An Automated Runtime Power-Gating Scheme

Mototsugu Hamada, Takeshi Kitahara, Naoyuki Kawabe, Hironori Sato, Tsuyoshi Nishikawa,
Takayoshi Shimazawa, Takahiro Yamashita, Hiroyuki Hara, and Yukihito Oowaki
Toshiba Corp.
Toshiyuki Furusawa
*Toshiba Microelectronics Corp.*

Abstract

An automated runtime power-gating scheme to reduce the leakage power in the active mode is presented in this paper. We propose a circuit that generates a sleep control signal from a clock-gating control signal automatically. By the combination of selective MT-CMOS scheme, the generated sleep control signal, and a novel flip-flop circuit with an additional latch function, a zero-wait transition from a sleep mode to an active mode is enabled. The additional latch function required for the zero-wait transition is achieved by only 6 transistors in addition to a conventional flip-flop. By the scheme, any design with the clock-gating scheme can be transformed automatically to a power-gated design while keeping the system operation the same in terms of the cycle accuracy. The scheme is applied to an MPEG4/H.264 audio/video codec and 21% power saving is achieved in the active mode while keeping the area overhead only 16% in a 90nm CMOS design.
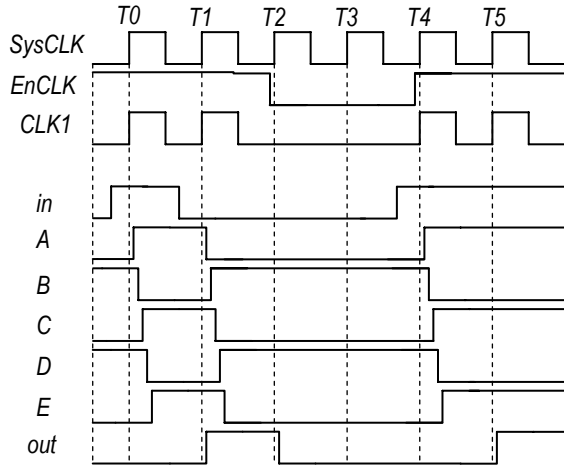
1. Introduction

Power reduction is one of the most critical issues with recent CMOS scaling. Especially on 90nm CMOS node and beyond, DC power reduction should be addressed as well as AC power reduction since the off-leak current of the transistors is becoming a dominant part of power dissipation[1]. Many techniques to reduce the leakage power have been proposed such as, MT-CMOS and VTCMOS[2][3]. While they are useful for alleviating the leakage issue when the system is idle, they are helpless when the system is active. When the DC power is comparable to or more than AC power, the issue is to reduce the DC power when the system is active. To reduce the active DC power, minimum and necessary parts of the circuits are to be made active for minimum and necessary duration. To achieve that, two types of granularity of the power-gating techniques have to be small, i.e. spatial granularity and temporal granularity. For example, MT-CMOS has a large switch transistor to shut down a certain functional block. The switch is on when the block is supposed to be active, and it is off when the block is idle. Generally the power domain is large and the control of the switch is done at the application level[4]. Therefore it takes some time to change a certain power domain from a sleep mode to an active mode and vice versa. By definition, the spatial granularity and the temporal granularity of conventional MT-CMOS are both coarse. Therefore MT-CMOS is not suitable for reducing "active" DC power.

Selective MT-CMOS scheme utilizes a small power switch in each high-performance cell on the critical path in multi-Vt designs[5]. In Reference [6], clock-gating control signal is used to shut down a certain power domain in combination with selective MT-CMOS scheme. However the scheme still requires several clock cycles to switch between active and sleep modes. In that sense, the cycle accuracy does not hold and therefore the temporal granularity is coarse though the spatial granularity may be fine.

