A nanoelectromechanical-switch-based thermal management for 3-D integrated many-core memory-processor system

Huang, Xiwei; Zhang, Chun; Yu, Hao; Zhang, Wei

2012

Huang, X., Zhang, C., Yu, H., & Zhang, W. (2012). A nanoelectromechanical-switch-based thermal management for 3-D integrated many-core memory-processor system. IEEE Transactions on Nanotechnology, 11(3), 588-600.

<https://hdl.handle.net/10356/95634>

<https://doi.org/10.1109/TNANO.2012.2186822>

© 2012 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The published version is available at: [DOI: <http://dx.doi.org/10.1109/TNANO.2012.2186822>].

Downloaded on 26 Apr 2023 20:13:12 SGT

A Nanoelectromechanical-Switch-Based Thermal Management for 3-D Integrated Many-Core Memory-Processor System

Xiwei Huang, Chun Zhang, Hao Yu, *Member, IEEE*, and Wei Zhang, *Member, IEEE*

Abstract—Tera-scale has become the recent interest for high-performance computing system. In order to increase bandwidth yet decrease power, 3-D integrated many-core memory-processor system is one of the most promising solutions. However, the increased power density and longer vertical heat-removal path in 3-D can result in thermal reliability concerns such as thermal runaway and thermal stability, which pose a significant barrier for tera-scale applications. Due to “green-switch” properties such as zero leakage current, infinite subthreshold slope, and temperature resilient behavior, nanoelectromechanical switches (NEMS) are explored in this paper to mitigate the thermal reliability issues for 3-D integrated many-core memory-processor system. The NEMS-based thermal management for 3-D integrated many-core memory-processor system is examined from device, circuit, and system levels, respectively. Moreover, one real-time thermal management is developed for improving system reliability with the use of NEMS-based thermal buffer and power gating. Experimental results show that our proposed approach can effectively prevent the thermal runaway and also maintain high thermal stability for 3-D integrated many-core memory-processor system.

Index Terms—Nanoelectromechanical switches (NEMS), thermal management, 3-D integrated many-core memory-processor system.

I. INTRODUCTION

WITH the rapid increasing demand of high-performance and high-throughput computing, tera-scale system has recently obtained wide interest [1], especially in the domains of server and modern data center. Tera-scale refers to the terabytes of data to be handled by computing system with teraflops of computation performance, which can be applied to real-time 3-D graphics, scientific modeling and simulation, artificial intelligence, etc. [2]. However, such a performance cannot be achieved by the traditional 2-D integrated system due to the lim-

itations of power and bandwidth [3]. Compared to the traditional 2-D integration, by stacking processors and memories vertically and connecting them using through-silicon-vias (TSVs), the 3-D integrated many-core memory-processor system [3], [4] can achieve shorter interconnect length with increased bandwidth and reduced power. As such, 3-D integration has recently become the trend for the integration of thousands of processors and memories within a single chip for tera-scale computing.

However, the primary design challenge for 3-D integrated many-core memory-processor system is the severe thermal concern from twofold. First, the power density is increased due to larger integration density leveraged from the third dimension as more processing components are squeezed into a relative smaller space. Second, the vertical heat-removal path from chip to heatsink has become longer. As such, heat is accumulated inside the chip and hence becomes difficult for dissipation. There are two primary thermal reliability concerns in 3-D integrated many-core memory-processor system, namely *thermal runaway* [8] and *thermal stability* [10]. Thermal runaway is mainly due to a positive feedback loop between leakage current and temperature with exponential dependence. Thermal stability is mainly due to the threshold voltage shifting under increased temperature, which can result in the reduced noise margin for memory cell.

In general, one needs to improve the heat-removal ability to resolve the aforementioned thermal reliability issues in 3-D integrated many-core memory-processor system. Many previous approaches, such as adding thermal TSVs [4], [6] or microfluidic channels [7], [8], have been proposed recently at device and circuit level. However, their primary limitation is the large overhead to signal routing. There also exist a number of thermal management schemes at system level [11], [12], but they are mainly performed in an off-line fashion and hence are unable to consider the real-time power management commonly deployed for modern computing system. In order to improve the thermal reliability of the 3-D integrated many-core memory-processor system, the first design consideration is to reduce the power consumption, i.e., the source of heat generation. Since leakage power contributes significantly to the total power consumption at advanced technology nodes [13], it is the primary power source to be reduced. The second design consideration is to create rapid heat-dissipation path. For example, one can interleave the location of memory and processor such that a large temperature gradient is formed for an efficient heat removal. Finally, the third design consideration is to perform real-time thermal control instead of the static thermal control.

Manuscript received September 29, 2011; accepted January 23, 2012. Date of publication; date of current version. This work was supported by the Singapore Ministry of Education (MOE) Grant TIER-2 ARC5/11 and Grant TIER-1 RG26/10. This paper was presented in part at the 2011 IEEE/ACM International Symposium on Nanoscale Architectures. The review of this paper was arranged by Associate Editor M. De Vittorio.

X. Huang, C. Zhang, and H. Yu are with the School of Electrical and Electronic Engineering, Nanyang Technological University, 639798 Singapore (e-mail: HUAN0234@e.ntu.edu.sg; chunzhang@ntu.edu.sg; haoyu@ntu.edu.sg).

W. Zhang is with the School of Computer Engineering, Nanyang Technological University, 639798 Singapore (e-mail: zhangwei@ntu.edu.sg).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNANO.2012.2186822

The use of “green-switch” by nanoelectromechanical switches (NEMS), which exhibits the infinite subthreshold slope and zero leakage properties, is one promising approach [14]–[16], [18] for leakage reduction recently. In this paper, we explore the design of hybrid NEMS-CMOS for 3-D integrated many-core memory-processor system, such as NEMS-based design for memory and power gating. Based on the unique thermal resilience characteristics of NEMS devices, one real-time thermal management scheme is further developed to improve the thermal reliability. Experimental results by simulations show that compared to the CMOS-based design, the hybrid CMOS-NEMS-based design can reduce the leakage power by 56–70%; and the thermal stability is also improved by 2.3–7 times. Moreover, with the proposed real-time thermal management, the hybrid CMOS-NEMS-based system is able to aggressively throttle the thermal runaway, and the starting-time for thermal reliability can be delayed 4–6.6 times when compared to the CMOS-based system. The real-time thermal management also leads to less power gating overhead and 8 °C lower steady temperature when compared with the static thermal management.

The remaining part of this paper is organized as follows. In Section II, the problem of thermal reliability is addressed and the hybrid CMOS-NEMS-based thermal management is proposed for 3-D integrated many-core memory-processor system. Then, the design considerations of the hybrid CMOS-NEMS-based thermal management are explored from device, circuit, and system levels in Sections III–V respectively. Experimental results are presented in Section VI and the paper is concluded in Section VII.

II. NEMS-BASED THERMAL MANAGEMENT FOR 3-D MANY-CORE MEMORY-PROCESSOR SYSTEM

In this section, the thermal reliability is first defined, and the architecture of NEMS-based thermal management is explained in detail for 3-D many-core memory-processor system.

