

An Oscillation-Based On-Chip Temperature-Aware Dynamic Voltage and Frequency Scaling Scheme in System-on-a-Chip

Katherine Shu-Min LI^{†a)}, Yingchieh HO^{††}, Yu-Wei YANG[†], Nonmembers, and Liang-Bi CHEN^{†*b)}, Member

SUMMARY The excessively high temperature in a chip may cause circuit malfunction and performance degradation, and thus should be avoided to improve system reliability. In this paper, a novel oscillation-based on-chip thermal sensing architecture for dynamically adjusting supply voltage and clock frequency in System-on-a-Chip (SoC) is proposed. It is shown that the oscillation frequency of a ring oscillator reduces linearly as the temperature rises, and thus provides a good on-chip temperature sensing mechanism. An efficient Dynamic Voltage-to-Frequency Scaling (DF2VS) algorithm is proposed to dynamically adjust supply voltage according to the oscillation frequencies of the ring oscillators distributed in SoC so that thermal sensing can be carried at all potential hot spots. An on-chip Dynamic Voltage Scaling or Dynamic Voltage and Frequency Scaling (DVS or DVFS) monitor selects the supply voltage level and clock frequency according to the outputs of all thermal sensors. Experimental results on SoC benchmark circuits show the effectiveness of the algorithm that a 10% reduction in supply voltage alone can achieve about 20% power reduction (DVS scheme), and nearly 50% reduction in power is achievable if the clock frequency is also scaled down (DVFS scheme). The chip temperature will be significant lower due to the reduced power consumption.

key words: DVFS, measurement and simulation of multiple-processor (multicores, MPSoC) systems, real-time distributed systems, reliability, thermal effect control scheme

1. Introduction

With rapid advance in the VLSI technology, the continuous device scaling leads to tremendous increase in device density and circuit speed. As a result, the high power consumption and power density become important design issues, as high power density raises temperature and may cause overheating [1]. The problem with high working temperature is that it reduces carrier mobility and threshold voltage; thus, it has a significant impact on circuit performance.

In addition to performance degradation, very high temperature also affects circuit reliability. A region with excessive heat is referred to as a hot spot. Hot spots may cause transient and/or permanent faults, which lead to incorrect operations. Various methods to deal with temperature-induced problems have been proposed, including software solution [2] and thermal-aware design techniques [3], [4].

On the other hand, a mechanism that can dynamically change power consumption so as to avoid overheating is very useful to improve system reliability and yield without elaborated design effort.

Traditionally, voltage scaling is the most effective way to combat rise in power, as both static (leakage) power and dynamic power in CMOS circuits will be reduced as the supply voltage is lower,

$$P_{static} = I_{leakage} \times V_{DD} \quad (1)$$

$$P_{dynamic} = k \times C_L \times V_{DD}^2 \times f_p \quad (2)$$

where C_L and f_p are load capacitance and operation frequency, respectively.

As it is more difficult to scale down supply voltage in the future, more sophisticated schemes have to be applied to achieve required system performance with low power consumption. In the dynamic voltage and frequency scaling (DVFS) technique [5], [6], the supply voltage level of the CPU is dynamically scaled to meet system computation demand with just enough circuit speed. Many modern microprocessors are designed with the DVFS functionality [7], [8].

With dynamically scalable voltage levels, it is possible to deal with overheating problem with on-chip thermal sensing mechanism. Whenever the chip temperature is higher than a given threshold, the supply voltage is changed to a lower level to reduce power consumption so that chip temperature will be lower. On the other way, at normal or lower temperature, a high supply voltage can be used to provide higher circuit performance. In this approach, it is easier to handle high chip temperature due to heavy computation in a short time without sacrificing overall system performance. Our approach is to apply the above dynamical voltage adjustment according to chip temperature.

In this paper, we propose a thermal-aware dynamic scaling framework in SoC, in which supply voltage and frequency can be dynamically adjusted to achieve good performance at restricted temperature level. First, we provide a simple *oscillator-based thermal sensor* that can be used to conduct on-chip temperature measurement.

Many different thermal sensors have been proposed [9]–[11]. In this paper, our goal is to provide a real-time on-chip thermal management mechanism for SoC architecture that can dynamically and adaptively adjust chip temperature. The temperature measurement is achieved with oscillator-based sensors that can be easily embedded

Manuscript received December 4, 2013.

Manuscript revised April 9, 2014.

[†]The authors are with the Department of Computer Science and Engineering, National Sun Yat-Sen University, Kaohsiung 80424, Taiwan.

^{††}The author is with the Department of Electrical Engineering, National Dong-Hwa University, Hualien, Taiwan.

*Presently, also with BXB Electronics Co., Ltd., Kaohsiung 80673, Taiwan.

a) E-mail: smli@cse.nsysu.edu.tw

b) E-mail: liangbi.chen@gmail.com

DOI: 10.1587/transinf.2013LOP0016

into cores. It is shown that this sensor provides good linearity [9] in the temperature range of interest, and thus provides good temperature estimation.

The sensors can be embedded into cores that may induce high power density and become hot spots. The sensor outputs are checked by an on-chip Dynamic Voltage Scaling or Dynamic Voltage and Frequency Scaling (DVS or DVFS) monitor to see if it is necessary to change the supply voltage level.

With a reduced supply voltage, both dynamic and static power can be reduced, and the chip temperature will be lower as a result. In the DVFS scheme, clock frequency can also be adjusted, the dynamic power can be further reduced, as Eq. (2) suggests. As a result, the working temperature of a chip can be adjusted on-line, which significantly improves the circuit reliability, performance, and yield.

The feasibility of the proposed method is evaluated through spice simulation of oscillator-based thermal sensors with TSMC 0.13 μ m and Berkeley Predictive Technology Model (BPTM) technology (ie. 90nm, 65nm, 45nm). Experimental results on SoC benchmark circuits show that a 10% reduction in supply voltage alone can achieve about 20% power reduction (DVS scheme), and nearly 50% reduction in power is achievable if the clock frequency is also scaled down (DVFS scheme). The chip temperature will be significant lower due to the reduced power consumption. Therefore, the proposed thermal sensor and the DVFS scheme together can provide an efficient on-chip temperature control mechanism.