Another scheme to employ the clock-gating control signal as an power-gating signal was presented in [7]. They make a certain part of combinational circuits, which they call "sleepy gates," shut down while the clock-gating control signal is low. However, the problem is that timing violations arise when "sleepy gates" wake up. It is because the clock-gating control signal itself is sometimes on a critical path and therefore the power switch controlled by the signal prolongs the critical path. They claim that the scheme is applicable only for BURN-IN TEST mode where the timing is not an issue. It is usable if some waits are allowed between active and sleep modes. They also presented a fine-grained power-gating scheme in [8]. However, since they employed a common "Enable" signal for clock-
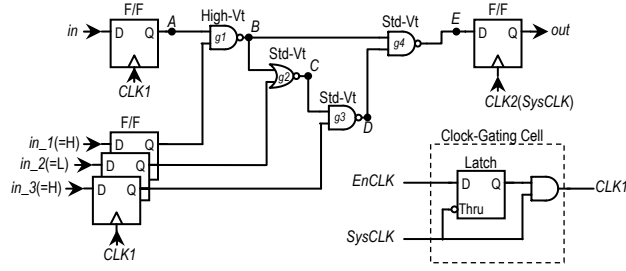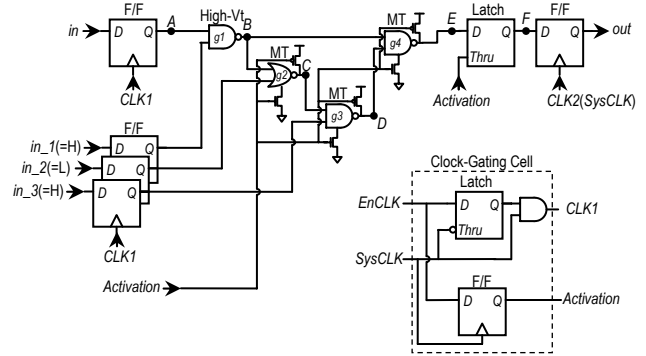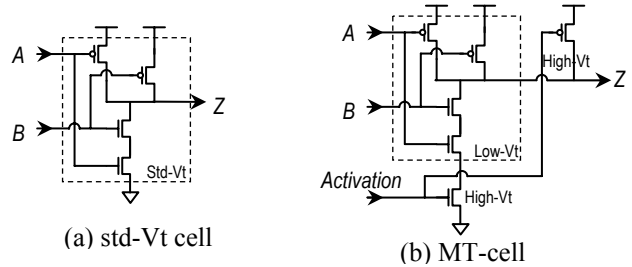
Figure 1. A conventional clock-gating design.

Figure 2. An operation waveform of the clock-gating design shown in Figure 1.

Figure 3. The proposed power-gating design.

(a) std-Vt cell   (b) MT-cell

Figure 4. Schematics of (a)std-Vt cell and (b)MT-cell of 2-input NAND.

Figure 5. An operation waveform of the power-gating design shown in Figure 3.

gating and power-gating, a zero-wait transition from the sleep mode to the active mode is impossible because "Enable" signal may change right before the rising edge of the system clock.   In the sense described above, the spatial granularity may be fine but the temporal granularity is still not fine enough in [7] and [8].
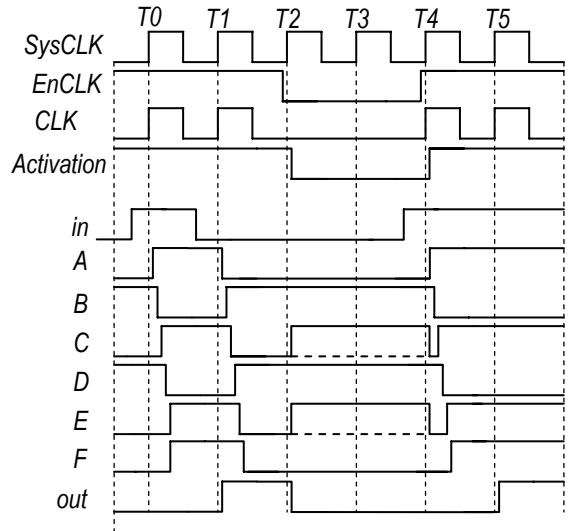
In this paper, we present a novel power gating scheme, in which both the spatial and temporal granularity are fine.   To achieve the fine spatial granularity, we adopt the selective MT-CMOS scheme. To achieve the fine temporal granularity we introduce an automated generation of power-gating control signals from the clock-gating control signals. On top of that, a novel flip-flop is also introduced. The flip-flop circuit having an additional latch function enables a zero-wait transition from/to a sleep mode to/from an active mode. The latch holds the output datum of the combinational logic circuits right before going into the sleep mode. When waking up, the flip-flop captures the datum stored in the latch at the first rising edge of the clock. Hence it can steal one clock-cycle waiting for the combinational circuits to wake up. From the second rising edge of the clock onward, the circuit operates as a conventional flip-flop. Owing to its zero-wait nature, any designs with the