A. Thermal Reliability in 3-D Many-Core Memory-Processor System

In 3-D many-core system, the power at the source which generates temperature T is defined as the thermal power P_{thermal} . As the temperature varies at ms-scale, P_{thermal} is a thermal-time-averaged value from the electrical-cycle-accurate power P_{cycle} . The thermal power can be expressed as the summation of dynamic power and leakage power

$$P_{\text{thermal}} = \frac{\int_{t_1}^{t_2} P_{\text{cycle}} dt}{(t_2 - t_1)} = P_{\text{dynamic}} + P_{\text{leakage}} \\ = P_{\text{dynamic}} + B \cdot e^{AT} \quad (1)$$

where A and B are the coefficients to describe the exponential dependence between leakage power and temperature from microarchitecture level. The thermal power could induce high-temperature crisis, which will further leads to the thermal reliability problems in 3-D.

In this paper, we focus on thermal reliability problems from two aspects, namely the leakage induced thermal runaway and

Fig. 1. (a) Positive feedback loop of thermal and leakage-power coupling for thermal runaway. (b) Concept of static-noise-margin (SNM) as the metric for thermal stability of SRAM.

thermal stability caused by threshold voltage shifting under increased temperature, as shown in Fig. 1(a). Thermal runaway is caused by the exponential dependence of leakage power and temperature. As such, if there is not enough heat-removal ability to dissipate the generated heat, the chip temperature can keep rising up and finally result in system malfunction [19]. As for the thermal stability, we mainly focus on the memory block due to its high integration density. For illustration purpose, SRAM is employed as the main memory in 3-D many-core memory-processor system. The stability of SRAM is usually evaluated by the static noise margin (SNM) [10], which is the maximum value of dc noise voltage tolerable by the SRAM cell without changing the stored data. SNM can be computed as the side length of the largest square that can be nested between the voltage transfer curves (VTCs) of two inverters in SRAM cell, as shown in Fig. 1(b). When temperature rises, the gain of the inverter in SRAM cell is degraded due to threshold voltage shifting and leads to a lowered SNM. When the external dc noise becomes larger than the SNM, the state of the SRAM cell changes and the stored data are lost.

B. NEMS

The NEMS is an electrostatically actuated mechanical switch whose state of operation is set by the voltage difference between a movable terminal and a fixed terminal [15]. Fig. 2(a) shows the scanning electron microscopy (SEM) picture of a fabricated lateral NEMS device [16]. The basic operation of the device can be illustrated using the three-terminal model shown in Fig. 2(b). The three terminals are *source* (S), *gate* (G) and *drain* (D), similar as CMOS device, with the beam physically connecting to the source. By applying a bias voltage between the gate and the source, i.e., V_{GS} , an electrostatic force will be generated and attracts the beam toward D , while the elastic force in the beam resists the beam from deflecting. The device is *turned ON* when the bias voltage goes beyond a certain threshold, i.e., *pull-in* voltage V_{pi} , such that the beam bends sufficiently to touch the drain and creates a physical conduction path from the source to drain. Similarly, V_{po} is the *pull-out* threshold voltage to disconnect the beam from the drain (i.e., *turn OFF* the device). Note that due to such physical disconnection, NEMS exhibits zero off-state leakage as there is no path for current to flow. Moreover, owing to the surface adhesion force at the switch

Fig. 2. (a) SEM image of a laterally actuated NEMS device [16]. (b) I_{DS} – V_{GS} characteristics to illustrate the hysteresis behavior of the NEMS device: the switch pulls in when $V_{GS} > V_{PI}$, and pulls out when $V_{GS} < V_{PO}$.

TABLE I
COMPARISON OF DIFFERENT TYPES OF NEMS

Number of terminals	2T [19]	3T [18]	4T [15]
Structure	suspended gate	cantilever	cantilever
Control signal	gate-source voltage	gate-source voltage	gate-body voltage
Pull in/out voltage	30/25V	0.8/0.2V	194/NA mV

beam contact region, V_{PO} is usually smaller than V_{PI} , i.e., the I – V characteristics of NEMS exhibit hysteresis. The V_{PI} and V_{PO} can be designed by adjusting the physical dimensions of the NEMS.

The fabrication of NEMS has been demonstrated in [16], [18]–[21] with a range of example circuits such as adder, flip-flop, SRAM, DRAM, DAC, and ADC. Moreover, the NEMS devices remained functional after more than 60 billion cycles. These validate the feasibility of implementing highly integrated hybrid CMOS-NEMS circuits for reduced leakage. Based on the number of terminals, NEMS can be classified in Table I. Comparing the pull in/out voltage, the three-terminal NEMS switch demonstrated in [18] is employed in this paper for thermal management with lowest delay (nanosecond) and area ($300 \sim F^2$).

C. NEMS-Based Thermal Management Solution

Our 3-D hybrid CMOS-NEMS thermal management architecture is shown in Fig. 3 to resolve the thermal reliability concerns. In this 3-D many-core architecture, each layer is composed of *processors* (green blocks) and *memories* (blue blocks). Note that although SRAM is used as the example of memory here, the proposed thermal management can be readily implemented for other types of memories as well. At device level,

Fig. 3. System architecture of 3-D many-core memory-processor with NEMS-based thermal buffer and power gating for both processor and memory at block level.

the NEMS is utilized in the hybrid CMOS-NEMS SRAM cell design to replace the two pull-down transistors, which could reduce the leakage power generation at the source. At circuit level, since the dissipated power from memory is relatively lower compared to the power from processors, memories are interleaved in between processors to act as the *thermal buffers* in both horizontal and vertical directions. Therefore, the relatively high-temperature processors are separated by the low-temperature memories. This leads to a large temperature gradient along the edge which improves the heat transfer and dissipation on chip.

Moreover, at system level, NEMS-based power switches are deployed together with the 3-D power-ground grid to further cut off the leakage current of memories and processors when needed. A real-time thermal reliability-aware controller is proposed to perform an effective power-gating control. Note that the real time here is at the thermal-time scale (i.e., ms) because preventing the thermal-runaway is our objective for power gating. This is different from the previous real-time power gating [22], [23] at the electrical-time scale (i.e., ns), which is not useful for thermal management. Details are discussed in Section V. Note that for the power gating, the retention blocks (black blocks in Fig. 3) are needed to hold the data for memory and processor blocks during power off mode.

III. DEVICE LEVEL—NEMS-BASED MEMORY CELL DESIGN

Due to the ever fining feature size, CMOS-based memory cell suffers from high-leakage induced thermal reliability issues as mentioned previously. Take the SRAM cell as an example, the traditional CMOS SRAM cell could also fail to fulfill the requirements posed by 3-D many-core memory-processor system. Alternatively, the “green switch” NEMS can be used to build the SRAM cell by substituting the two pull-down CMOS transistors [18] as shown in Fig. 3. In this section, we analyze both the advantages and overheads of such hybrid CMOS-NEMS SRAM cell regarding thermal reliability.

