperature, reducing the speed increase that FBB can provide. Although the frequency loss is reduced to about 12MHz from 125°C to 150°C when the supply voltage is 1.0V, applying FBB can only provide about 6MHz performance gain, still not enough to overcome this thermal impact.

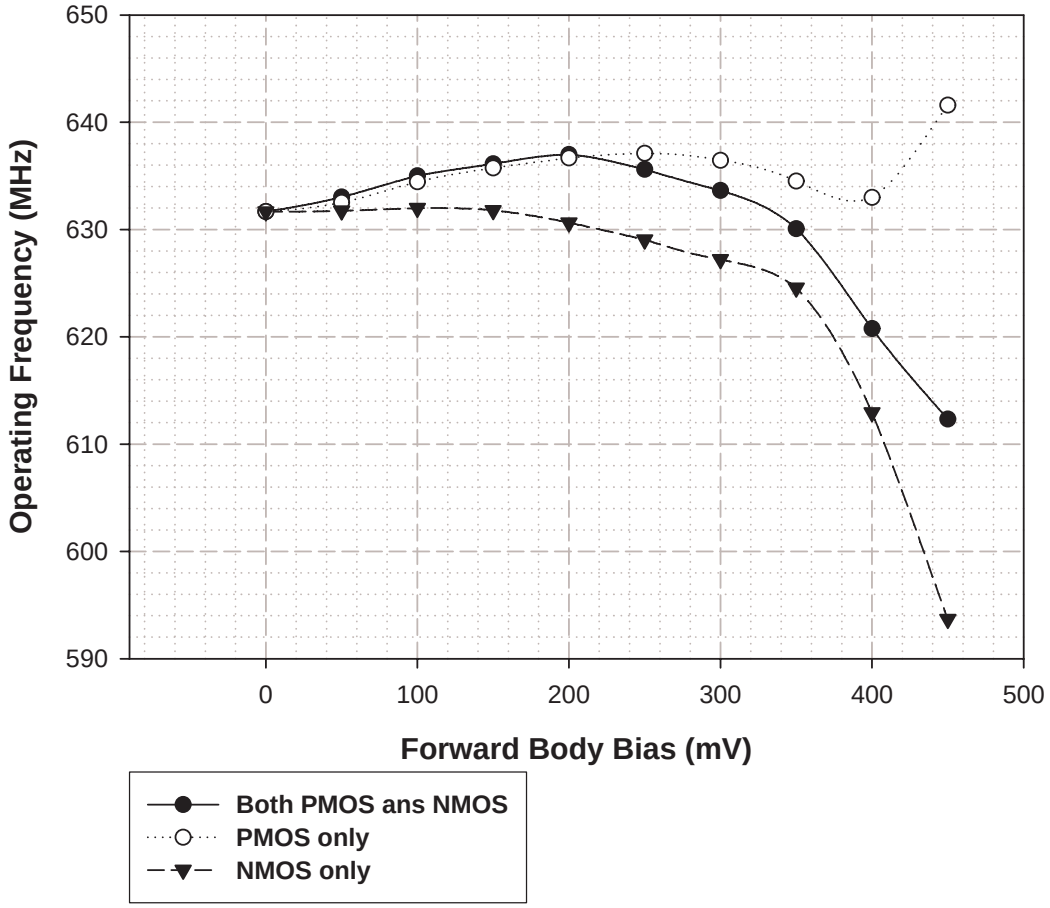

Figure 4.15 Relationship between operating frequency and forward body bias at 125°C in the high V_{DD} zone. Supply voltage is 1.0V.

Concerning about the middle V_{DD} zone, however, Figure 4.16 presents the power of forward body bias to increase speed. The frequency loss is about 8MHz from 25°C to 50°C when the supply voltage is 750mV, and applying FBB can provide about 30MHz performance gain to compensate this loss. Besides, the applicable FBB range is about the same with the high V_{DD} zone case, revealing that it is less dependent on the supply voltage than operating temperature.