This paper organizes as follows. Section 2 describes the proposed thermal management mechanism, and Sect. 3 presents the design of on-chip thermal sensor. Section 4 gives the experimental results, and Sect. 5 concludes this work.

2. Global Architecture

This section describes the proposed temperature management framework. This mechanism is based on the linear relationship among temperature, frequency and power, as discussed in Sect. 2.1. The proposed SoC architecture is presented in Sect. 2.2.

2.1 Temperature-Aware Dynamic Scaling

The basic idea behind the proposed temperature-aware dynamic scaling framework is shown in Fig. 1. A ring oscillator (RO) is used as a thermal sensor. The propagation delay of a logic gate increases as the ambient temperature becomes higher and thus the oscillation frequency (f_{osc}) of the ring oscillator will be lower.

The oscillator output is fed to a counter, which is enabled for a fixed period of time once it is necessary to measure the working temperature. If the counter content is smaller than a predetermined value, which means f_{osc} is lower than a given threshold, the chip is overheating and a reduced level of power consumption is required to cool

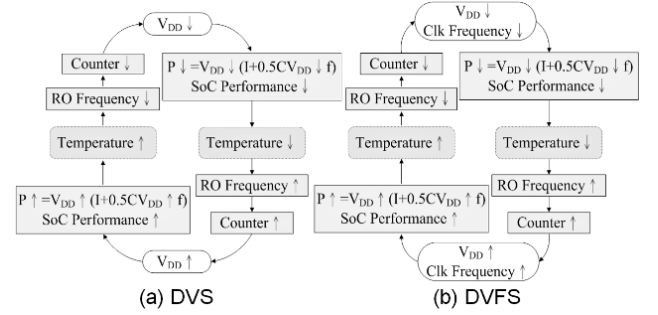


Fig. 1 Overview of the proposed method with ring oscillators as thermal sensors: (a) DVS (Dynamic Voltage Scaling), and (b) DVFS (Dynamic Voltage and Frequency Scaling).




Fig. 2 Relationship between temperature, frequency, and power in a 9-inverter ring oscillator.

down the chip. As indicated by Eq. (1) and (2), both leakage and dynamic power will be reduced as the supply voltage becomes lower, which in turn causes chip temperature to be lower.

At lower temperature, the gate delay is shorter, which renders the oscillation frequency f_{osc} higher. Once f_{osc} is high enough, indicated by a number in the counter that is larger than a given value, the supply voltage can be increased to provide a higher system performance. This scenario is illustrated in Fig. 1 (a) for Dynamic Voltage Scaling (DVS), if only V_{DD} is scaled, and Fig. 1 (b) for Dynamic Voltage and Frequency Scaling (DVFS), as both V_{DD} and clock frequency f are scaled. In general, gate delay at a lower V_{DD} is longer and thus the clock frequency has to be reduced as well.

The effectiveness of the thermal management methodology outlined in Fig. 1 depends on whether on-chip temperature measurement can be carried out easily and precisely. The ring oscillator is a good candidate to achieve this goal, as shown in Fig. 2, where the oscillation frequency decreases almost linearly as the temperature increases from 0°C to 125°C.

These data are obtained through SPICE simulation of a 9-inverter ring oscillator with TSMC 0.13 μ m technology, and the straight lines in Fig. 2 are regression lines. Also shown in the figures are power vs. frequency and power vs. temperature plots. Since the dynamic power is proportion to the frequency (Eq. (2)), it is not surprising that the power consumed by an oscillator grows lineally as the oscillation frequency f_{osc} increases. The relation between power consumption and temperature is also shown in the figure.

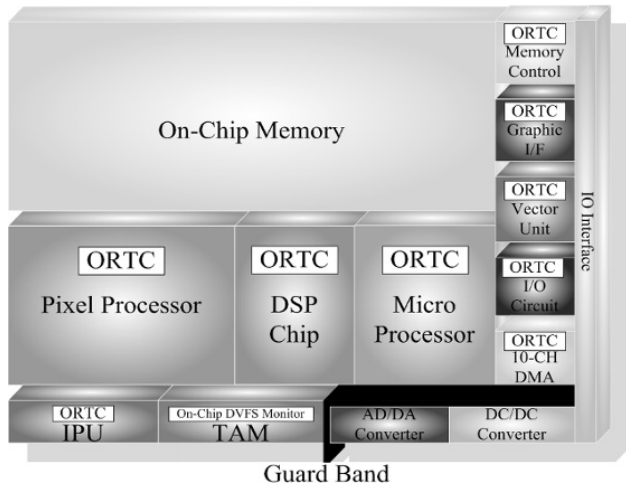


Fig. 3 System architecture of DVS/DVFS schemes for real-time frequency measurement under thermal effects and process variation to dynamically adjust V_{DD} .

2.2 SoC System Architecture for DVS and DVFS

In order to implement the proposed schemes, the Oscillation Ring based Thermal Sensor (ORTS) has to be embedded in all cores that may be overheating, as shown in Fig. 3. All the sensor results are collected in a centralized on-chip DVS or DVFS monitor, including a global reset control. Whenever a core with high temperature is detected, the supply voltage is scaled down to reduce power consumption and working temperature. The supply voltage and operation clock frequency is restored to higher level only if temperatures in all cores fall back to the normal level.

The sensors have to be activated periodically to make sure that the temperature will not be higher than a given level for an extended period of time. On the other hand, the thermal sensors themselves also induce power consumption and thus should not be activated too often. The exact period of temperature sampling is thus a trade-off between these two factors. The details of the ORTS design will be discussed in the next section.

3. Thermal Sensor Design

This section discusses how to design the thermal sensor and the on-chip DVS/DVFS monitor.