Fig. 4. Leakage power dissipations for 45 and 22 nm for CMOS-based SRAM cell versus hybrid CMOS-NEMS-based SRAM cell.

TABLE II
CMOS TRANSISTOR PARAMETERS USED IN LEAKAGE SIMULATION

CMOS Transistor	45nm		22nm	
	W (nm)	L (nm)	W (nm)	L (nm)
Pull-up PMOS	405	45	240	22
Pull-down NMOS	135	45	80	22
Access NMOS	135	45	80	22
V_{dd} (V)	1.2		1	

A. Device-Level Leakage Reduction

Due to its pure mechanical operation mechanism, NEMS behaves drastic leakage current reduction over CMOS transistor. When it is in the off state, there is no path for current to flow from source to drain owing to the physical separation. In other words, the NEMS device exhibits temperature-independent zero off-state leakage current and experiences an abrupt turn-on, i.e., an extremely steep subthreshold slope.

As shown in Fig. 3, in the hybrid CMOS-NEMS SRAM cell where two CMOS pull-down transistors are replaced by NEMS, the physical conduction path from V_{dd} to ground is broken at hold mode, which leads to a dramatic static power dissipation reduction compared to the CMOS SRAM cell. As shown in Fig. 4, 70% and 56% leakage power reduction is achieved for 45- and 22-nm technology, respectively, at room temperature. In our experiment, the Verilog-A behavior model of NEMS comes from [18], and the CMOS transistor model comes from predictive technology model (PTM) [24]. Note that different power supplies are used and proper CMOS transistor sizes are chosen for stable operation as in Table II. The simulation is performed using Cadence Spectre [25].

B. Device-Level Stability Improvement

Compared to the traditional CMOS SRAM cell, the hybrid CMOS-NEMS SRAM cell also greatly improves the thermal stability due to the high-temperature resilience of NEMS. The temperature resilience comes from the mechanical property as the surface adhesion force at the contact region of the switch beam causes difference in V_{pi} and V_{po} and hence the hysteresis behavior. Fig. 5 provides the graphical representation of SNM for both 45- and 22-nm technology during hold and read state at 27 °C. The steep subthreshold slope of NEMS makes the two voltage transition curves form a larger rectangle area, while its hysteresis characteristic of different V_{pi} and V_{po} moves the

Fig. 5. SNM plots for CMOS and hybrid CMOS-NEMS-based SRAM cells in (a) 45-nm hold state, (b) 45-nm read state, (c) 22-nm hold state, and (d) 22-nm read state.

intersection point of the two VTCs to one side, enabling a bigger square to be fitted in. Therefore, NEMS-based SRAM can withstand bigger noise input. As shown in Fig. 5, compared to CMOS SRAM cell, the hybrid CMOS-NEMS SRAM cell improves the stability by 2.3 and 4 times for 45- and 22-nm CMOS in hold state, and also improves the stability by 5 and 7 times in read state.

In addition, the CMOS SRAM cell is more vulnerable at reading state [26] because it must retain the state in presence of the bit-line precharge voltage. When the read operation is performed as in Fig. 5(b) and (d), the access NMOS transistor and the pull-down NMOS transistor form a voltage divider, and the zero-storage node rises. If this read disturb voltage rises larger than the trip voltage of the inverter at the other half of cell, the cell would flip. Therefore, compared to hold state, CMOS SRAM cell has 3 and 2 times poorer SNM in read state for 45- and 22-nm CMOS, and the SNM reduction at read state for CMOS-NEMS design is 17% and 11%, respectively.

C. NEMS Device Overhead and Solution

The usual concern for NEMS-based design is the area and delay overhead due to mechanical beam switching and contact resistances. The area overhead can be alleviated by fabricating NEMS on top of CMOS device through CMOS-compatible fabrication process [18]. For the imbalance between electrical and mechanical delay, it can be minimized by an optimized relay-based design which triggers all mechanical movement simultaneously within each circuit as explained in [20]. As one can fabricate a global switch to turn ON and OFF different NEMS at different SRAM cells simultaneously, the delay overhead actually is acceptable because the primary showstopper such as

Fig. 6. 2-D checkerboard-like architectures with 4-processor and 16-processor per-layer.

the leakage power and thermal runaway can be healed by using NEMS. In addition, owing to the advancing technology, the electrical delay can be further minimized with process scaling. Some recent studies of scaling effect [27], [28] have already reported the obvious mechanical delay reduction. For example, in scaled NEMS circuits, the performance is dominated by the electrical delay (due to high contact resistances) rather than the mechanical delay [28].

IV. CIRCUIT LEVEL—NEMS-BASED THERMAL BUFFER

In circuit level, the interleaved structure of memory and processor is shown to improve the heat dissipation by creating temperature gradient, where the NEMS-based memory is employed as thermal buffer.

A. Thermal Buffer

To reduce the thermal problem in 3-D many-core memory-processor system, besides the power reduction at the source, the creation of heat dissipation path is also necessary. In the many-core system, each processor may form one peak-temperature region or hotspot, and these regions further merge into one huge hotspot across large area. This is not desirable as heat cannot be effectively removed along the lateral layer due to the small temperature gradient. To improve the lateral spreading of heat, one can use the relatively low-power memory, called *thermal buffer*, to separate these processors and therefore reduce the chances for thermal reliability problems. Fig. 6 shows the 2-D checkerboard-like floor-plan suggested in [29]. The processors are interleaved by memories as thermal buffers for easier heat dissipation and can be extended in a 3-D interleaving fashion.

B. Circuit-Level Leakage Reduction

To effectively apply thermal buffer, a relative low power ratio, or thermal gradient, between memory and processor is required. However, with the significant increase of leakage power for technology beyond 45 nm, the leakage power of memory may present the same significance as the dynamic power of processors. Therefore, the feasibility of memory as thermal buffer requires a re-examination.

Without losing generality, we still take the SRAM as an example to study the impact of thermal buffer at circuit level. As illustrated in Fig. 7(a), which presents the power ratio between memory and processor for SPEC-2000 benchmarks under different technologies, the average memory versus processor power ratio for 45-nm CMOS-based memory is 0.24. However, the ratio for 22-nm CMOS-based memory increases to 0.7, 3 times

Fig. 7. (a) Averaged memory/processor power ratio for 45 and 22 nm by CMOS and NEMS technologies. (b) Leakage power of 22-nm memory by CMOS and NEMS technologies at different temperatures.

larger than the 45-nm CMOS technology. Such a high ratio can no longer guarantee the feasibility of thermal buffer. Hence, to enable the continuous usage of memory as thermal buffer for small technology node, a new design of low leakage power is needed.

