

Power-Saving Techniques and Future Design of SPARC64 V/VI

● Aiichiro Inoue

(Manuscript received May 31, 2007)

The performance of high-end microprocessors has increased along with improvements in semiconductor integration. Unfortunately, this increase in performance has been accompanied by an increase in power consumption. This paper describes the techniques we used to save power in the SPARC64 V/VI microprocessor for mission-critical servers. First, it describes how we reduced dynamic power consumption with designs that reduced circuit switching. Then, it describes the circuit management techniques employed to reduce static power consumption due to transistor leakage current in a static state. The paper finally analyzes the power consumption of the SPARC64 V/VI with these technologies applied, and discusses the future direction of design for this microprocessor.

1. Introduction

The semiconductor manufacturing process has evolved to achieve advantages in miniaturization that dramatically improve microprocessor performance. At the same time, higher levels of performance have resulted in a steady increase in power consumption. Due to recently imposed physical restrictions, semiconductor design is undergoing a paradigm shift toward focusing on power use. Fujitsu has utilized its own production facilities to develop high-performance microprocessors that consume low power. The power consumption performance of these microprocessors is among the best in the world.

This paper describes the characteristics of semiconductor power consumption, outlines the techniques that we applied to the SPARC64 V/VI processors¹⁾⁻⁷⁾ for reducing power consumption without sacrificing performance, and then discusses the future direction of microprocessor design.

2. Characteristics of power consumption

CMOS semiconductors are characterized by low power consumption and high integration. Power consumption (P) can be divided into a dynamic component (P_a) caused by circuit switching and a static component (P_s) as expressed below in Equation (1).

$$P = P_a + P_s \quad (1)$$

2.1 Dynamic power consumption

Since dynamic power component P_a is caused by electric charging and discharging in circuit switching, P_a can be expressed as shown below in Equation (2), where C_{eff} denotes the chip-equivalent capacitance, V_{dd} the supply voltage, and f the operating frequency.

$$P_a = C_{eff} \times V_{dd}^2 \times f \quad (2)$$

As expressed in Equation (2), lowering both the supply voltage and frequency reduces power consumption, but also adversely affects perfor-

Table 1
Power-saving techniques for SPARC64 V/VI.

	Main purpose		SPARC64 V (130 nm)	SPARC64 V (90 nm)	SPARC64 VI (90 nm)
	Dynamic power saving	Static power saving			
Power-saving latch	○	–	△	△	△
Latch-IH control	○	–	○	◎	◎
RAM-CE control	○	–	○	○	○
Chopper-IH control	○	–	–	–	△
Multi-Vth	–	○	△	○	○
Capacitor-cell leakage reduction	–	○	–	○	○
Supply voltage V_{dd} control of each chip	○	○	–	–	○

◎: Widely applied ○: Applied △: Partially applied

mance. C_{eff} depends on the circuit scale and semiconductor manufacturing process employed. When the same circuit is used, C_{eff} decreases by 0.7 to 0.8 times through miniaturization with each progressive semiconductor generation. Not all circuits are switched based on the clock in a given cycle; therefore, C_{eff} is lower than total the circuit capacitance.

Reducing C_{eff} is consequently important in reducing power consumption without affecting performance. Specifically, the latest semiconductor process must be employed, efficient and compact circuits should be designed, and the circuits switched only when necessary.

2.2 Static power consumption

Leakage current that flows without circuit operation while the power is on causes static power component P_s . Moreover, the static power consumption in recent generations of miniaturized semiconductors (especially the 90-nm circuit generation and later) is simply too high to ignore.