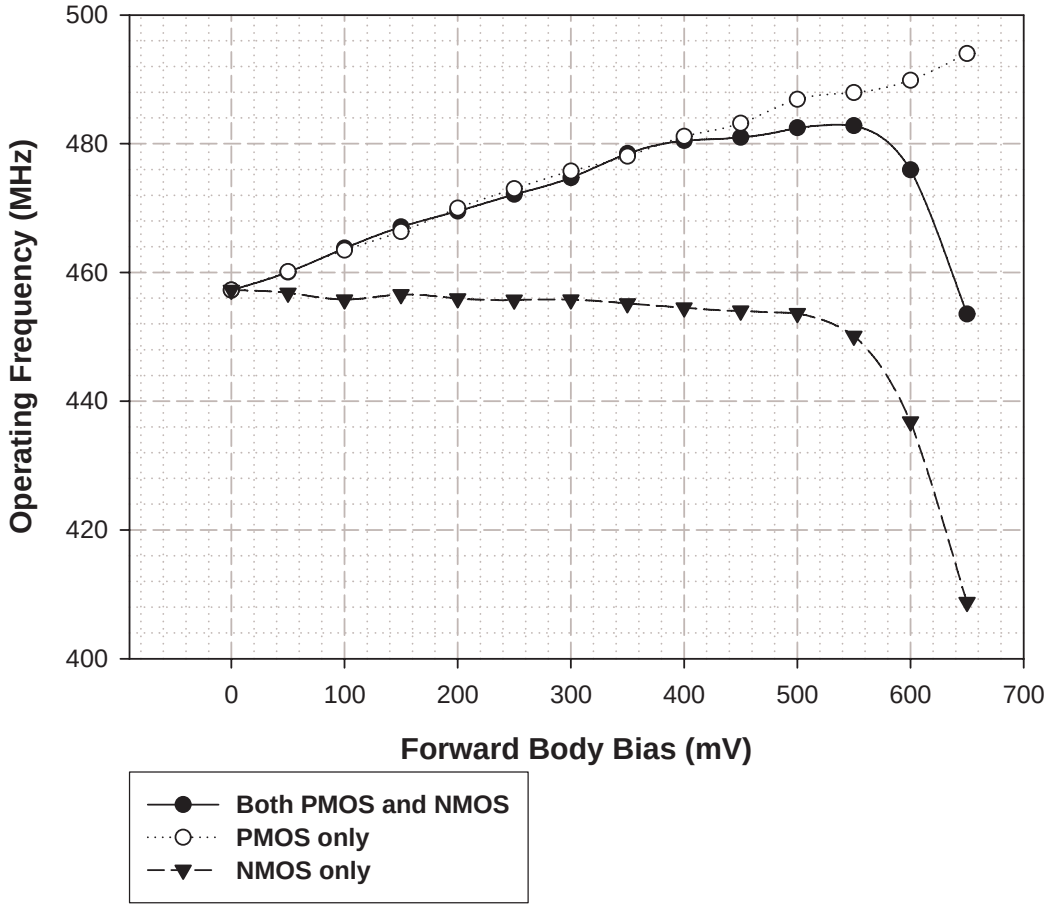

Figure 4.16 Relationship between operating frequency and forward body bias at 25°C in the middle V_{DD} zone. Supply voltage is 750mV.

According to Figure 4.14 to Figure 4.16, we can draw Figure 4.17 to conclude the FBB characteristics. The lower the supply voltage and operating temperature are the better performance improvement it has. Therefore, when operating in the high V_{DD} zone, DVS is used alone to ensure performance coherence between functional blocks. In the middle V_{DD} zone, FBB can co-work with DVS to solve the problem. However, it is unnecessary to adopt FBB in the low-V_{DD} operation, since the frequency, since there is no performance coherence problem in this situation.