From circuit level, the leakage reduction effect of NEMS-based memory compared with CMOS-based memory can be observed from Fig. 7(b), which compares the leakage power (Y-axis) of one 4MB SRAM using the 22-nm CMOS and 22-nm hybrid CMOS-NEMS technology, respectively, under different temperatures (X-axis). Note that we also use the term of 22-nm NEMS for the same meaning as 22-nm hybrid CMOS-NEMS. From the figure, we observe that NEMS-based memory design can reduce the leakage power of memory by 60% in average. In addition, Fig. 7(a) shows that at 22-nm technology node, the NEMS-based memory significantly reduces the power ratio and makes the memory versus processor power ratio comparable to that of 45-nm CMOS technology. As such, the reduction in memory leakage power by NEMS enables its utilization as thermal buffer in the 22-nm technology when temperature becomes a major concern for 3-D many-core memory-processor system design.

Now we explore the effectiveness of NEMS-based thermal buffer when integrated into a 16-processor per chip architecture. The checkerboard-like floor-plan as in Fig. 6 is assumed. Different processors are running identical SPEC-2000 application such as *gcc*. Fig. 8(a) shows that although the 45-nm CMOS interleaved memory can still marginally work as thermal buffer to form temperature gradient between processors and memories, there already begins to appear a consolidated region of temperature gradient in the center. On the other hand, it is clear that due to the high memory versus processor power ratio for 22-nm CMOS technology as shown in Fig. 7(a), the memory block can no longer work as a guard-ring to create the temperature gradient from processors and the merged hotspot region occupies almost half of the floor-plan area as shown in Fig. 8(b). As such, the heat spreading ability is poor. The effect of using 22-nm NEMS memory as thermal buffer is shown in Fig. 8(c). Compared to the CMOS-based design in Fig. 8(a) and (b), NEMS-based memory has a much more uniformly distributed temperature profile. Hence, the peak temperature in Fig. 8(c) is much lower. This indicates a good heat-spreading ability, which is very effective to avoid thermal reliability problems.

Fig. 8. Steady-state temperature profiles for 16-processor per-layer: (a) 45-nm CMOS-based memory; (b) 22-nm CMOS-based memory; and (c) 22-nm NEMS-based memory.

Fig. 9. Averaged processor/memory steady-state temperatures for 22-nm CMOS and 22-nm NEMS in (a) 4-processor per-layer and (b) 16-processor per-layer.

Fig. 9 compares the average steady state temperature profiles for processors and memories when running different applications for 4-processor and 16-processor architectures, respectively. We can observe that by using NEMS, the final average temperature difference between memory and processor is 2.6–3.6 times larger than CMOS. The whole average temperature of NEMS-based design is approximately 20 °C and 7 °C lower in 16-processor and 4-processor chip, respectively. Both results verify the effectiveness of NEMS-based thermal buffer.

C. Circuit-Level Stability Improvement

To consider the circuit-level memory stability, the temperature effect on the SRAM cell stability under 45- and 22-nm process is first analyzed. The cell setting is the same as the discussion in previous section. From the results in Fig. 9, the hold and read SNM of hybrid CMOS-NEMS design almost keeps steady at a high value larger than 600 mV. This is because the pull-in voltage and pull-out voltage is only determined by the NEMS beam dimension, contact area, surface properties, etc. [18]. As such, these physical properties are less vulnerable to

Fig. 10. SRAM SNM variations with temperatures in hold and read states.

Fig. 11. SNM distribution profiles for 16-processor per-layer with: (a) 45-nm CMOS SRAM, (b) 22-nm CMOS SRAM, and (c) 22-nm NEMS SRAM.

temperature variations, which is a prominent thermal resilience advantage for NEMS-based memory design. On the contrary, as the CMOS threshold voltage depends heavily on temperature, for CMOS SRAM cell design, the SNM are reducing at a speed of 3.2 mV/10 °C and 4.2 mV/10 °C for 45- and 22-nm process in hold state, and 4.9 mV/10 °C and 3.1 mV/10 °C in read state.

Based on the temperature-aware SNM variation and the thermal distribution profile in Fig. 8, the read SNM distribution profile for CMOS and hybrid CMOS-NEMS SRAM memory are plotted in Fig. 11. The stability of processor is measured as the SNM of its internal SRAM cache memory. Not surprisingly, the hybrid CMOS-NEMS SRAM design has much larger SNM than CMOS design in both 22- and 45-nm technology, which guarantees the stable working region of the 3-D many-core memory-processor system at high temperature.

D. Circuit-Level NEMS Overhead and Solution

The upcoming 3-D integration technology leaves great potential to reduce the overhead of NEMS. For example, although NEMS may consume larger area than CMOS, it does not directly translate into larger memory size under 3-D integration, because NEMS can be fabricated on top of the CMOS device with back-end-of-line compatible technology [18]. In addition, the newly introduced vertical connections with TSVs reduce the delay and increase the communication bandwidth between memories and processors. For example, the processors are able

Fig. 12. (a) Fine-grain header-type power gating. (b) Fine-grain footer-type power gating. (c) Coarse-grain (distributed) power gating.

to concurrently access multiple banks of the memory, which is infeasible in 2-D structures due to the overwhelming signal routing overhead. Combined with the global switch technique described in Section IV-C, the delay overhead of NEMS device can be greatly alleviated at circuit level in 3-D many-core memory-processor system.

V. SYSTEM LEVEL—NEMS-BASED THERMAL MANAGEMENT

Although the use of thermal buffer can distribute the heat uniformly in lateral space, this approach might still not be sufficient to prevent thermal reliability problems for 3-D many-core memory-processor system due to the poor vertical heat-removal path. Thus, it is needed to further look into the solution for leakage power reduction. One common practice for this purpose is to apply power gating [22], [23]. From the 3-D system level, the NEMS-based thermal management utilizes this power gating technique to avoid thermal runaway and keep memories at high stability.

A. NEMS-Based Power Gating

As shown in Fig. 12, power gating, or power switch-OFF technique, usually refers to placing low-leakage PMOS sleep transistors as header switches or NMOS as footer switches on-chip to selectively turning OFF current supply based on the application requirement. Designers are usually employing two types of power gating, namely fine-grain power gating and coarse-grain power gating.

However, the use of traditional CMOS transistor for power gating has certain limitations. First, the CMOS sleep-transistor leaks current by itself even when switched OFF. Moreover, a larger CMOS transistor creates a smaller voltage drop, but also leaks more. In contrast, NEMS sleep transistor has a much higher I_{on}/I_{off} ratio, and when NEMS sleep transistors are sized up, the on resistance can be identical with CMOS but with zero leakage. Hence, NEMS ideally works for the power gating to turn off the current leaking paths. In our architecture, the NEMS-based power-gating switch is designed in a distributed fashion for functional blocks at microarchitecture level. As shown in Fig. 3, NEMS switches are applied for each block of memory and processor at different layers in a vertically distributed manner with data retention blocks.

To effectively control thermal reliability problems, the NEMS switches are controlled by a proposed real-time power-gating controller at thermal-constant scale (ms), thus the delay (ns)

overhead from NEMS is negligible. Next, the criteria to determine thermal runaway and thermal stability are further derived in the following.