To analyze static power consumption, we calculated leakage current (I_{leak}) through experiments conducted using the SPARC64 V/VI built based on Fujitsu's 90-nm manufacturing process. This calculation method is explained in Reference 8).

$$\begin{aligned}
 P_s &= V_{dd} \times I_{leak} \\
 &= V_{dd} \times I_{ddq} \times e^{\{\alpha(V_{dd} - V_0) + \beta(t - t_0)\}}
 \end{aligned} \quad (3)$$

where, I_{ddq} denotes the leakage current measured under the reference conditions, V_{dd} the supply voltage, V_0 the reference voltage at I_{ddq} measurement, t the junction temperature, t_0 the reference temperature at I_{ddq} measurement, α the voltage coefficient, and β the temperature coefficient.

The actual measurement results of the SPARC64 V are $\alpha = 3.06$ and $\beta = 0.0157$ at $V_0 = 1V$ and $t_0 = 85^\circ C$. The actual measurement results of the SPARC64 VI are $\alpha = 2.86$ and $\beta = 0.0168$ at $V_0 = 1V$ and $t_0 = 85^\circ C$.

As shown in Equation (3), leakage current I_{leak} increases exponentially when supply voltage V_{dd} or temperature t increases. The actual measurement results show that I_{ddq} ranges from several to several tens of amperes. The variation, depending on individual differences, tends to increase with each advance made in the semiconductor manufacturing process. Reducing this variation poses a problem in miniaturization techniques. To reduce the leakage current, rises in temperature during operation must be curtailed.

3. Power-saving techniques for SPARC64 V/VI

Table 1 lists the application states of the power-saving techniques used for the SPARC64 V/VI. These techniques are classified as dynamic power-saving techniques and static power-saving techniques. Note that different




Figure 1
Circuit diagram of power-saving latch.

power consumption results are obtained depending on how the chips are used, even for the same individual chip.

3.1 Reducing dynamic power consumption

1) Power-saving latch^{note 1)}

Figure 1 shows an example of a circuit diagram for a power-saving latch. In this type of latch, the clock pulses usually distributed to both the master and slave sides are supplied only to the master side, and thus reduce the clock load capacity. Using this latch reduces power consumption by 26% as compared to the conventional system shown in **Figure 2**. Except in scan operation, data on the master side is always propagated to the slave side; therefore, the output data from the slave side can be used as slightly delayed master output data. Since the timing of slave output of the power-saving latch differs

from that of regular latches, the power-saving latch cannot necessarily replace every regular latch. However, power-saving latches are adapted where appropriate in the SPARC64 V/VI.

2) Latch-IH control

As shown in **Figure 1**, the latch has an Inhibit (IH) pin to suppress internal clock propagation. The IH pin reduces the power consumption of a single latch circuit by 70% to 80% as shown in **Figure 2**. This also suppresses logical operation of the output circuit and therefore reduces the power consumption of adjacent circuits.

When the IH pin is enabled, the latch retains the original value. As a result, when the IH pin of all the latches is enabled, the CPU will not operate. There are many cases where the CPU need not operate, such as when data is held waiting for synchronization and when data need not be replaced until the next operation begins. In such cases, the IH pin of each circuit can be kept enabled. In the SPARC64 V/VI, IH-pin control applies to at least 50% of all the latches.

note 1) A circuit that has two stable binary states (0 and 1) receives one bit or several bits of input data at a clock edge, and holds the data until the next clock edge is detected.




Figure 2
Power consumption of circuits using latches.

3) RAM-CE control

The RAM macro^{note 2)} has a Chip Enable (CE) pin to regulate the activation of internal circuits. When data is to be retained but need not be written or read, the CE pin is turned off to save power. In the SPARC64 V/VI, the second cache circuit is divided into blocks, and each RAM macro circuit is only activated when necessary.

4) Chopper-IH control

A chopper circuit comprises the last part of the clock distribution circuits. Clock pulses narrowed by chopping are supplied to the latch. This chopper has an IH pin for suppressing operation. One chopper supplies clock pulses to more than one latch. When the latches need not be used, their clock pulses can be stopped at the same time, and the IH pin turned on to reduce power consumption. Unlike the latch-IH control described above, chopper-IH control cannot operate depending on individual logic conditions. There are also restrictions on the delay of signals sent to the IH-pin. Conversely, more effective

note 2) Static, volatile random access memory not requiring refresh operation.

power savings are made possible by chopper-IH control than by latch-IH control. For these reasons, chopper-IH control is sparingly used in the SPARC64 VI.