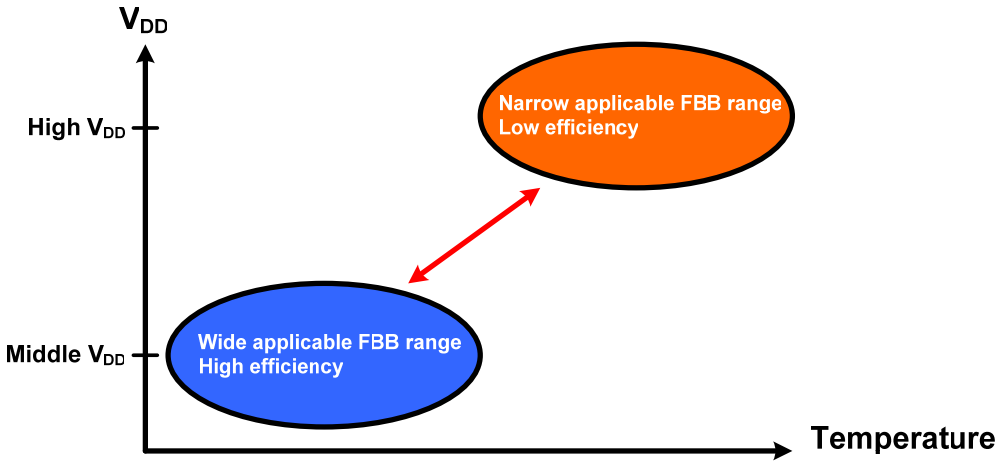

Figure 4.17 Illustration of FBB application efficiency. FBB is suitable in middle-V_{DD} and low-temperature environments.

At last, one thing should be clarified. Even though FBB is applicable in the middle V_{DD} zone, we don't take this option eventually. Instead, DVS is still used alone to face the frequency loss for simplicity. As shown in Figure 4.18, although applying FBB can reduce V_{DD} at the same target frequency, the power consumption is higher because the circuit is leakier due to the decrease of threshold voltage. Raising a margin V_{DD} is more efficiency to ensure performance coherence in this technology.