B. Thermal Runaway Criteria

A number of previous works [4], [6], [30] have revealed the thermal-runaway phenomenon from temperature-aware microarchitecture level. However, no quantitative criterion is obtained so far that can guide when thermal runaway happens. In the following, a thermal-runaway threshold temperature $T_{threshold}$ is derived for this purpose.

The thermal power $P_{thermal}$ shown in (1) is already exponentially dependent on temperature with parameters A and B . For a given many-core system specification, off-line characterization can be performed to extract the corresponding A and B using the microarchitecture-level leakage power estimation with McPAT [31] and CACTI [32].

Then, at microarchitecture level, we assume that each block has thermal power generation $P_{thermal}$, thermal capacitance C , temperature T , and K heat-removal neighbors. The temperature increase is the result of interaction between the thermal power generation and the heat flux flowing from the neighbors:

$$C \frac{dT}{dt} = P_{thermal} - \sum_{j=1}^K \frac{T - T_j}{R_j}. \quad (2)$$

Note that R_j is the thermal resistance from j th heat-removal neighbor and is obtained from the used thermal model [33].

One observation in the thermal-runaway region is that the transient temperature becomes monotonic (up to second-order derivative) with respect to time [30], i.e.,

$$\frac{dT}{dt} > 0, \quad \frac{d^2T}{dt^2} > 0. \quad (3)$$

Differentiating (2) with respect to time t , one can obtain

$$C \frac{d^2T}{dt^2} = BAe^{AT} \frac{dT}{dt} - \left(\sum_{j=1}^K \frac{1}{R_j} \right) \frac{dT}{dt}. \quad (4)$$

With further consideration of (3), one can have

$$BAe^{AT} > \sum_{j=1}^K \frac{1}{R_j} \quad (5)$$

where BAe^{AT} denotes the leakage-power generation ability, and $\sum_{j=1}^K 1/R_j$ denotes the heat-removal ability. Then, one can define a threshold temperature for the thermal runaway

$$T_{threshold} = T_{threshold}^{runaway} = \frac{1}{A} \ln \left(\frac{1}{AB} \sum_{j=1}^K \frac{1}{R_j} \right). \quad (6)$$

Note that the dynamic power does not depend on T and it is not a regular function of t neither.

Fig. 13. Thermal management with closed feedback control loop.

C. Thermal Stability Criteria

Increased temperature also degrades the thermal stability of CMOS memory. Take SRAM as an example, the SNM of CMOS SRAM decreases due to the threshold voltage shifting at high temperature. Although such degradation will not cause permanent device damage as thermal runaway, it potentially causes system errors and reduces overall system reliability. As such, the suitable temperature management scheme is still needed to maintain the thermal stability for memory.

By calculating the stability-temperature relationship $S(T)$ similar to the SNM curve shown in Fig. 10, we can calculate the temperature threshold $T_{\text{threshold}}^{\text{cell}}$ when one single memory cell enters the unstable region:

$$T_{\text{threshold}}^{\text{cell}} = \text{maximum } T \quad \forall S(T) > S' \quad (7)$$

where S' is the user-defined spec such as minimal acceptable SNM.

However, this threshold temperature $T_{\text{threshold}}^{\text{cell}}$ at device level cannot be directly used in system-level thermal management. Due to the temperature distribution within one memory block, different memory cells vary in stability. As the thermal management is performed at memory block level, it is more reasonable to measure the overall stability of a memory block in a statistical way. As such, here we employ the strategy that evaluates the stability of memory as the total percentage $\varphi(T)$ of unstable cells within it:

$$\varphi(T) = \frac{N'}{N} \quad (8)$$

where N' indicates the number of memory cells whose temperature is larger than $T_{\text{threshold}}^{\text{cell}}$, and N is the total number of memory cells. Then, the temperature threshold to keep one whole memory block stable is defined as

$$T_{\text{threshold}}^{\text{stability}} = \text{maximum } T \quad \forall \varphi(T) < \varphi' \quad (9)$$

where φ' can be decided by the user based on the specific stability requirements for different applications.

D. Real-time Controller for Thermal Management

This paper develops a real-time power gating for thermal management to prevent thermal runaway and thermal stability problems at thermal-time-constant scale (i.e., ms), which is different from the traditional power gating at electrical-time-constant scale (i.e., ns). Given the derived thermal-reliability criteria, potential thermal reliability problems such as thermal runaway or memory instability can be predicted. Based on the prediction, the power gating control loop is depicted in Fig. 13. The input is

Fig. 14. Transient temperature increased with thermal reliability concern, and real-time prediction based on a small moving window.

the thermal-reliability threshold temperature obtained from (6) or (9) off-line. The predicted temperature is compared with the threshold to determine if power gating needs to be performed. Although the predictor-based thermal management has been addressed for dynamic frequency voltage scaling [34], [35], it has not been explored yet for power gating, specifically for reducing leakage-power-induced thermal reliability problems.

As illustrated in Fig. 14, our real-time power gating is performed as follows. Temperature prediction is performed inside one specific time window which is also defined as the control cycle. It consists of a number of sampling intervals, all at ms-scale. Inside one time window at thermal-time-constant scale, auto regressive model $AR(p)$ [36] is employed for the transient temperature prediction. AR model usually works in time series analysis, attempting to predict an output T_t of a system based on the recent historical outputs obtained from temperature sensors: $T_{t-1}, T_{t-2}, \dots, T_{t-p}$, because T_t is linearly regressive against $T_{t-1}, T_{t-2}, \dots, T_{t-p}$:

$$T_t = \sum_{i=1}^p a_i \cdot T_{t-i} + \varepsilon(t). \quad (10)$$

Here, p is the order of the model, usually much less than the length of the series, and the choice of p is a tradeoff between accuracy and computation effort: $a_1, a_2, \dots, a_p$ are the auto regression coefficients; $\varepsilon(t)$ is the uncorrelated error. The parameters can be calculated using least square method and updated in every control cycle. Note that the length of the temperature series is corresponding to the number of sampling intervals in the time window. Since the temperature change is not rapid at thermal-time-constant scale and is partially correlated due to the correlation of workloads, the prediction by windows can still be fast and highly accurate.

In the experiments, to emulate the real-time controller, the temperature change for each control cycle can be estimated using

$$\Delta T_i = \frac{P_i \cdot \Delta t}{c_i} + \frac{T_i \cdot \Delta t}{R_i \cdot C_i} \quad (11)$$

where P_i is the thermal power estimated by power model, Δt is the control cycle, and R_i and C_i are the thermal resistance and capacitance from thermal model.

Fig. 15. Flow-chart of the real-time power-gating control.

Fig. 16. Real-time power gating control with two thresholds.

The real-time power gating control flow chart for each memory and processor block is illustrated in Fig. 15. According to the flow, the procedure of the real-time power gating is summarized as follows.