3.1 General Architecture of an Oscillation Ring Triggered Thermal Sensors (ORTS)

The sensor architecture is shown in Fig. 4. An odd number of inversions are connected into a feedback structure that can generate an oscillation signal, as shown in Fig. 4(a). The period of the oscillation signal is determined by the propagation delay in the loop, and the delay is affected by the ambient temperature. When the sensor is activated, the oscillation signal is used to trigger a counter for a predetermined period of time. Obviously, the content in the counter

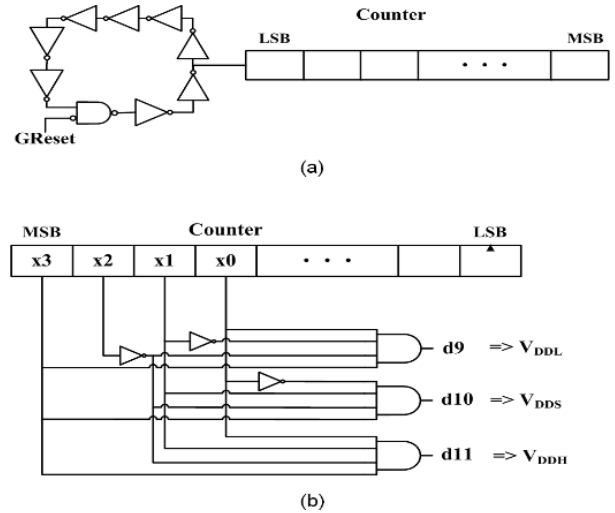


Fig. 4 An oscillation-based thermal sensor example.

is affected by the chip temperature. At a higher temperature, the propagation delay will increase, which makes the period of the oscillation signal longer and thus the counter will record a smaller number in a predetermined sampling period. The content of the counter is then decoded to select an appropriate supply voltage. The length of the counter is determined by the time period used to sample the oscillation signal as well as the maximum expected frequency. For example, if the sampling time is 0.1ms and the maximum frequency is 500MHz, a 16-bit counter should be enough.

In general, many different supply voltage levels can be used in a DVS and/or DVFS system. To simplify discussion, the design example shown in Fig. 4(b) assumes that three supply voltages are used: standard V_{DD} (V_{DDS}), high V_{DD} (V_{DDH} , which is 110% of V_{DDS}), and low V_{DD} (V_{DDL} , which is 90% of V_{DDS}). A priority decoder is used to select the supply voltage. When the oscillation frequency f_{osc} is too low, a very small value will appear in the counter and thus signal d_9 is asserted, which requests the on-chip DVS and/or DVFS monitor to reduce the supply voltage to V_{DDL} . On the other hand, if f_{osc} is too high, signal d_{11} is asserted to indicate that V_{DDH} can be used to provide better system performance for this core. Otherwise, d_{10} is asserted to request a standard V_{DDS} . In reality, signal d_{10} is not necessary, as we shall see later in this section.


3.2 Feasibility of ORTS Thermal Sensor under Process Variation

A ring oscillator [9], [10] is a feedback loop consisting of odd number of inverters. A ring oscillator with nine inverters is shown in Fig. 4. The NAND gate is inserted in the ring so that the ring can be disabled by a global reset signal (G_{Reset}). Once the reset signal is not asserted ($G_{Reset} = 0$), the odd number of inversion in the feedback loop will create an oscillation signal.

In order to consider process variation effect on the ring architecture, we have conducted Monte Carlo simulation on

# simulations	Max Freq. (MHz)	Min Freq. (MHz)	Max Power (μ W)	Min Power (μ W)
20	476	526	12.44	11.18
200	476	556	12.45	10.56
2000	476	556	12.46	10.56

(a)



(b)

Fig. 5 Monte Carlo simulation of the 9-inverter ring oscillator under process variation: (a) frequency and power range, (b) simulated waveforms for (1) 20, (2) 200, (3) 2000 runs.

a 9-inverter ring oscillator with TSMC 0.13 μ m technology. Each feature size is assumed to be Gaussian distributed with $3\sigma = 20\%$ of the nominal value, and 20, 200 and 2000 simulation runs are executed. The results are shown in Fig. 5. Figure 5 (a) provides the ranges of oscillation frequency and power, and the Fig. 5 (b) shows that the variation of oscillation frequency f_{osc} is very small even under $3\sigma = 20\%$ process variation.

3.3 On-Chip DVFS Monitor

The overall temperature-aware DVS/DVFS system works under the following assumptions.

1) Three supply voltage levels are used: $V_{DDL} < V_{DDS} < V_{DDH}$.

2) In a DVFS system, the corresponding clock frequencies are $F_{clkL} < F_{clkS} < F_{clkH}$.

3) For a ring oscillator, the oscillation frequency (f_{osc}) is used to measure the ambient temperature. Two threshold frequencies are used to indicate the temperature level: F_{oscL} and F_{oscH} . Note that f_{osc} decreases as the temperature rises, as shown in Fig. 2. Thus the condition $f_{osc} < F_{oscL}$ indicates a high temperature and thus requires a lower supply voltage (V_{DDL}) and clock frequency (F_{clkL}). On the other hand, the condition $f_{osc} > F_{oscH}$ indicates a low temperature and thus one may apply higher supply voltage (V_{DDH}) and faster clock (F_{clkH}) to boost performance. For $F_{oscL} < f_{osc} < F_{oscH}$, standard supply voltage (V_{DDS}) and clock (F_{clkS}) are used.

The threshold frequencies F_{oscL} and F_{oscH} are used to determine the threshold values in the counter in Fig. 4. The oscillation signal is used to trigger the counter for a given period of time, and let the counter content be n at the end of the sampling period. Obviously, a higher oscillation frequency creates a larger counted number n . Let the counter contents corresponding to F_{oscL} and F_{oscH} be N_L and N_H respectively. As a result, $n < N_L$ indicates a high temper-




Fig. 6 A typical DVS/DVFS instance.

ature and $n > N_H$ implies a low temperature. Therefore, the temperature detection can be realized by a priority decoder, as shown in Fig. 4 (b). When $n < N_L$ is true, signal d_9 is asserted to indicate that a low supply voltage V_{DDL} is requested. On the other hand, if $n > N_H$ holds, signal d_{11} is asserted to indicate that a high supply voltage V_{DDH} can be used. Otherwise, signal d_{10} is asserted. Please note that DVS scheme works in the same way as DVFS does, only without frequency scaling.

In this paper we assume a synchronous design. For simplicity, it is assumed that the supply voltage is the same in all cores. In other words, the supply voltage of the whole chip is reduced whenever one module is overheating.