clock-gating scheme can be transformed automatically to power-gating designs while keeping the system operation the same in terms of the cycle accuracy. Design methodology will also be shown. By applying this scheme to an audio/video codec in a 90nm CMOS design, 46% leakage reduction is achieved without a performance overhead.

2. Proposed Power Gating Scheme

The goal is to provide a simple design methodology that enables a power-gating design to reduce the active DC power while the RTL descriptions and the software environment of the SoC remain the same. To that end, the circuits made sleep have to wake up instantly after the wake-up instructions so that the system can operate with zero wait.

There are three enabling techniques; (1) cell-level power-gating technique with a pull-up pMOS transistor, (2) automatic generation of power-gating control signal, and (3) compact implementation of a flip-flop with an additional latch.

2.1 Conventional Gated-clock Design

Figure 1 shows a conventional design with the clock gating technique and Figure 2 shows its operation waveform. The design is a dual-threshold voltage (dual-Vt) design so that it can reduce the leakage current. The gate *g1* is a high-Vt cell while *g2*, *g3*, and *g4* are standard(std)-Vt cells to meet the timing constraints. For simplicity, the signals in_1, in_2, and in_3 are fixed to 1, 0, and 1, respectively.
Note that the flip-flops launching the input data to the combinational logic gates are triggered by the gated clock signal *CLK1*. The clock signal, *CLK1*, is gated for two cycles from *T2* to *T4*, which is generated by the clock-gating cell consisting of a latch and an AND gate.

2.2 Cell-Level Power Switch Insertion

Next, the design is converted into a power-gated design. Figure 3 shows the proposed power-gating design. The std-Vt cells are replaced by cells consisting of a low-Vt transistors with a high-Vt power switch, namely multi-threshold cells(MT-cells), where |low-Vt|<|std-Vt|<|high-Vt|. Figure 4 shows a schematic of a MT-cell for a 2-input NAND gate. Delay performance of the MT-cell is designed to be almost the same as its standard-Vt counterpart. It also has a pMOS pulling-up the output to VDD. Without this, the output may be in a high-impedance state in the sleep mode and therefore the gate in the next stage may draw a significant amount of current.

### 2.3 Power-Gating Control Signal

The issue is how to control the power switches in each MT-cell. We utilize the control signal of the clock gating system. The basic idea is that circuits do NOT need to be powered when the clock signal to the circuits is gated. However, for the sake of runtime power gating,

there is still an issue. As long as we use the clock-gating information for the power-gating, we cannot gate the power of the circuits before we know whether the clock is gated in the circuits, by the law of causality. The clock-gating control signal may toggle in the last minutes before the rising edge of the system clock as is the case with *EnCLK* signal shown in Figure 2. We cannot wake up the circuits in time for the rising edge of the clock if we start to power the circuits after we know the clock is not gated at the next rising edge, which is almost already there. Then, what we do is to synchronize the clock-gating control signal by the system clock to generate the control signal of the power switch, Activation, as shown in the clock-gating cell of Figure 3. Thus we can generate the signal automatically. Since the circuits are powered from the beginning of the clock cycle, we can avoid making the power-gating control signal on a critical path. However we have a different problem, that is, we are one-clock-cycle-behind.
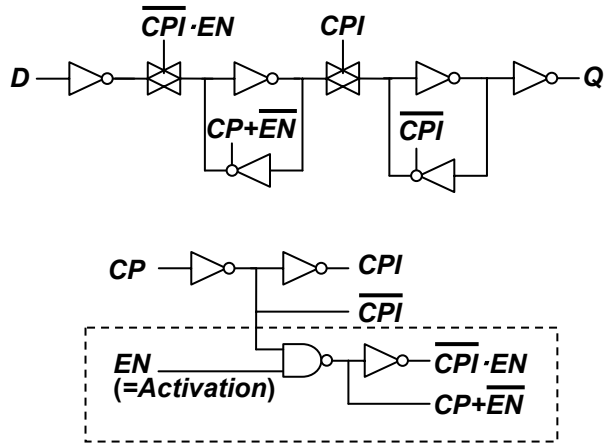
2.4 Flip-Flop with an Additional Latch