- 1) The control starts with the initial power input to obtain the transient temperature T .
- 2) At each time window, update the obtained transient temperature from last interval.
- 3) Temperature predictions are made using AR model based on the previously recorded temperature series.
- 4) Check if the block is already power gated or not. If not, the predicted temperature is checked against the thermal-reliability criteria. The block will be power gated if it offends the thermal-reliability threshold.
- 5) If the block is already power gated, we check and wake it up if the predicted temperature goes lower enough to the safe region under T_{safe} .
- 6) Based on thermal management decisions made in steps 4 and 5, the system run to the next time window and all block temperatures will be updated and recorded.

Note that the two temperatures $T_{threshold}$ and T_{safe} turn the real-time thermal control into a hysteresis control, as shown in Fig. 16. As the temperature changes with the generated power at ms-scale, the effect of temperature variation could be later than the power variation at electrical-cycle level. Considering this and that the power of system workload usually changes stochastically, the margin between $T_{threshold}$ and T_{safe} can prevent the

Fig. 17. (a) Thermal-runaway threshold temperature prediction based on the derived criteria from training data. (b) Thermal-runaway threshold temperature simulation result extracted from Hotspot.

overhead induced by too frequent sleep/wake-up transition. The result is verified in the experiment section.

VI. EXPERIMENTS BY SIMULATION

To evaluate the effectiveness of NEMS-based thermal buffer and power gating in keeping thermal reliability, the 3-D many-core memory-processor system with CMOS-NEMS-based thermal management circuit is built as in Figs. 3 and 6. All NEMS devices are based on the model in [18], and all CMOS devices (45 and 22 nm) are based on the PTM [24]. In our numerical experiments, SRAM is used as an example of memory. Modified CACTI [32] and McPAT [31] are used to model both the CMOS- and NEMS-based memories and processors with consideration of leakage power. Each SRAM block is assumed to be 4 MB, the typical design specification for current many-core memory-processor architectures. The Wattch [37] simulator is also modified to consider the advanced CMOS technology and generate dynamic power trace. The resulting platform is further integrated together with HotSpot [33] for electrical-thermal coupling simulation and thermal resistance and capacitance calculation. The real-time temperature prediction is implemented with AR model by assuming the given measured power traces. As such, one can evaluate the proposed hybrid CMOS-NEMS platform for 3-D many-core memory-processor system with thermal management.

A. Real-time Thermal Management for Thermal Runaway

We demonstrate the effectiveness of the real-time power gating in preventing the thermal runaway reliability problem for 3-D many-core memory-processor system in this part. As one illustration example, we set the heat-spreader size to be 3 cm × 3 cm, and the heat-sink size to be 6 cm × 6 cm. Based on the thermal model used in [33], the equivalent heat-removal resistance R is calculated as 0.365 K/W, i.e., the heat-removal ability is 2.74 W/K. In addition, from off-line leakage power characterization, $A = 0.0666$, $B = 0.0024$ for 45-nm CMOS with 16 processors. As such, the system has large power dissipation but poor thermal removal ability. As shown in Fig. 17(a), the thermal-runaway threshold temperature $T_{threshold}$ is

Fig. 18. Transient temperatures of thermal runaway under the dynamic thermal management for 22-nm CMOS- and 22-nm NEMS-based designs.

Fig. 19. Thermal runaway starting time with respect to power traces for CMOS- and NEMS-based designs.

Fig. 20. Transient temperature comparison for real-time power-gating control with SNM-constraint for CMOS-based SRAM, and thermal-runaway-constraint for NEMS-based SRAM.

approximately 418 K (145 °C), which agrees with the detailed thermal transient simulation by HotSpot [33] as in Fig. 17(b).

With $T_{\text{threshold}} = T_{\text{threshold}}^{\text{runaway}} = 145^\circ\text{C}$, we go on to do the real-time thermal management for thermal runaway. For a fast simulation, the initial temperatures for all the blocks are set to be 100°C . The real-time power gating is applied at thermal-time constant level with 0.06-ms sampling period in a time window of 1.2 ms for the prediction. T_{safe} is empirically set as 90% of the $T_{\text{threshold}}$ to leave a control margin, i.e., 130.5°C . In practice, this percentage could be specified by the user.

The targeted structure is a four-layer with 16-processor per-layer 3-D many-core system running SPEC2000 applications *gcc*. We investigate the real-time thermal-runaway control of one block which is going to the status of thermal runaway. For 22-nm CMOS and 22-nm NEMS, the results in Fig. 18 are obtained. It shows that the thermal runaway happens at 145°C when without thermal control. The transient temperature originally will get into the runaway region for both CMOS and NEMS when there is no thermal-runaway control. However, with the real-time power gating, the temperature can be successfully throttled back. As we have employed the hysteresis control, the power gating will be stopped unless T_{safe} is arrived. Thus, the transient temperature increases again, yet always below the thermal runaway region.

In addition, considering the starting time of thermal runaway, 22-nm NEMS is later than 22-nm CMOS as seen in Fig. 18. Based on the results of more benchmarks *gzip*, *vortex*, *crafty*, *mesa*, and *parser*, the average NEMS-based thermal runaway starting time is approximately 4 times later than that of CMOS as shown in Fig. 19. This longer reliable running time also testifies the advantage of NEMS for leakage and thermal-runaway reduction in 3-D many-core memory-processor systems.

B. Real-time Thermal Management for Thermal Stability

In this part, we consider the real-time power gating for keeping the thermal stability of memories. The 22-nm SRAM is chosen as an example, since it is most vulnerable to noise interference in read state. According to the SNM plot in 22-nm CMOS read state, the curve fitting is done to get the function between

SNM (V) and temperature ($^\circ\text{C}$), i.e.,

$$S(T) = (-0.0003 \cdot T + 0.0982)V. \quad (12)$$

From the analysis in Section IV-C, we assume the minimum stable SNM threshold to be 60 mV. Then, the threshold temperature $T_{\text{threshold}}^{\text{cell}} = 127^\circ\text{C}$ can be obtained for the 22-nm CMOS SRAM cell. We further define the thermal stability threshold $T_{\text{threshold}}^{\text{stability}}$ for one memory block as the temperature causing more than 10% memory cells to enter unstable region.

Note that as thermal runaway tends to physically damage the device permanently, its criteria from (6) must still be satisfied. Thus, the overall thermal management threshold temperature becomes

$$T_{\text{threshold}} = \min\{T_{\text{threshold}}^{\text{stability}}, T_{\text{threshold}}^{\text{runaway}}\}. \quad (13)$$

This is important for NEMS-based memory, since it has high enough SNM, which implies that the $T_{\text{threshold}}^{\text{stability}} \gg T_{\text{threshold}}^{\text{runaway}}$. Therefore, its thermal control threshold temperature is still dominated by $T_{\text{threshold}}^{\text{runaway}}$.