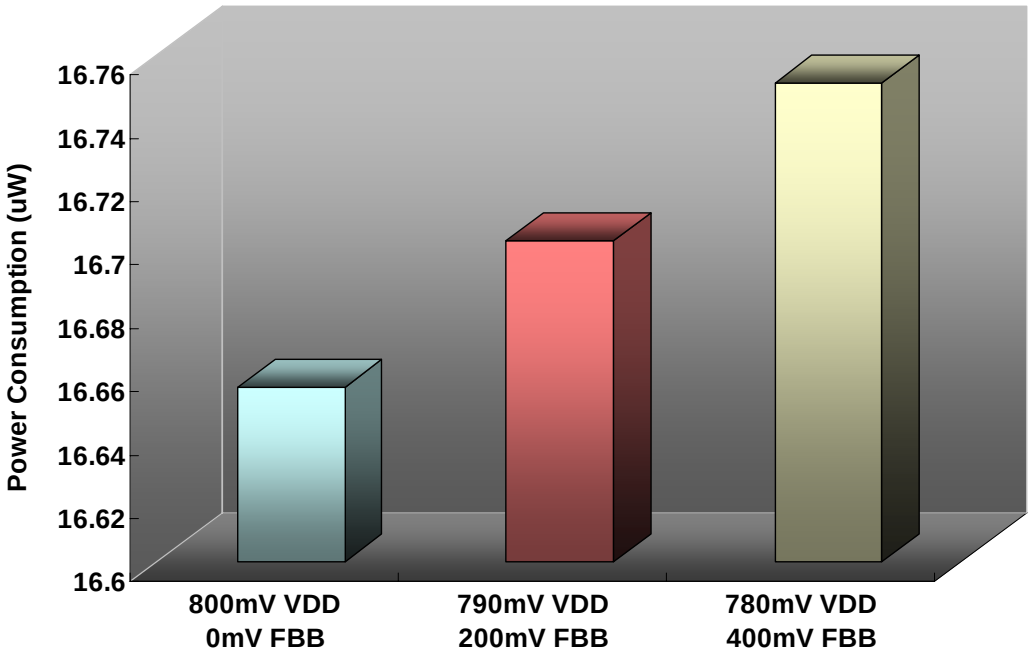

Figure 4.18 Comparison of power consumptions at 500MHz between different V_{DD} and FBB voltages. FBB is applied on PMOS.

4.5 Discussion

In our analysis flow, every design decision is based on its own trade-offs for each low power technique. For example, we consider area overhead, virtual ground noise and power consumption when choosing sleep transistor type for power gating, but we are not conscious that this earlier decision has impact for other techniques. In fact, technique-to technique trade-offs need to be taken into account when we make a design decision, and sleep transistor type selection is coincidentally the most complicated case among all design decisions.

Table 4.4 shows the advantages of choosing NMOS or PMOS to be sleep transistors corresponding to power gating and reverse body bias. Since both techniques deal with leakage reduction, they have to compromise to complement each other for some design constraints. For example, if the triple-well technology is not supported by the foundry, inevitably we can't use NMOS sleep transistors for small area overhead as we wish. Therefore, sometimes we need to "bottom-up" to change or reconsider previous design decisions.

Table 4.4 Comparison of different sleep transistor types concerning about both power gating and reverse body bias.

Decision / Technique	NMOS sleep transistor	PMOS sleep transistor
Power Gating	Small area overhead	Flexibility of sleep transistor sizing
	Less power consumption	Less virtual supply noise
Reverse Body Bias	Large leakage suppression	Low RBB for leakage convergence
	Small loading for charge-pump circuits	No need of triple-well technology

Besides, the type of technology process has great impact for architecture decisions. Modern nano-scale CMOS technology platforms typically provide up to six different types of logic core devices with different oxide thickness and threshold voltages for high performance, low leakage power, and low operating power requirements [paper and TSMC]. The effect of body bias is just restricted by oxide thickness and threshold voltages. Although not useful in our simulations based on TSMC 100nm CMOS technology, FBB is more powerful for devices with thicker oxides and larger threshold voltages [31]. On the contrary, RBB is useful for devices with thicker oxides and low threshold voltages. Namely, FBB is applicable for low power technologies while RBB is suitable for high speed technologies.

Therefore, for simplicity we give up FBB in our CTAPM system based on used technology. However, if the target technology changes and FBB enables enough performance gain through the same analysis flow, then it is still a solution to ensure performance coherence or compensate process variation in other power management designs. RBB also has the possibility to not be adopted for low leakage technology since threshold voltages are larger to reduce its effect. The full version of CTAPM system has been proposed in Chapter 3, Figure 4.19 now shows an example of simplified version of CTAPM system, which only uses DVS to adjudge performance in the active mode and power gating to suppress leakage power in the standby mode, but skip body bias for simplicity. The analysis flow is still the same for different technologies, while the selection of low-power techniques will be case by case.