A typical DVS/DVFS sampling process is illustrated in Fig. 6. The temperature sampling process is carried out at fixed time interval to ensure that the chip temperature is kept under control. The sampling process starts by deasserting GReset to enable the ring oscillator; the counter also has to be cleared before the sampling process. Once the counter starts counting, it is enabled for a given sampling period T_{period} . Assume that there are M cores in a chip, and let the d_9 , d_{10} , and d_{11} signals from core i be denoted as $d_{9,i}$, $d_{10,i}$, and $d_{11,i}$, respectively. The supply voltage of the whole chip is adjusted according to the decoder outputs of all modules as follows.

Let the global supply-voltage level control signals be D_9 , D_{10} , and D_{11} . These three signals can be determined according to the following Boolean equations.

$$D_9 = d_{9,1} + d_{9,2} + \dots + d_{9,M} \quad (3)$$

$$D_{11} = d_{11,1} \cdot d_{11,2} \cdot \dots \cdot d_{11,M} \quad (4)$$

$$D_{10} = \overline{D_9} \cdot \overline{D_{11}} \quad (5)$$

Note that the above equations hold as $d_{9,i}$, $d_{10,i}$, and $d_{11,i}$ form an 1-out-of-3 code for each i ; in other words, only one of the three bits is a logic 1. Thus, the supply voltage should be lowered to V_{DDL} if at least one core is overheating, as Eq. (3) shows. The supply voltage can be raised to V_{DDH} for better performance if the temperature in each and every core is low, as shown in Eq. (4). Otherwise, the standard supply voltage V_{DDS} is used.

As shown in Fig. 6, applying DVS/DVFS schemes involved idle time when the chip has to be stopped. The time penalty required to change supply voltage may actually degrade the overall system performance if the voltage changes occur too frequently. The frequency of power supply level

changes depends on the distribution of workload, the threshold F_{oscH} , and rate of heating/cooling once the supply voltage is raised/reduced. On the other hand, very high temperature may cause reliability problem and should be avoided; therefore, F_{oscL} should be selected according to the projected temperature profile.

3.4 General Design Procedure

The thermal management mechanism described in this section can be easily extended to a system with any number of supply voltage levels. The general design procedure for a system with L different supply voltage levels is outlined as follows.

Let the supply voltage levels be labeled as $V_{DD,1}, V_{DD,2}, \dots, V_{DD,L}$, with $V_{DD,1} < V_{DD,2} < \dots < V_{DD,L}$. In order to determine the appropriate supply voltage level of a core, the oscillation frequency of core i , denoted as $f_{osc,i}$ is divided into L ranges by $L-1$ threshold frequencies $F_{osc,1} < F_{osc,2} < \dots < F_{osc,L-1}$. The boundary frequencies $F_{osc,0}$ and $F_{osc,L}$ are assumed to be 0 and ∞ , respectively. When $F_{osc,j-1} < f_{osc,i} < F_{osc,j}$ is true, the highest supply voltage that can be used in core i is the j -th supply level $V_{DD,j}$.

As we have seen in this section, the range of $f_{osc,i}$ can be determined by comparing the content of the i -th counter ni with respect to a set of known threshold values $N_1 < N_2 < \dots < N_{L-1}$, where N_j is the counter content when the oscillator frequency f_{osc} is equal to $F_{osc,L}$. According to the counter content, L signals $d_{1,i}, d_{2,i}, \dots, d_{L,i}$, which form a 1-out-of- L code, are decoded for core i . If $d_{j,i}$ is 1, the acceptable voltage levels to core i will be $V_{DD,1}, V_{DD,2}, \dots, V_{DD,j}$.

The global supply-voltage level control signals are D_1, D_2, \dots, D_L , where D_i indicates supply level $V_{DD,i}$ should be used. The global signals are determined by the following equations. In these equations, the summation and product represent logic OR and AND operations, respectively.

$$D_1 = \sum_{i=1}^M d_{1,i} \quad (6)$$

$$D_j = \left\{ \prod_{l=1}^{j-1} \overline{D_l} \right\} \cdot \left\{ \sum_{i=1}^M d_{j,i} \right\}, \text{ where } 1 < j \leq L \quad (7)$$

4. Experimental Results

The effectiveness of the proposed thermal management mechanism is verified by SPICE simulation, and the results are presented in this section.

4.1 Feasibility of ORTS Thermal Sensor

The thermal sensor is the most important part of the temperature measurement mechanism. Therefore, it is very important to ensure that the sensor design is useful for various structures and technologies.

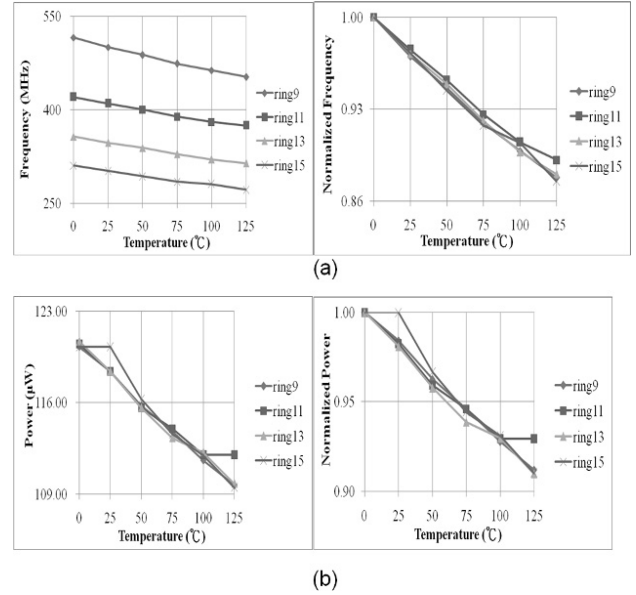


Fig. 7 Simulation of various ring lengths: (a) frequency vs. temperature, (b) power vs. temperature.

4.1.1 Effects of Ring Lengths

It is shown in Fig. 2 that the oscillation frequency and power consumption change almost linearly as the temperature increases in a 9-inverter ring oscillator. In general, the linear relation holds for all ring oscillators, as shown in Fig. 7, where rings consisting of 9, 11, 13 and 15 inverters are SPICE simulated with TSMC 0.13 μ m technology.