Using the same benchmark and system configuration as in the previous section, and also set T_{safe} here as 90% of the $T_{\text{threshold}}$, the results for one block are shown in Fig. 20. We can see the 22-nm CMOS-based memory has temperature successfully kept below 127°C , which is the stable working region. On the contrary, NEM-based memory is able to work at higher temperature

Fig. 21. Thermal stability starting time with respect to power traces for CMOS- and NEMS-based designs.

Fig. 22. Transient temperature comparison for real-time power-gating control and static power-gating control for 22-nm NEMS-based designs.

region due to its much higher stability (i.e., SNM), and is only limited by thermal runaway problem.

Similarly, the starting time for thermal stability problem is shown in Fig. 21, with the average NEMS-based thermal reliability problem starting time 6.6 times later than that of CMOS. The even larger ratio compared to that in Section VI-A comes from the fact that NEMS-based design is much less vulnerable to thermal stability problems defined in Section V-C. Thus, the use of NEMS in memory design is verified to have better performance in thermal stability improvement.

C. Comparison of Real-time and Static Control

Moreover, to testify the effectiveness of the real-time power gating, we compare the transient temperature profiles between real-time control and static control. The static controller works by simply turning OFF the processor and memory block whenever its temperature is above $T_{\text{threshold}}$, and turning it on once temperature goes below $T_{\text{threshold}}$. The result for 22-nm NEMS-based design is shown in Fig. 22. We can see the transient temperature for real-time control varies between the two previously defined $T_{\text{threshold}}$ and T_{safe} , while the temperature for static control keeps fluctuating around the 145°C threshold. And the average temperature for the real-time control is 8°C smaller than that of static control.

Fig. 23. Power-on block ratio comparison for real-time power-gating control and static power-gating control.

Fig. 23 further illustrates the ratio of power-on blocks in both control schemes. Obviously, the frequent fluctuations for static control stand for the high frequency of block ON/OFF switching, which results in large power gating and data retention overhead. On the contrary, due to the hysteresis control scheme, the real-time control greatly reduces the power gating frequency, which will further translated into better system performance.

VII. CONCLUSION

The increased leakage power by technology scaling and longer heat removal path by vertical layer stacking can introduce significant thermal reliability problems in 3-D integration, namely thermal runaway and thermal stability. In this paper, we have explored the solution by the NEMS-based thermal management with the proven results by extensive experimental validations from device, circuit, and system levels, respectively. Experimental results show that compared to the CMOS-based design: at device level, the hybrid CMOS-NEMS-based design can reduce the leakage power by 56–70% and can maintain 2.3–7 times larger thermal stability; at circuit level, the hybrid CMOS-NEMS-based design can maintain the effectiveness of memory as thermal buffer for better heat dissipation; and at system level, the hybrid CMOS-NEMS-based real-time thermal management is able to aggressively throttle the thermal runaway and maintain the thermal reliable running time for 4–6.6 times longer. The real-time thermal management also leads to less power gating overhead and 8°C lower steady temperature when compared to the static thermal management. All these perspectives have demonstrated the advantage when applying NEMS for thermal management in the future 3-D many-core memory-processor system at tera-scale.

ACKNOWLEDGMENT

The authors would like to express their greatest gratitude to the reviewers, who spent effort in reviewing and made great suggestions for the improvement of this work.