3.2 Reducing static power consumption

1) Adopting multiple threshold voltages (Multi-V_{th})

Semiconductor miniaturization has reduced both supply voltage (V_{dd}) and threshold voltage (V_{th}). V_{th} must be kept low to achieve high operating speed. However, a low V_{th} increases leakage current (i.e., current that flows between the source and drain even with no voltage applied between both). In contrast, a high V_{th} reduces leakage current as well as operating speed. To maintain both high performance and low power consumption, various transistors having different V_{th} values should therefore be used according to the circuit conditions. The SPARC64 V/VI employs transistors having three types of V_{th} for the logical circuits.

2) Reducing leakage in capacitor cells

The internal voltage of high-speed LSIs

varies locally according to operation. Arranging capacitor cells in the LSI can prevent such variations in voltage. Moreover, the gate oxide thickness must be reduced to increase capacitor cell capacity. In other words, the thinner the gate oxide, the greater the flow of gate leakage (i.e., leakage current caused by electrons passing through the gate oxide via the tunnel effect). Therefore, to reduce the gate static power consumption, the arrangement area of capacitor cells must be widened, with a thick gate oxide used for each capacitor cell. In the SPARC64 V, 78% of the capacitor cells on the chip are the low-leakage type. In the SPARC64 VI, 90% of the capacitor cells on the chip are the low-leakage type.

3.3 Power saving by use conditions: supply voltage V_{dd} control of each chip

There are process variations in the manufacture of semiconductors. Advances made in miniaturization increase such variations. Chips operating at higher speeds generally have higher static power consumption, while chips operating at lower speeds tend to have lower static power consumption. The supply voltage exponentially affects static power consumption, and directly determines the operating speed. Therefore, setting the supply voltage uniquely to a chip is effective in maintaining high operating speed and low power consumption. For example, the operating voltage can be kept low to reduce both static power and dynamic power for a chip that has high static power consumption and a sufficient margin of operating speed. The operating voltage can also be raised to increase the operating speed for a chip that has low static power consumption and an insufficient margin of operating speed. Consequently, a system that centrally manages information about each chip in the manufacturing process was constructed for the SPARC64 VI.

3.4 Other power-saving techniques

We also employed various power-saving

techniques relative to clock distribution, latch driving, and design techniques, in order to more easily acquire a higher frequency without unnecessary power consumption. These techniques are detailed in Reference 9).

4. Power consumption and analysis of SPARC64 V/VI

Figure 3 shows the power consumption of the SPARC64 V/VI. In this figure, MAX indicates the chip having the highest power consumption in the standard shipping range; TYP indicates the chip positioned at the center of manufacturing variations. These chips were installed in a system and a benchmark test was performed to measure the highest power consumption at maximum temperature in the use environment. This tested the worst-case scenario for power consumption.

The SPARC64 V has power consumption of 65 W, which is lower than that of other high-end processors. Thus, the SPARC64 V offers power consumption performance among the best in the world. The SPARC64 VI, however, consumes a much higher 120 W of power. Conversely, its dynamic power consumption is approx. 55 W or about 1.4 times that of the SPARC64 V. Both the SPARC64 V and VI utilize 90-nm generation semiconductor techniques. The SPARC64 VI contains a dual CPU core developed based on enhancements to the SPARC64 V. Therefore, the power-saving techniques of the SPARC64 V are expanded when applied to the SPARC64 VI. Again, this is a worst-case scenario, since chips from the bottom-end of a manufacturing lot are tested under the most extreme temperature and workload conditions. Therefore, actual power consumption is much lower than shown in this figure.

Figure 4 shows the simulated results of the power consumption components of the SPARC64 VI chip positioned at the center of manufacturing variations. Even at the highest power consumption, the logic power consump-




Figure 3
Power consumption of SPARC64 V/VI.