Figure 7 shows that the linear relationship exists in all rings. Longer rings have larger delays and thus lower frequencies, as shown in the left column of Fig. 7 (a). However, the percentage of the frequency with respect to the temperature is almost the same in all rings, as shown in the ring column of Fig. 7 (a). The relationship between power consumption and temperature is illustrated in Fig. 7 (b).

4.1.2 Effects of Process Variation

The performance of a ring oscillator is subject to the effect of process variation. Thus, it is important to show that the linear relationship among temperature, frequency, and power is not adversely affected by process variation.

In order to show how process variation affects the thermal sensor, we have conducted Monte Carlo simulation on a 9-inverter ring oscillator with TSMC 0.13 μ m technology. As in Sect. 3.2, each feature size is assumed to be Gaussian distributed with $3\sigma = 20\%$ of the nominal value, and 20 simulation runs are executed. The results are shown in Fig. 8. It is evident that the variation of oscillation frequency f_{osc} and power consumption is very small even under $3\sigma = 20\%$ process variation.

4.1.3 Effects of Technology Change

The linear relationship between frequency and temperature




Fig. 8 Relationship among temperature, frequency, and power for the 9-inverter ring oscillator under process variation.




Fig. 9 Frequency vs. temperature of the 9-inverter ring oscillator in various process technologies.

also holds in various technologies, as shown in Fig. 9, where the performance of the 9-inverter ring oscillator is plotted for 90, 65, and 45nm technologies using BPTM (Berkeley Predictive Technology Model). The oscillation frequency of the ring oscillator always decreases as the temperature rises no matter which technology is used; as a result, the proposed sensor design is still applicable as processing technology moves forward.

4.2 DVS/DVFS Schemes

In order to show that voltage and frequency scaling can effectively reduce power consumption and temperature, we verify the power and temperature distribution in five SoC test cases through simulation, and the results are presented in this section. The test cases are synthetic SoC circuits applied with various workloads among cores. The first two test cases (SoC1, SoC2) are multi-core systems where each core is an EDK3.2 compatible Microblaze core. The next two test cases (SoC3, SoC4) are identical multi-core systems where each core is an 8051 processor. However, different workloads are executed. The last test case is a multi-core system with all digital phase-lock loops (ADPLL). All core designs are available from OpenCores; the number of gates of the

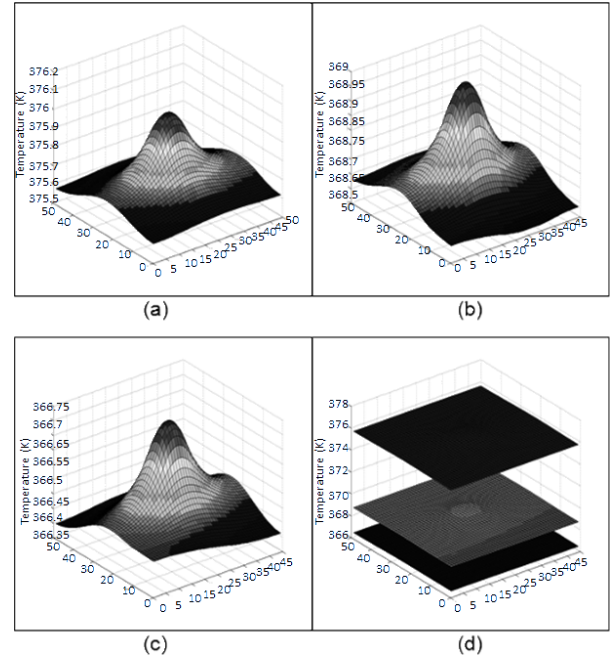


Fig. 10 Thermal profiles of an SoC circuit with: (a) V_{DDH} (1.98V), (b) V_{DDS} (1.8V), (c) V_{DDL} (1.62V), (d) three voltages under DVS scheme.

five designs are 6716250, 4298400, 1876800, 1876800, and 581250, respectively.

The circuits are designed with TSMC 0.18 μ m technology with 1.8 V standard supply voltage (V_{DDS}). The power, delay, and temperature are estimated by Synopsys PrimePower, Synopsys PrimeTime, and HotSpot [12], respectively. Two sets of experiments have been conducted: (1) only supply voltage is scaled (DVS), and (2) both supply voltage and clock frequency are scaled (DVFS).

In the DVS scheme, the clock frequency f_{clk} is fixed, with the period large enough to accommodate the worst-case delay in the highest acceptable temperature. In the DVFS scheme, the clock frequency is also scaled to achieve the best performance, with the highest clock frequency f_{clkH} so that the chip can work properly.

In order to show that supply voltage scaling indeed can reduce chip temperature, the temperature profiles of an SoC circuit with three different supply voltages are shown in Fig. 10. The temperature distributions on chip surface for the three different supply voltages are illustrated in Fig. 10(a), Fig. 10(b) and Fig. 10(c), respectively, while in Fig. 10(d) they are compared in the same figure.

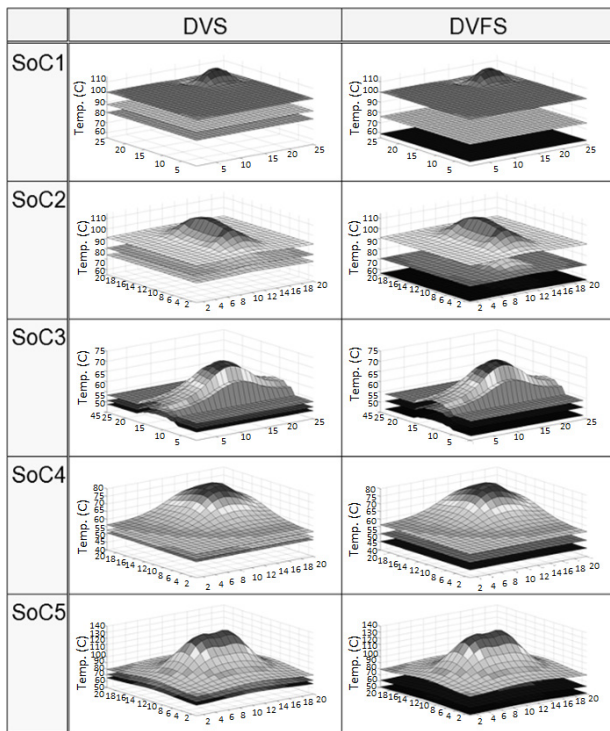
The effects of DVS and DVFS schemes on temperature and power are illustrated in Fig. 11 and Fig. 12. In both figures, five circuit examples are used for illustration. In the left column of Fig. 11, the temperature distributions on the chip surface are shown for the three different supply voltages without frequency scaling (DFS), while in the right column it is assumed that the clock frequency is also scaled proportionally (DVFS). Similarly, the power distribution is shown in Fig. 12. It can be seen that chip temperature and power consumption are both reduced as the supply voltage level

Table 1 Temperature and Power of Dynamic Voltage Scaling (DVS).