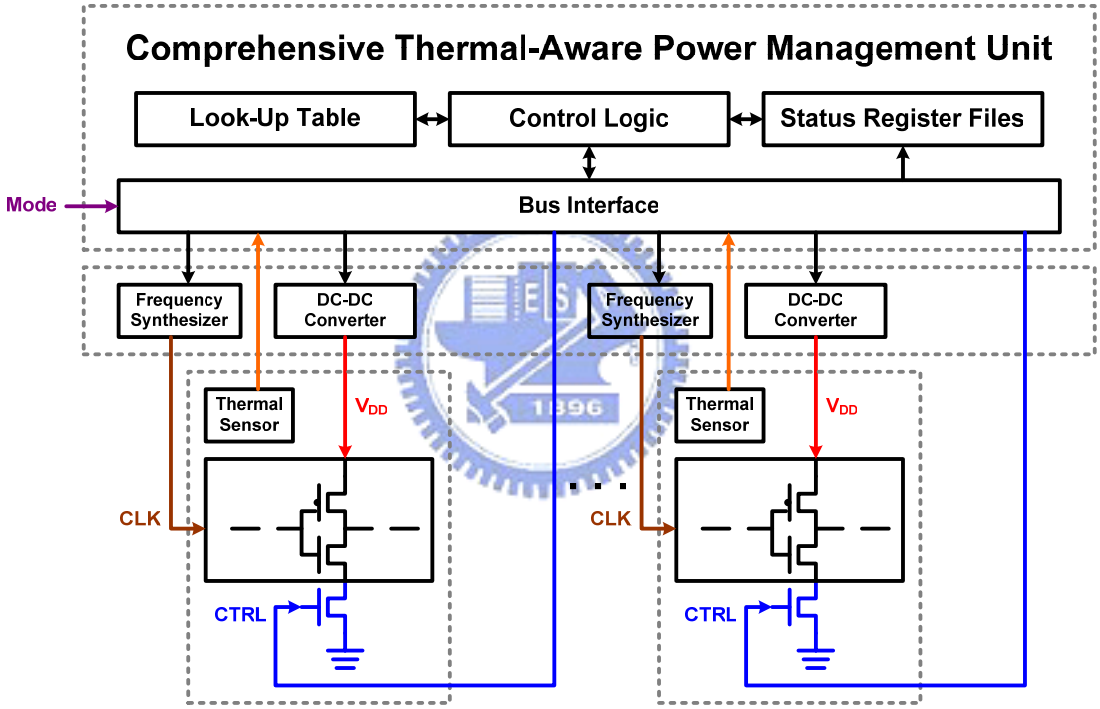

Figure 4.19 Simplified version of the CTAPM unit architecture. DVS, clock gating and power gating still remain to maintain the functionality, whereas charge-pump circuits for body bias are removed.

4.6 Summary

Utilizing power gating technique will reduce leakage power from 2.3uW to 12.66nW for our test vehicle, and combining power gating with 0.3V reverse body bias can further reduce leakage power to 7.2nW at room temperature. More than 90% leakage saving can be achieved applying these techniques. However, because the lack of BSIM4 model, we can not obtain the optimized RBB values at different temperatures for the used technology.

Beside power control mechanism, we utilize DVS on the temperature axis to meet performance requirements. The experimental results show that the proposed CTAPM architecture has potential to reduce more than 70% of dynamic power dissipation in regular CMOS systems and ensure performance coherence at the same time. Nevertheless, forward body bias is not powerful as we expected based on used technology. Performance gain it provided is hard to compensate frequency loss when the temperature gets higher, especially at high supply voltage level.

Finally, design trade-offs between different techniques and technologies are discussed in this chapter. The type of sleep transistors has the influence on both power gating and reverse body bias. Besides, FBB may be useful in low power technologies while RBB is powerful in high speed technologies. The proposed full-version CTAPM architecture can be modified or simplified while maintaining the same functionality.

Chapter 5

Conclusion

5.1 Conclusion

Power dissipation in modern VLSI designs has become the most critical issue in System-on-Chip (SoC) era. Modern SoC integrations and mobile systems have emphasized low power techniques due to shortage of battery life. Besides, thermal impact is another reason to address the importance of power management because of high costs of packaging and cooling requirements for reliability.

In this thesis, a block-level optimization of comprehensive thermal-aware power management (CTAPM) is presented. The proposed CTAPM system adopts several circuit-level low power techniques to take care of both dynamic power and leakage power consumption in power phase, and deal with performance coherence between functional blocks in thermal phase as well. Dynamic voltage scaling (DVS) is utilized on both performance and temperature axis to reduce dynamic power and ensure all voltages islands are synchronous in the active mode. Clock gating is applied to reduce unnecessary clock switching power in the standby mode. Power gating and reverse body bias is aimed at leakage power suppression. However, forward body bias is not powerful to use.