We insert an additional latch between node E and the input of F/F as shown in Figure 3 to solve the problem above. Figure 5 shows an operational waveform of the proposed power-gated design with the three enabling techniques into account. Since an MT-cell has a pull-up pMOS to fix the floating node issue, the power-gated combinational logic gates may change their output value when going into the sleep mode like nodes C and *E* in Figure 5, where the broken lines show the original values when the circuits are not power-gated. Note that the node *E* is the input of the flip-flop in the original clock-gating design in Figure 1. The additional latch holds the value to be input to the flip-flop while the circuits are in the sleep mode. At the rising edge of the clock right after the clock-gating control signal rises to high(*T4*), the combinational logic gates and the additional latch are still in the sleep mode. Therefore the data stored in the additional latch (node *F*) is made captured by the flip-flop at *T4*. Thereafter the latch is in the transparent mode. As we can see in Figures 2 and 5, the function of the proposed power-gated design is exactly the same as the conventional clock-gated design in the register transfer level.

The remained problem is an area overhead of the additional latch. Since a conventional flip-flop consists of two latches (master/slave latches), the addition of the inserted latch means the area penalty of almost 50% with respect to the area occupied by flip-flops. Therefore we prepare a new circuit for the novel flip-flop, which is depicted in Figure 6 with comparing a conventional flip-flop. The additional latch function is implemented by
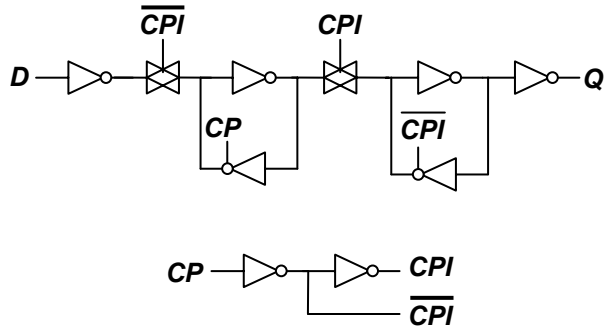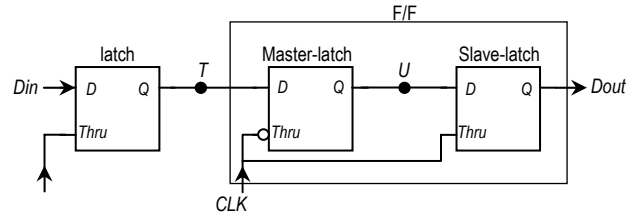
(a) Conventional F/F

(b) Proposed F/F

Figure 6. Schematics of (a) a conventional F/F and (b) the proposed F/F.

State	EN	CLK	Din	latch	T	M-latch	U	S-latch	Dout
(1)	H	H	Q	Thru	Q	Hold	R	Thru	R
(2)	H	L	Q	Thru	Q	Thru	Q	Hold	R
(3)	L	H	Q	Hold	R	Hold	S	Thru	S
(4)	L	L	Q	Hold	R	Thru	R	Hold	S