Figure 4
Power consumption components of SPARC64 VI.

tion is 11% and latch power consumption 17% thanks to the low operation rate. Conversely, RAM power consumption is somewhat high (at 26%), although half is simulated without using RAM-CE control.

5. Future research

Power-saving results depend on actual operating conditions. Therefore, greater power savings must be achieved as a design target by checking the actual results.

Reducing the junction temperature during operation can effectively reduce static power consumption that accounts for at least half of the total power consumption of the SPARC64 VI. Moreover, a balance between device power consumption and the installation environment must be carefully considered.

At the same time, we will continue enhancing our multi-core chips and the system-on-a-chip (with system functions built into processor chip capabilities). The power-saving performance of the server unit will also be improved by utilizing a high-performance system consisting of a smaller quantity of chips.

6. Conclusion

This paper described the characteristics of semiconductor power consumption and the power-saving techniques employed in the SPARC64 V/VI processors. Computational

power affects human life through such equipment and devices as PCs, servers, automobiles, home electric appliances, and cellular telephones. The degree of semiconductor integration is also expected to increase in the future. Semiconductors will continue being the preferred platform for innovation, and are clearly superior to all other industrial products.

Processors, on the other hand, have the highest power consumption per unit volume as compared with other industrial products. Although future increases in computational power are expected, increases in power consumption pose a serious problem that may affect the future of the earth and human life.

The techniques described in this paper may appear to be an accumulation of simple efforts. However, the high computational power being provided by silicon technology will not be easily replaced by another form. Silicon technology will become increasingly important in the future, and we must continue conducting research and development of processors to accumulate even more useful technical knowledge.



Aiichiro Inoue, *Fujitsu Ltd.*

Mr. Inoue received the B.S. degree in Mechanical Engineering from the University of Tokyo in 1980, and then joined Fuji Photo Film Co., Ltd. In 1983, he moved to Fujitsu Ltd., where he was engaged in development of M-series mainframes. He has been involved in the CPU designs of most Fujitsu and Amdahl mainframes. Since 2000, he has also been engaged in

development of SPARC-based microprocessors.

References

- 1) A. Inoue: SPARC64 V Processor for UNIX Servers. (in Japanese), *FUJITSU*, **53**, 6, p.450-455 (2002).
- 2) A. Inoue: Fujitsu's new SPARC64™ V for Mission Critical Servers. Microprocessor Forum, October 15, 2002.
- 3) T. Maruyama: SPARC64™ VI: Fujitsu's Next Generation Processor. Microprocessor Forum, October 14, 2003.
- 4) A. Inoue: SPARC64™ V/VI for Mission-Critical Servers. Fall Processor Forum, October 5, 2004.
- 5) T. Maruyama: SPARC64™ VI/VI + Next Generation Processor. Fall Microprocessor Forum, October 25, 2005.
- 6) A. Inoue: SPARC64™ VI: A State of the Art Dual-Core Processor. Fall Microprocessor Forum, October 10, 2006.
- 7) M. Sakamoto, A. Katsuno, A. Inoue, T. Asakawa, H. Ueno, K. Morita, and Y. Kimura: Microarchitecture and Performance Analysis of a SPARC-V9 Microprocessor for Enterprise Server Systems. Proceedings of HPCA9, February 2003, p.141-152.
- 8) A. Inoue: High-performance, High-reliability Technology of SPARC64 V/VI. Scientific System Users Group, October 31, 2006.
- 9) H. Ando, Y. Yoshida, A. Inoue, I. Sugiyama, T. Asakawa, K. Morita, T. Muta, T. Motokurumada, S. Okada, H. Yamashita, Y. Satsukawa, A. Konmoto, R. Yamashita, and H. Sugiyama: A 1.3 GHz Fifth Generation SPARC64 Microprocessor. ISSCC2003, February 11, 2003.