Basic concepts including architecture design, control state machine, and look-up table design are also introduced after the overview of power source categories and these low power techniques. Through proposed analysis flow and test vehicle, the experimental results show that the proposed CTAPM architecture has potential to reduce more than 70% of dynamic and 90% of static power dissipation in regular CMOS systems. It also brings a significant improvement in system stability. The results are based on TSMC 100nm technology. Last but not least, the analysis flow is the same for different technologies, while the selection of low-power techniques can be case by case.

5.2 Future Works

Although transient analysis is not covered in this thesis, it is important for hardware or software scheduling to place the CTAPM into a real SoC system. Transition time and power need to be studied for each adopted low power technology. For example, if using power gating saves 10uW of leakage power in the standby mode, but it consumes extra 20uW transition power to close sleep transistors, then the system should keep sleep transistors OFF for at least two seconds to gain net power reduction. For another example, if it needs 10ms to change the supply voltage to different level, then the system has to be stopped for the same interval during the transition. Such transient information can be known by simulations of peripheral circuits or direct physical measurements. As it should be, peripheral circuits, such as DC-DC converters for DVS, charge-pump circuits for body bias, and thermal sensors for software feedback, are also future works to realize the CTAPM system. Moreover, since the analysis flow is compatible, a standardized test key for individual optimization of chips can be developed in the future.

Noise margin consideration is another important issue for future researchers. The ideally optimized V_{DD} look-up tables may not suitable for some noise-sensitive circuits, since the supply voltage may be scaled down a lot to destroy noise margins. Besides, sleep transistors acting as minor power-rail resistance contribute virtual supply noise in normal operation. These drawbacks should be examined to maintain system stability.

Finally, the cell-based design flow may have to be modified to implement these low power techniques. However, it is not an issue for full-custom design. Since the hardware-language coding can't generate sleep transistors for power gating, it should be done by either CAD tool support during the place-and-route (P&R) stage or post-layout custom modification. Besides, substrate/well contacts of original standard cells must be removed and re-create properly to isolate the body of transistors for body bias technique [32]. Voltage separation also needs to be done to apply DVS individually for voltage islands [33]. These back-end design processes have to be rearranged into the conventional cell-based design flow.

Bibliography

[1] M. Powell, S.-H. Yang, B. Falsafi, K. Roy, and T. N. Vijaykumar, "Gated-Vdd: A Circuit Tecnhique to Reduce Leakage in Deep-Submicron Cache Memories," in *International Symposium on Low Power Electronics and Design*, 2000.

[2] M. Nakai, S. Akui, K. Seno, T. Meguro, T. Seki, T. Kondo, A. Hashiguchi, H. Kawahara, K. Kumano, and M. Shimura, "Dynamic Voltage and Frequency Management for a Low-Power Embedded Microprocessor," *IEEE Journal of Solid-State Circuits*, vol. 40, pp. 28-35, January 2005.

[3] J. T. Kao, M. Miyazaki, and A. P. Chandrakasan, "A 175-mV Multiply-Accumulate Unit Using an Adaptive Supply Voltage and Body Bias Architecture," *IEEE Journal of Solid-State Circuits*, vol. 37, pp. 1545-1554, November 2002.

[4] K. Roy, S. Mukhopadhyay, and H. Mahmoodi-Meimand, "Leakage Current Mechanisms and Leakage Reduction Techniques in Deep-Submicrometer CMOS Circuits," *Proceedings of the IEEE*, February 2003.

[5] J. T. Kao and A. P. Chandrakasan, "Dual-Threshold Voltage Techniques for Low-Power Digital Circuits," *IEEE Journal of Solid-State Circuits*, vol. 35, pp. 1009-1018, July 2000.