Figure 7. A block diagram and a truth table of a F/F with a latch.
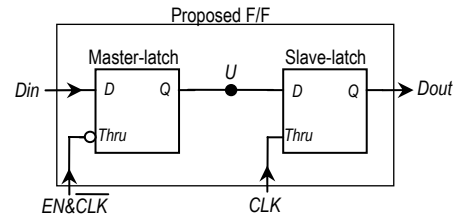
State	EN	CLK	Din	latch	T	M-latch	U	S-latch	Dout
(1)	H	H	Q	Thru	Q	Hold	R	Thru	R
(2)	H	L	Q	Thru	Q	Thru	Q	Hold	R
(3)	L	H	Q	Hold	R	Hold	S	Thru	S
(4)	L	L	Q	Hold	R	Hold	R	Hold	S

Figure 8. A block diagram and a truth table of the proposed F/F.

adding only 6 transistors. Moreover, since there is no additional transistor in the signal path, which means D-to-Q path, the delay of the proposed flip-flop is almost the same as the conventional one.

Next, the reason why the proposed flip-flop circuit has the same functions as a flip-flop with an additional latch will be described. Figure 7 is a block diagram of a flip-flop with an additional latch and its truth table. The point is that there are only two or three different values for four nodes, *Din*, *T*, *U*, and *Dout*. By utilizing this redundancy, we can save a latch by merging the front-end latch and the master latch. Figure 8 shows a block diagram of the new flip-flop and its truth table. The difference is the clock signal of the master latch. In the state (4), both of the latches in the new flip-flop can be in HOLD state where both *EN* and *CLK* are Low, achieving the three different values for three nodes.

2.5 Summary of the Proposed Scheme

Starting from a clock-gating design, we now have a power-gating design. Comparing the behavior of the clock-gating design shown in Figure 1 and the power-gating design shown in Figure 3, we can see two designs are exactly same in terms of the register transfer level functions. On the other hand, the leakage current of the std-Vt cells can be reduced by introducing the power-gating design.

3. Design Methodology

Figure 9 shows a design flow to implement the proposed scheme. First, an RTL description is synthesized to generate an initial netlist and placement by using std-Vt cells and high-Vt cells (dual-Vt design). Then std-Vt cells are to be replaced by MT-cells. However in the following two cases, the std-Vt cell is not replaceable. The first case is that the std-Vt cell has a fan-in propagated from the output of flip-flop whose clock signal is not gated(e.g. The cell g_A in Figure 10 cannot be replaced.). The other case is that the output of the cell is an asynchronous set/reset input of a flip-flop. (e.g.

385
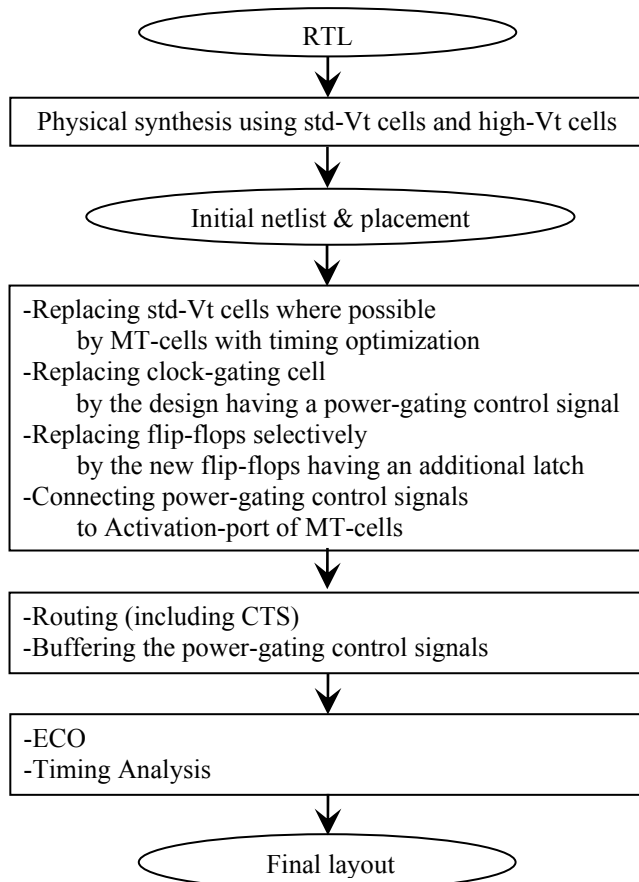
Figure 9. The design flow of the proposed power-gating design.
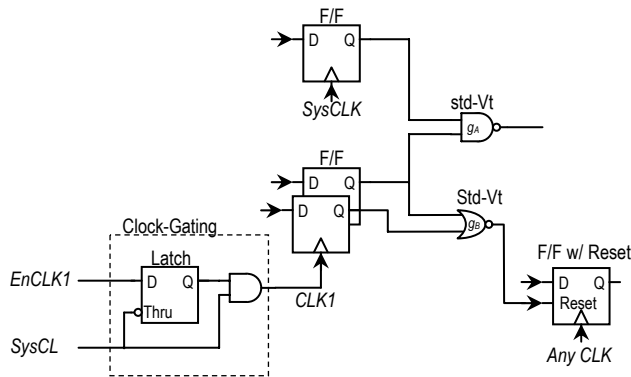