REFERENCES

- [1] Intel Research: Tera-Scale. [Online]. Available: <http://techresearch.intel.com/ResearchAreaDetails.aspx?Id=27>
- [2] J. Bautista, "Tera-scale computing and interconnect challenges," in *Proc. Design Autom. Conf.*, 2008, pp. 665–667.
- [3] F. Li, C. Nicopoulos, T. Richardson, Y. Xie, N. Vijaykrishnan, and M. Kandemir, "Design and management of 3D chip multiprocessors using network-in-memory," in *Proc. Int. Symp. Comput. Archit.*, 2006, pp. 130–141.
- [4] K. Banerjee, S. J. Souri, P. Kapur, and K. C. Saraswat, "3D ICs: A novel chip design for improving deep submicron interconnect performance and systems-on-chip integration," *Proc. IEEE*, vol. 89, no. 5, pp. 602–633, May 2001.
- [5] H. Yu, J. Ho, and L. He, "Allocating power ground vias in 3D ICs for simultaneous power and thermal integrity," *ACM Trans. Design Autom. Electron. Syst.*, vol. 14, no. 3, pp. 1–31, 2009.
- [6] H. Yu, Y. Y. Shi, L. He, and T. Karnik, "Thermal via allocation for 3D ICs considering temporally and spatially variant thermal power," *IEEE Trans. Very Large Scale Integr. Syst.*, vol. 16, no. 12, pp. 1609–1619, Dec. 2008.
- [7] J. M. Koo, S. J. Im, L. Jiang, and K. E. Goodson, "Integrated microchannel cooling for three-dimensional electronic circuit architectures," *J. Heat Transfer*, vol. 127, no. 1, pp. 49–58, 2005.
- [8] H. Qian, X. Huang, H. Yu, and C. H. Chang, "Cyber-physical thermal management of 3D multi-core cache-processor system with microfluidic cooling," *J. Low Power Electron.*, vol. 7, no. 1, pp. 110–121, 2011.
- [9] A. Vassighi and M. Sachdev, "Thermal runaway in integrated circuits," *IEEE Trans. Device Mater. Rel.*, vol. 6, no. 2, pp. 300–305, Jun. 2006.
- [10] E. Seevinck, F. J. List, and J. Lohstroh, "Static-noise margin analysis of MOS SRAM cells," *IEEE J. Solid-State Circuits*, vol. SSC-22, no. 5, pp. 748–754, Oct. 1987.
- [11] A. Kumar, L. Shang, L.-S. Peh, and N. K. Jha, "System-level dynamic thermal management for high-performance microprocessors," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 27, no. 1, pp. 96–108, Jan. 2008.
- [12] J. Donald and M. Martonosi, "Techniques for multi-core thermal management: Classification and new exploration," in *Proc. Int. Symp. Comput. Archit.*, 2006, pp. 78–88.
- [13] H. F. Dadgour and K. Banerjee, "Hybrid NEMS-CMOS integrated circuits: A novel strategy for energy-efficient designs," *IET Comput. Digital Tech.*, vol. 3, pp. 593–608, 2009.
- [14] Y. Zhou, S. Thekkel, and S. Bhunia, "Low power FPGA design using hybrid CMOS-NEMS approach," in *Proc. Int. Symp. Low Power Electron. Design*, 2007, pp. 14–19.
- [15] F. Chen, H. Kam, D. Markovic, T. K. Liu, V. Stojanovic, and E. Alon, "Integrated circuit design with NEM relays," in *Proc. Int. Conf. Comput. Aided Design*, 2008, pp. 750–757.
- [16] H. F. Dadgour, M. M. Hussain, C. Smith, and K. Banerjee, "Design and analysis of compact ultra energy-efficient logic gates using laterally-actuated double-electrode NEMS," in *Proc. Design Autom. Conf.*, 2010, pp. 893–896.
- [17] K. Akarvardar, C. Eggimann, D. Tsamados, Y. S. Chauhan, G. C. Wan, A. M. Ionescu, R. T. Howe, and H.-S. P. Wong, "Analytical modeling of the suspended-gate FET and design insights for low-power logic," *IEEE Trans. Electron Devices*, vol. 55, no. 1, pp. 48–59, Jan. 2008.
- [18] S. Chong, K. Akarvardar, R. Parsa, J.-B. Yoon, R. T. Howe, S. Mitra, and H.-S. P. Wong, "Nanoelectromechanical (NEM) relays integrated with CMOS SRAM for improved stability and low leakage," in *Proc. Int. Conf. Comput. Aided Design*, 2009, pp. 478–484.
- [19] A. M. Ionescu, V. Pott, R. Fritsch, K. Banerjee, M. J. Declercq, P. Renaud, C. Hibert, P. Fluckiger, and G. A. Racine, "Modeling and design of a low-voltage SOI suspended-gate MOSFET (SG-MOSFET) with a metal-over-gate architecture," in *Proc. Int. Symp. Quality Electron. Design*, 2002, pp. 496–501.
- [20] M. Spencer, F. Chen, C. Wang, R. Nathanael, H. Fariborzi, A. Gupta, H. Kam, V. Pott, J. Jeon, T.-J. King Liu, D. Markovic, E. Alon, and V. Stojanovic, "Demonstration of integrated micro-electro-mechanical relay circuits for VLSI applications," *IEEE J. Solid-State Circuits*, vol. 46, no. 1, pp. 308–320, Jan. 2011.
- [21] T.-H. Lee, S. Bhunia, and M. Mehregany, "Electromechanical computing at 500 °C with silicon carbide," *Science*, vol. 329, no. 5997, pp. 1316–1318, Sep. 2010.
- [22] H. Xu, W.-B. Jone, and R. Vemuri, "Accurate equivalent energy breakeven time estimation for power gating," in *Proc. Int. Conf. Comput. Aided Design*, 2008, pp. 161–168.
- [23] A. Agarwal, H. Li, and K. Roy, "DRG-cache: a data retention gated-ground cache for low power," in *Proc. Design Automation Conf.*, 2002, pp. 473–478.
- [24] PTM [Online]. Available: <http://ptm.asu.edu/>
- [25] *Virtuoso Spectre Circuit Simulator User Guide*, ver. 5.1.41, Cadence Design System Inc., San Jose, CA.
- [26] A. J. Bhavnagarwala, X. Tang, and J. D. Meindl, "The impact of intrinsic device fluctuations on CMOS SRAM cell stability," *IEEE J. Solid-State Circuits*, vol. 36, no. 4, pp. 658–665, Apr. 2001.
- [27] H. F. Dadgour, M. M. Hussain, A. Cassell, N. Singh, and K. Banerjee, "Impact of scaling on the performance and reliability degradation of metal-contacts in NEMS devices," in *Proc. IEEE Int. Rel. Phys. Symp.*, 2011, pp. 280–289.
- [28] H. Kam, T.-J. King Liu, V. Stojanovic, D. Markovic, and E. Alon, "Design, optimization, and scaling of MEM relays for ultra-low-power digital logic," *IEEE Trans. Electron Devices*, vol. 58, no. 1, pp. 236–250, Jan. 2011.
- [29] K. Sankaranarayanan, S. Velusamy, and K. Skadron, "Microarchitectural floorplanning for thermal management: A technical report," Dept. Comput. Sci., Univ. of Virginia, Charlottesville, Tech. Rep. CS-2005-08, May 2005.
- [30] W. Liao, L. He, and K. M. Lepak, "Temperature and supply voltage aware performance and power modeling at microarchitecture level," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 24, no. 7, pp. 1042–1053, Jul. 2005.
- [31] McPAT [Online]. Available: <http://www.hpl.hp.com/research/mcpat/>
- [32] CACTI [Online]. Available: <http://www.hpl.hp.com/research/cacti/>
- [33] HotSpot [Online]. Available: <http://lava.cs.virginia.edu/HotSpot/>
- [34] A. Coskun, T. S. Rosing, and K. Gross, "Proactive temperature management in MPSoCs," in *Proc. Int. Symp. Low Power Electron. Design*, 2008, pp. 165–170.
- [35] Y. Wang, K. Ma, and X. Wang, "Temperature-constrained power control for chip multiprocessors with online model estimation," in *Proc. Int. Symp. Comput. Archit.*, 2009, pp. 1681–1696.
- [36] J. D. Hamilton, *Time Series Analysis*. Princeton, NJ: Princeton Univ. Press, 1994.
- [37] Wattch [Online]. Available: <http://www.eecs.harvard.edu/~dbrooks/wattch-form.html>

Xiwei Huang received the B.Eng. degree from Beijing Institute of Technology, Beijing, China, in 2009. He is currently working toward the Ph.D. degree in the School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore.

His research interests include nanoelectromechanical systems-based and microfluidics-based thermal management in 3-D many-core on-chip system.

Chun Zhang received the B.S. and Ph.D. degrees in microelectronics from Fudan University, Shanghai, China, in 2005 and 2011, respectively.

During 2010 and 2011, he was a Visiting Student in the Department of Electrical Engineering, University of California, Los Angeles and the Department of Electrical and Computer Engineering, University of Alberta. Since July 2011, he has been a Research Fellow at the School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore. His primary research interests are 3-D

computing systems and designs at nanotera scale, cyber-physical energy management systems, design automation and novel architectures of field-programmable gate array.

Hao Yu (S'02–M'06) received the B.S. degree from Fudan University, Shanghai, China, in 1999, and the M.S. and Ph.D. degrees in electrical engineering from University of California, Los Angeles, in 2007, with major of integrated circuit and embedded computing.

He was a Senior Research Staff at Berkeley Design Automation. Since October 2009, he has been an Assistant Professor at the School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore. He has 58 peer-reviewed publications. He is the Associate Editor and Technical Program Committee Member of several journals and conferences. His primary research interests include 3-D computing system and designs at nanotera scale.

Dr. Yu received the Best Paper Award in *ACM Transactions on Design Automation of Electronic Systems*, Best Paper Award nominations in the Design Automation Conference, the International Conference on Computer Aided Design, and the Asia and South Pacific Design Automation Conference, and Inventor Award from Semiconductor Research Cooperation.

Wei Zhang (M'xx) received the Bachelor's and Master's degrees in electrical engineering from the Harbin Institute of Technology, Harbin, China, in 1999 and 2001, respectively, and the Ph.D. degree in computer engineering from Princeton University, Princeton, NJ, in 2009.

She is currently an Assistant Professor at the School of Computer Engineering, Nanyang Technological University, Singapore. Her research interests include embedded systems, reconfigurable computing, nanotechnology, electronic design automation, and network-on-chip.

Dr. Zhang received a Princeton Research Prize and a Best Paper Award in 2007 and 2009, respectively. She is a member of the Association for Computing Machinery.

Q3

IEEE
Proof