Circuit	Average Temperature (°C)			Average Core Power (mW)		
	V_{DDH}	V_{DDs}	V_{DDL}	V_{DDH}	V_{DDs}	V_{DDL}
SoC1	100.49	88.34	80.33	417.12	333.14	277.71
normalized	125.10%	109.97%	100.00%	150.20%	119.96%	100.00%
SoC2	94.25	83.35	76.17	449.45	358.97	299.24
normalized	123.74%	109.43%	100.00%	150.20%	119.96%	100.00%
SoC3	56.78	53.44	51.13	11.05	8.82	7.3
normalized	111.05%	104.52%	100.00%	151.37%	120.82%	100.00%
SoC4	60.92	56.41	54	41.38	32.39	27.59
normalized	112.81%	104.46%	100.00%	149.98%	117.40%	100.00%
SoC5	95.43	83.81	76.32	102.13	80.66	66.82
normalized	125.04%	109.81%	100.00%	152.84%	120.71%	100.00%
Comp.	120.69%	108.11%	100.00%	150.46%	119.94%	100.00%


Table 2 Experimental Results on Temperature Distribution Improvement due to Dynamic Voltage Scaling (DVS).

Circuit	Peak Temperature (°C)			Bottom Temperature (°C)			Temperature Range (°C)			Average Temperature (°C)		
	V_{DDH}	V_{DDs}	V_{DDL}	V_{DDH}	V_{DDs}	V_{DDL}	V_{DDH}	V_{DDs}	V_{DDL}	V_{DDH}	V_{DDs}	V_{DDL}
SoC1	113.10	98.50	88.80	97.40	85.90	78.30	15.70	12.60	10.50	100.49	88.34	80.33
SoC2	112.70	98.10	88.40	91.30	81.00	74.20	21.40	17.10	14.20	94.25	83.35	76.17
SoC3	72.00	65.60	61.20	53.30	50.60	48.80	18.70	15.00	12.40	56.78	53.44	51.13
SoC4	78.00	69.80	65.40	55.50	52.10	50.40	22.50	17.70	15.00	60.92	56.41	54.00
SoC5	147.20	124.70	110.20	79.60	71.30	65.90	67.60	53.40	44.30	95.43	83.81	76.32
Comp.	126.33%	110.31%	100.00%	118.73%	107.34%	100.00%	151.35%	120.12%	100.00%	120.69%	108.11%	100.00%

**Fig. 11** Temperature distribution under DVS/DVFS.

becomes lower in the DVS scheme, while the reduction is more significant in the DVFS scheme. Obviously, the DVFS scheme is more effective than DVS scheme.

Tables 1, 2, and 3 present the experimental results of the DVS scheme. Table 1 provides the general information in the five circuits, including average temperature (°C) and average power in each core. Data shown in the gray area are the normalized comparison. Table 2 gives details about the temperature distribution on chip surface under these three

**Fig. 12** Power distribution under (a) DVS, (b) DVFS.

supply voltages. Table 3 shows the simulation results of the power distribution. It can be seen that the power consumed under V_{DDH} and V_{DDs} are about 1.5x and 1.2x of the power consumed under V_{DDL} . In other words, if we raise the supply voltage from V_{DDs} (1.8 V) to V_{DDH} (1.98 V) (i.e., a 10% increase in supply voltage), the power consumption increases about 25%. On the other hand, if the supply voltage is reduced from V_{DDs} to V_{DDL} (1.62 V) (i.e., a 10% decrease in supply voltage), the reduction in power consumption is about 20%. The change in temperature range due the sup-

Table 3 Experimental Results on Power Distribution Improvement for Dynamic Voltage Scaling (DVS).

Circuit	Peak Power (mW)			Bottom Power (mW)			Power Range (mW)			Average Power (mW)		
	V _{DDH}	V _{DDs}	V _{DDL}	V _{DDH}	V _{DDs}	V _{DDL}	V _{DDH}	V _{DDs}	V _{DDL}	V _{DDH}	V _{DDs}	V _{DDL}
SoC1	826.30	660.10	550.20	227.20	181.50	151.30	599.10	478.60	398.90	417.12	333.14	277.71
SoC2	1032.90	825.10	687.70	413.20	330.00	275.10	619.70	495.10	412.60	449.45	358.97	299.24
SoC3	35.18	28.10	23.24	8.80	7.02	5.81	26.39	21.07	17.43	11.05	8.82	7.30
SoC4	162.37	127.12	108.28	29.00	22.70	19.34	133.38	104.42	88.94	41.38	32.39	27.59
SoC5	596.40	471.00	390.20	59.64	47.10	39.02	536.76	423.90	351.18	102.13	80.66	66.82
Comp.	150.78%	119.99%	100.00%	150.40%	119.93%	100.00%	150.93%	120.02%	100.00%	150.46%	119.94%	100.00%

Table 4 Temperature and Power of Dynamic Voltage and Frequency Scaling (DVFS).

Circuit	Average Temperature (°C)			Average Core Power (mW)		
	V _{DDH}	V _{DDs}	V _{DDL}	V _{DDH}	V _{DDs}	V _{DDL}
SoC1	100.49	76.31	59.17	417.12	249.95	131.54
	169.83%	128.97%	100.00%	317.11%	190.02%	100.00%
SoC2	94.25	72.56	57.21	449.45	269.32	141.73
	164.74%	126.83%	100.00%	317.12%	190.02%	100.00%
SoC3	56.78	52.56	47.67	11.05	8.25	5.01
	119.11%	110.26%	100.00%	220.56%	164.67%	100.00%
SoC4	60.92	53.51	46.85	41.38	26.61	13.36
	130.03%	114.22%	100.00%	309.73%	199.18%	100.00%
SoC5	95.43	70.61	55.41	102.13	56.27	28.19
	172.23%	127.43%	100.00%	362.29%	199.61%	100.00%
Comp.	153.16%	122.24%	100.00%	319.27%	190.85%	100.00%

Table 5 Experimental Results on Temperature Distribution Improvement due to Dynamic Voltage and Frequency Scaling (DVFS).