Figure 10. Non-replaceable cases.

The cell g_B in Figure 10 cannot be replaced.) In these cases, we may keep the cell as it is, accepting the leakage due to the cell. After the replacement is completed, the design is just a cell-based design. Therefore no major enhancement in CAD-tools is required to apply the scheme.

Table 1: Area comparison of the two designs.

	Core A	Core B	Core C	Core D	Total
Clock-Gating	0.178	0.174	0.091	0.557	1.000
Power-Gating	0.234	0.205	0.109	0.612	1.160
Area	31.1%	18.3%	19.4%	9.9%	16.0%

Table 2. Power ratio of the power-gating to the clock-gating.

	Core A	Core B	Core C	Core D	Total
AC	1.001	1.047	1.002	1.014	1.020
DC	0.424	0.617	0.631	0.567	0.536
Total	0.537	0.856	0.763	0.860	0.787
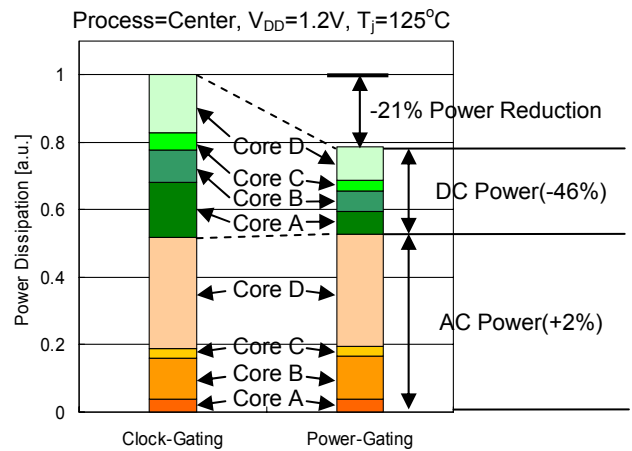
Figure 11. Power dissipation comparison.

For the power simulation, MPEG4 video encoding is carried out in both designs. The video size is VGA(640x480) and the frame rate is 30fps. All the four cores run at 150MHz. Simulation condition is 125°C and 1.2V supply. The process condition is nominal (center). Figure 11 and Table 2 show power dissipation comparison of the two designs. The proposed scheme enables 21% power reduction. As for AC power, the power-gating design is 0.1%-4.7% larger due to the same reason as the area overhead. On the other hand, DC power is reduced by 36.9%-57.6%. In this example, AC/DC ratio is still comparative. However, if the process is drifted to a fast corner, which in turn means larger DC current, DC power will be increased by a factor of 10 and will dominate the total power. In that case, the total power reduction(=AC+DC) would reach to 42%.

5. Conclusions

An automated runtime power-gating scheme is proposed. By the combination of three enabling techniques, such as, cell-level power-gating technique, automatic generation of power-gating control signal, and compact implementation of a flip-flop with an additional latch, the design can resume from the sleep mode without a performance penalty. Any designs with the clock-gating scheme can be transformed automatically to power-gated designs while keeping the system operation the same in terms of the cycle accuracy. By applying the scheme to an audio/visual codec SoC, 46% leakage reduction and 21% total power reduction are achieved in a 90nm CMOS process. The scheme is very effective for 90nm and beyond where DC power will dominate the total power rather than AC power.