[6] J. W. Tschanz, S. G. Narendra, Y. Ye, B. A. Bloechel, S. Borkar, and V. De, "Dynamic Sleep Transistor and Body Bias for Active Leakage Power Control of Microprocessors," *IEEE Journal of Solid-State Circuits*, vol. 38, pp. 1838-1845, November 2003.

[7] S. Lee and T. Sakurai, "Run-Time Voltage Hopping for Low-Power Real-Time Systems," in *37th Conference on Design Automation*, 2000.

[8] E. Rotem, A. Naveh, M. Moffie, and A. Mendelson, "Analysis of Thermal Monitor Features of The Intel Pentium M Processor," in *TACS Workshop at the 31st Annual International Symposium on Computer Architecture*, 2004.

[9] T. D. Burd, T. A. Pering, A. J. Stratakos, and R. W. Broderson, "A Dynamic Voltage Scaled Microprocessor System," *IEEE Journal of Solid-State Circuits*, vol. 35, pp. 1571-1580, November 2000.

[10] L. T. Clark, E. J. Hoffman, J. Miller, M. Biyani, Y. Liao, S. Strazdus, M. Morrow, K. E. Velarde, and M. A. Yarch, "An Embedded 32-b Microprocessor Core for Low-Power and High-Performance Applications," *IEEE Journal of Solid-State Circuits*, vol. 36, pp. 1599-1608, November 2001.

[11] N. J. Rohrer, C. Lichtenau, P. A. Sandon, P. Kartschoke, E. Cohen, M. G. Canada, T. Pfluger, M. I. Ringler, R. B. Hilgendorf, S. Geissler, and J. S. Zimmerman, "A 64-bit Microprocessor in 130-nm and 90-nm Technologies with Power Management Features," *IEEE Journal of Solid-State Circuits*, vol. 40, pp. 19-27, January 2005.

[12] Q. Wu, M. Pedram, and X. Wu, "Clock Gating and Its Application to Low Power Design of Sequential Circuits," in *IEEE Transactions on Circuit and Systems I: Fundamental Theory and Applications*, vol. 47, March 2000, pp. 415-420.

[13] W. Ye and M. J. Irwin, "Power Analysis of Gated Pipeline Registers," *Proceedings of the 20th Annual IEEE International ASIC/SOC Conference*, September 1999.

[14] A. Amirabadi, J. Jafari, A. Afzali-Kusha, M. Nourani, and A. Khaki-Firooz, "Leakage Current Reduction by New Technique in Standby Mode," in *14th ACM Great Lakes Symposium on VLSI*, 2004.

[15] C.-H. Hua and W. Hwang, "A Power Gating Structure with Concurrent Data Retention and Intermediate Modes in 100nm CMOS Technology," in *15th VLSI Design/CAD Symposium*, 2004.

[16] B. H. Calhoun and A. P. Chandrakasan, "Standby Power Reduction Using Dynamic Voltage Scaling and Canary Flip-Flop Structures," *IEEE Journal of Solid-State Circuits*, vol. 39, pp. 1504-1511, September 2004.

[17] J.-H. Choi, Y. Xu, and T. Sakurai, "Statistical Leakage Current Reduction in High-Leakage Environments Using Locality of Blocks Activation in Time Domain," *IEEE Journal of Solid-State Circuits*, vol. 39, pp. 1497-1503, September 2004.

[18] T. Kuroda, T. Fujita, S. Mita, T. Nagamatsu, S. Yoshioka, K. Suzuki, F. Sano, M. Norishima, M. Murota, M. Kako, M. Kinugawa, M. Karumu, and T. Sakurai, "A 0.9-V, 150-MHz, 10-mW, 4mm^2, 2-D Discrete Cosine Transform Core Processor with Variable Threshold-Voltage (VT) Scheme," *IEEE Journal of Solid-State Circuits*, vol. 31, pp. 1770-1779, November 1996.