Circuit	Peak Temperature (°C)			Bottom Temperature (°C)			Temperature Range (°C)			Average Temperature (°C)		
	V _{DDH}	V _{DDs}	V _{DDL}	V _{DDH}	V _{DDs}	V _{DDL}	V _{DDH}	V _{DDs}	V _{DDL}	V _{DDH}	V _{DDs}	V _{DDL}
SoC1	113.10	83.90	63.20	97.40	74.50	58.20	15.70	9.40	5.00	100.49	76.31	59.17
SoC2	112.70	83.60	63.00	91.30	70.80	56.30	21.40	12.80	6.70	94.25	72.56	57.21
SoC3	72.00	63.90	54.60	53.30	49.90	46.10	18.70	14.00	8.50	56.78	52.56	47.67
SoC4	78.00	64.50	52.40	55.50	50.00	45.10	22.50	14.50	7.30	60.92	53.51	46.85
SoC5	147.20	99.10	69.70	79.60	61.90	51.00	67.60	37.20	18.70	95.43	70.61	55.41
Comp.	172.66%	130.41%	100.00%	146.90%	119.63%	100.00%	315.80%	190.26%	100.00%	153.16%	122.24%	100.00%

Table 6 Experimental Results on Power Distribution Improvement for Dynamic Voltage and Frequency Scaling (DVFS).

Circuit	Peak Power (mW)			Bottom Power (mW)			Power Range (mW)			Average Power (mW)		
	V _{DDH}	V _{DDs}	V _{DDL}	V _{DDH}	V _{DDs}	V _{DDL}	V _{DDH}	V _{DDs}	V _{DDL}	V _{DDH}	V _{DDs}	V _{DDL}
SoC1	826.30	495.10	260.60	227.20	136.20	71.70	599.10	358.90	188.90	417.12	249.95	131.54
SoC2	1032.90	618.90	325.70	413.20	247.60	130.30	619.70	371.30	195.40	449.45	269.32	141.73
SoC3	35.18	26.25	15.95	8.80	6.56	3.99	26.39	19.69	11.96	11.05	8.25	5.01
SoC4	162.37	104.44	52.42	29.00	18.65	9.36	133.38	85.79	43.06	41.38	26.61	13.36
SoC5	596.40	328.60	164.60	59.64	32.86	16.46	536.76	295.74	148.14	102.13	56.27	28.19
Comp.	323.84%	192.04%	100.00%	318.30%	190.62%	100.00%	326.04%	192.60%	100.00%	319.27%	190.85%	100.00%

ply voltage level is roughly the same as that of the power consumption, as shown in Table 2.

Tables 4, 5, and 6 present the results of the DVFS scheme in the same way as shown in the previous three tables. With both supply voltage and clock frequency scaled, we can achieve even higher power and temperature reduction. The average power consumed under V_{DDH} and V_{DDs} are about 3.2x and 1.9x of the power consumed under V_{DDL} , as shown in Table 6. Therefore, a raise of supply voltage from V_{DDs} to V_{DDH} increases the power consumption by 68% on the average, while reducing supply voltage from V_{DDs} to V_{DDL} reduces power consumption by 48% on the average. The change in temperature distribution is similar to that of the power distribution.

The above data are graphically presented in Fig. 13 and Fig. 14. The temperatures in the experimented circuits are

given Fig. 13, where Fig. 13 (a) shows the average temperature and Fig. 13 (b) gives the peak temperature. In each case, six different voltage and frequency combinations are experimented: DVS with V_{DDL} , V_{DDs} , V_{DDH} , and DVFS with V_{DDL} , V_{DDs} , V_{DDH} . From this figure, it can be seen that circuits targeted with the highest performance may create excessive heat and very high temperature. This problem can be effectively avoided with the proposed on-chip thermal management mechanism.

Figure 14 illustrates the results of power consumption in the same way as those shown in Fig. 13.

5. Conclusion

This paper presents an on-chip thermal sensing architecture for DVS and DVFS schemes to reduce power and restrict




Fig. 13 Temperatures in the circuits under various supply voltage and frequency combinations: (a) average, (b) peak.

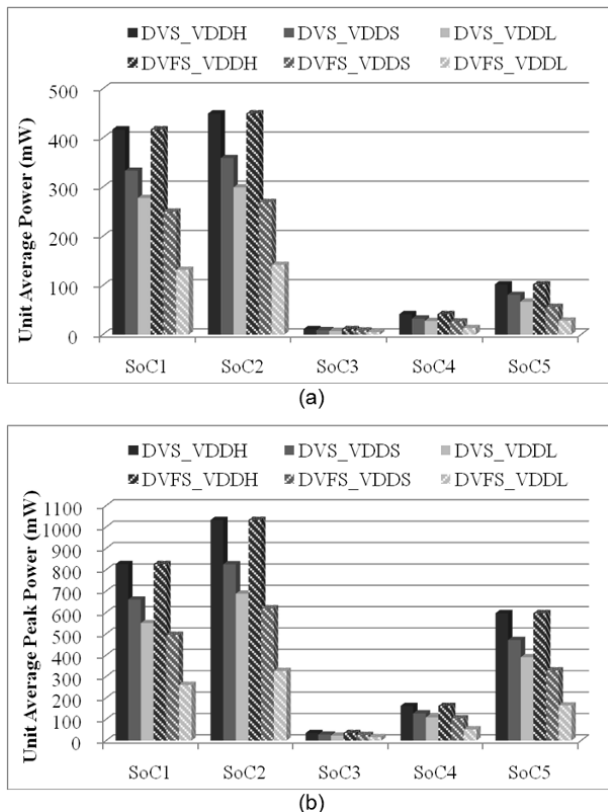


Fig. 14 Temperatures in the circuits under various supply voltage and frequency combinations: (a) average, (b) peak.

temperature in SoC. The main idea is to achieve temperature management through dynamically adjusted supply voltage level, which effectively balances power consumption and performance of a system chip. It is shown that the counter-based thermal sensor is a cheap yet effective way to carry out on-chip real-time temperature measurement. Experimental results confirm that the proposed DVS/DVFS can reduce the power consumption with a lower supply voltage, and the chip temperature will be lower as a result. Since high working temperature is a major source of circuit malfunction and degradation, the proposed method can improve system reliability by restricting the peak temperature allowed in a chip.