6. References

[1] T. Sakurai, "Perspective on Power-Aware Electronics," *IEEE International Solid-State Circuits Conference (ISSCC2003)* Digest of Tech. Papers, pp26-29, Feb. 2003.

[2] S. Mutoh, T. Douseki, Y. Matsuya, T. Aoki, S. Shigematsu, and J. Yamada, "1-V Power Supply High-Speed Digital Circuit Technology with Multithreshold-Voltage CMOS," IEEE Journal of Solid-State Circuits, vol. 30, no. 8, pp847-854, Aug. 1995.

[3] T. Kuroda, T. Fujita, S. Mita, T. Nagamatsu, S. Yoshida, F. Sano, M. Norishima, M. Murota, M. Kako, M. Kinugawa, M. Kakumu, and T. Sakurai, "A 0.9V 150MHz 10mW 4mm2 2-D Discrete Cosine Transform Core Processor with Variable-Threshold-Voltage Scheme, " *IEEE International Solid-State Circuits Conference (ISSCC'96)* Digest of Tech. Papers, pp166-167, Feb. 1996.

[4] Y. Kanno, H. Mizuno, Y. Yasu, K. Hirose, Y. Shimazaki, T. Hoshi, Y. Miyairi, T. Ishii, T. Yamada, T. Irita, T. Hattori, K. Yanagisawa, and N. Irie, "Hierarchical Power Distribution with 20 Power Domains in 90-nm Low-Power Multi-CPU Processor," *IEEE International Solid-State Circuits Conference (ISSCC2006)* Digest of Tech. Papers, pp540-541, Feb. 2006.

[5] K. Usami, N. Kawabe, M. Koizumi, K. Seta, and T. Furusawa, "Automated selective multi-threshold design for ultra-low standby applications," *Proceedings of the 2002 International Symposium on Low Power Electronics and Design (ISLPED2002)*, pp.202-206, Aug. 2002

[6] G. Uvieghara, M-C. Kuo, J. Arceo, J. Cheung, J. Lee, X. Niu, R. Sankuratri, M. Severson, O. Arias, Y. Chang, S. King, K-C. Lai, Y. Tian, S. Varadarajan, J. Wang, K. Yen, L. Yuan, N. Chen, D. Hsu, D. Lisk, S. Khan, A. Fahim, C-L. Wang, J. Dejaco, Z. Mansour and M. Sani, "A Highly-Integrated 3G CDMA2000 1X Cellular Baseband Chip With GSM/AMPS/GPS/Bluetooth/Multimedia Capabilities And ZIF RF Support," *IEEE International Solid-State Circuits Conference (ISSCC2004)* Digest of Tech. Papers, pp422-423, Feb. 2004.

[7] K. Usami and H. Yoshioka, "A scheme to reduce active leakage power by detecting state transitions," *The 2004 47th Midwest Symposium on Circuits and Systems (MWSCAS'04)*, pp.493-496, July 2004.

[8] K. Usami and N. Ohkubo, "A Design Approach for Fine-grained Run-Time Power Gating using Locally Extracted Sleep Signals," *XXIV IEEE International Conference on Computer Design (ICCD2006)*, Oct. 2006.

[9] T. Fujiyoshi, S. Shiratake, S. Nomura, T. Nishikawa, Y. Kitasho, H. Arakida, Y. Okuda, Y. Tsuboi, M. Hamada, H. Hara, T. Fujita, F. Hatori, T. Shimazawa, K. Yahagi, H. Takeda, M. Murakata, F. Minami, N. Kawabe, T. Kitahara, K. Seta, M. Takahashi, and Y. Oowaki, "A 63-mW H.264/MPEG-4 audio/visual codec LSI with module-wise dynamic Voltage/frequency scaling,**"** *IEEE International Solid-State Circuits Conference (ISSCC2005)* Digest of Tech. Papers, pp132-133, Feb. 2005.