[19] B. Chatterjee, M. Sachdev, S. Hsu, R. Krishnamurthy, and S. Borkar, "Effectiveness and Scaling Trends of Leakage Control Techniques for Sub-130nm CMOS Technologies," in *International Symposium on Low Power Electronics and Design*, 2003.

[20] A. Keshavarzi, S. Ma, S. Narendra, B. Bloechel, K.Mistry, T. Ghani, S. Borkar, and V. De, "Effectiveness of Reverse Body Bias for Leakage Control in Scaled

Dual Vt CMOS ICs," in *International Symposium on Low Power Electronics and Design*, 2001.

[21] T.-S. Cheng and W. Hwang, "On-Chip Dual-Level Body-Bias Generators Applied to SRAM Cell Array Design for Leakage Suppression," in *15th VLSI Design/CAD Symposium*, 2004.

[22] G. Ono, M. Miyazaki, H. Tanaka, N. Ohkubo, and T. Kawahara, "Temperature Referenced Supply Voltage and Forward-Body-Bias Control (TSFC) Architecture for Minimum Power Consumption," in *30th European Solid-State Circuit Conference*, 2004.

[23] K. Nose, M. Hirabayashi, H. Kawaguchi, S. Lee, and T. Sakurai, "VTH-Hopping Scheme to Reduce Subthreshold Leakage for Low-Power Processors," *IEEE Journal of Solid-State Circuits*, vol. 37, pp. 413-419, March 2002.

[24] K. Usami and M. Horowitz, "Clustered Voltage Scaling Technique for Low-Power Design," in *International Symposium on Low Power Electronics and Design*, 1995.

[25] D. E. Lackey, P. S. Zuchowski, T. R. Bednar, D. W. Stout, S. W. Gould, and J. M. Cohn, "Managing Power and Performance for System-on-Chip Designs Using Voltage Islands," in *International Conference on Computer Aided Design*, 2002.

[26] R. Puri, L. Stok, J. Cohn, D. Kung, D. Pan, D. Sylvester, A. Srivastava, and S. Kulkarni, "Pushing ASIC Performance in a Power Envelope," in *40th Design Automation Conference*, 2003.

[27] C. Chen, A. Srivastava, and M. Sarrafzadeh, "On Gate Level Power Optimization Using Dual-Supply Voltages," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 9, pp. 616-629, October 2001.

[28] Intel® Pentium® M Processor on 90 nm Process with 2-MB L2 Cache Datasheet. http://www.intel.com/design/mobile/datashts/302189.htm.

[29] M. Yoshimoto and K. Kawakami, "A Feed-Foreward Dynamic Voltage Frequency Management by Workload Prediction for a Low Power Motion Video Compression," in *12th Workshop on Synthesis and System Integration of Mixed Information Technologies*, 2004.

[30] G. Patounakis, Y. W. Li, and K. L. Shepard, "A Fully Integrated On-Chip DC-DC Conversion and Power Management System," *IEEE Journal of Solid-State Circuits*, vol. 39, pp. 443-451, March 2004.

[31] K. V. Arnim, E. Borinski, P. Seegebrecht, H. Fiedler, R. Brederlow, R. Thewes, J. Berthold, and C. Pacha, "Efficiency of Body Biasing in 90-nm CMOS for

Low-Power Digital Circuits," *IEEE Journal of Solid-State Circuits*, vol. 40, pp. 1549-1556, 2005.

[32] M. Hirabayashi, K. Nose, and T. Sakurai, "Design Methodology and Optimization Strategy for Dual-V_{TH} Scheme Using Commercially Available Tools," in *International Symposium on Low Power Electronics and Design*, 2001.

[33] J. Wang, “The Design of a Parallel Data Path with Dynamic Power Management,” *Master Thesis*, National Chiao Tung University, 2004.