Acknowledgments

The authors wish to thank Dr. W.-C. Wu, Industrial Technology Research Institute of Taiwan for his helpful discussion. The authors also appreciate the support from Feeling Technology Corp., Taiwan. This work was partially supported by Ministry of Science and Technology of Taiwan, under Grants NSC-102-2218-E-110-003 and NSC-102-2221-110-86.

References

- [1] "International technology roadmap for semiconductors," 2007, <http://public.itrs.net>
- [2] S.W. Chung and K. Skadron, "A novel software solution for localized thermal problems," *Lect. Notes Comput. Sci.*, vol.4330, pp.63–74, Springer Berlin Heidelberg, 2006.
- [3] Y. Wang, K. Ma, and X. Wang, "Temperature-constrained power control for chip multiprocessors with online model estimation," *Proc. 36th Annual International Symp. Computer Architecture*, pp.314–324, 2009.
- [4] D. Wolpert and P. Ampadu, "Exploiting programmable temperature compensation devices to manage temperature-induced delay uncertainty," *IEEE Trans. Circuits Syst. I: Regular Papers*, vol.59, no.4, pp.735–748, 2012.
- [5] M. Horowitz, T. Indermaur, and R. Gonzalez, "Low-power digital design," *Proc. IEEE Symp. Low-Power Electron.*, pp.8–11, 1994.
- [6] K. Choi, R. Soma, and M. Pedram, "Fine-grained dynamic voltage and frequency scaling for precise energy and performance tradeoff based on the ratio of off-chip access to on-chip computation times," *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.*, vol.24, no.1, pp.18–28, Jan. 2005.
- [7] Intel 80200 Processor Based on Intel XScale Microarchitecture. <http://developer.intel.com/design/iio/manuals/273411.htm>
- [8] Crusoe SE Processor TM5800 Data Book v2.1.
- [9] K. Arabi and B. Kaminska, "Built-in temperature sensors for on-line thermal monitoring of microelectronic structures," *Proc. Int. Conf. Comput.-Design*, pp.462–467, 1997.
- [10] S. Lopez-Buedo, J. Garrido, and E. Boemo, "Dynamically inserting, operating, and eliminating thermal sensors of FPGA-based systems," *IEEE Trans. Compon. Packag. Technol.*, vol.25, no.4, pp.561–566, Dec. 2002.
- [11] S.A. Bota, M. Rosales, J.L. Rosselló, and J. Segura, "Smart temperature sensor for thermal testing of cell-based ICs," *Proc. IEEE/ACM DATE*, pp.464–465, 2005.
- [12] W. Huang, S. Ghosh, S. Velusamy, K. Sankaranarayanan, K. Skadron, and M.R. Stan, "HotSpot: A compact thermal modeling methodology for early-stage VLSI design," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol.14, no.5, pp.501–513, 2006.

- [13] Y.-W. Yang and K.S.-M. Li, "Temperature-aware dynamic frequency and voltage scaling for reliability and yield enhancement," *Proc. IEEE/ACM ASP-DAC*, pp.49–54, 2009.



Katherine Shu-Min Li received the B.S. degree from Rutgers University, New Brunswick, NJ, and the M.S. and Ph.D. degrees from National Chiao Tung University, Hsinchu, Taiwan, in 2001 and 2006, respectively. She is currently an Associate Professor with Department of Computer Science and Engineering, National Sun Yat-Sen University, Kaohsiung, Taiwan. Her research interests include interposer test, 3D IC test & optimization, small delay defect test (SDD), crosstalk effects, smart

grids, fault tolerance, network-on-chips, signal & power integrity, SOC testing, floorplanning and routing for testability and yield enhancement, design for manufacturing, design for yield, transition faults, scan reordering, scan routing, low-power design techniques (DVFS), low-power scan techniques, particularly on oscillation ring test scheme, and interconnect optimization in deep sub-micron and nanotechnology. Dr. Li is a Senior Member of the IEEE Circuits and Systems Society, and a Member of the Association for Computing Machinery (ACM) and the ACM/Special Interest Group on Design Automation.



Yingchieh Ho received the B.S. and M.S. degrees in electronic engineering from National Central University, Chung-Li, Taiwan in 1999 and 2001, respectively. He received the Ph.D. degree at National Chiao-Tung University, Hsin-Chu, Taiwan in 2012. Since Feb. 2013, he has been on the faculty at the Department of Electrical Engineering, National Dong-Hwa University, where he is an Assistant Professor. His research interests are circuit design for low-voltage and biomedical circuit design.



Yu-Wei Yang received the B.S. and M.S. degrees in Computer Science and Engineering from the National Sun Yat-Sen University in 2008 and 2010, respectively. Her research interests are thermal-aware DVFS schemes analysis and design.



Liang-Bi Chen received the B.S. and M.S. degrees in Electronic Engineering from the National Kaohsiung University of Applied Sciences, Kaohsiung, Taiwan, in 2001 and 2003, respectively, and is a Ph.D. Candidate in Department of Computer Science and Engineering at the National Sun Yat-Sen University, Kaohsiung, Taiwan. From August 2008 to September 2008, he had an internship of Department of Computer Science at the National University of Singapore, Singapore. He was also a visiting

researcher in Department of Computer Science at the University of California, Irvine, CA, U.S.A. during September 2008 to August 2009 and in Department of Computer Science and Engineering at the Waseda University, Tokyo, Japan between July 2010 and August 2010. Since 2012, he joined BXB Electronics Co., Ltd, Kaohsiung, Taiwan, as a R&D Engineer. His research interests include power-aware embedded systems design, low-power systems design, system integration, and digital audio processing.