Technical Report GIT-CERCS

Thermal Field Management for Many-core Processors

Minki Cho, Nikhil Sathe, Sudhakar Yalamanchili and Saibal Mukhopadhyay
School of Electrical and Computer Engineering
Georgia Institute of Technology, Atlanta, Georgia, U.S.A
10 April 2009

Table of Contents

Abstract

This paper first presents an analysis of the global thermal field in many core processors in deep nanometer (to 16nm) nodes under power and thermal budget. We show that the thermal field can have significant spatiotemporal non-uniformity along with high maximum temperature. We propose spatiotemporal power multiplexing as a proactive method to reduce spatial and temporal temperature gradients. Several power multiplexing policies are evaluated for a 256 core many-core processor in 16nm nodes which demonstrate that the simple cyclic core-activation can achieve highly uniform thermal field with low maximum temperature.

1. Introduction

The semiconductor and computing industry is progressing towards chips with hundreds of on-chip cores. Moore's Law will be sustained by growth in the number of cores rather than frequency scaling. [1-2]. However, the increase in number of cores can far outweigh the reduction in power-per-core in successive technology generations, resulting in unsustainable growth in chip-power and die-temperature. Thus we find that the power and thermal budget will limit the number of cores that can be simultaneously active. To overcome these limits requires a comprehensive look at the process of power consumption, heat generation, and management of thermals across the die – thermal field management.

The management of maximum on-chip temperature is an integral part of integrated circuits (IC) in sub-100nm nodes [3-7]. A higher on-chip temperature can cause functional failures and/or degrade circuit reliability [3, 8]. ICs in nanometer nodes experiences significant dynamic variation in the power dissipation due to workload variations or run-time power control methods. This results in temporal variation in temperature or thermal cycles. Large and fast temporal variation in temperature (i.e. faster thermal cycles) is also detrimental to circuit reliability [3]. Finally, on-chip spatial variation in power results in spatial gradients in on-chip temperature. A higher spatial gradient results in different operating parameters (e.g. different delay) for circuit blocks within a chip causing functional failures. Further, the higher maximum temperature and spatiotemporally non-uniform temperature distribution also increase the cooling energy [9]. Hence, reducing the maximum temperature and maintaining the spatiotemporal uniformity of on-chip temperature is very important for energy-efficient and reliable computing.

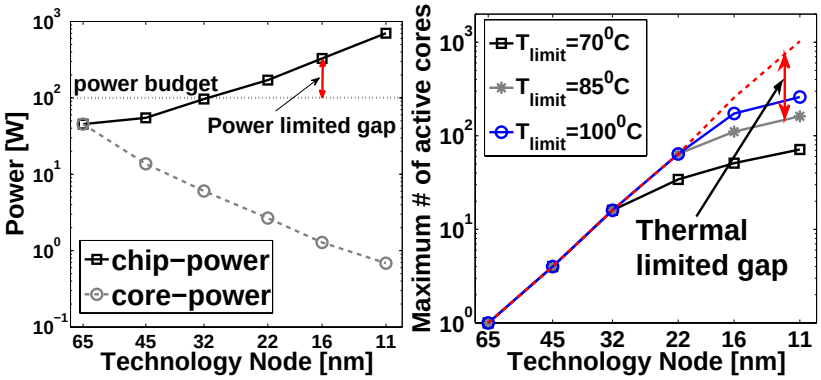

(a) (b)

Fig. 1: Limits to active cores (a) power, and (b) temperature.

The management of the maximum temperature in single or multi-core processors has received significant attention over the years. The basic principle that has been explored is to distribute the heat dissipation in space or time. For single or few-core systems with pre-defined locations of hotspots (i.e. regions with higher temperature) thermal-aware floorplanning methods have been developed [10]. On the other hand, dynamic thermal management (DTM) and dynamic voltage frequency scaling approaches (DVFS) have been explored to distributed heat generation in time to reduce the *maximum temperature* [3-7]. For DTM, power-aware scheduling techniques have been considered for single processor systems [4]. Chaparro et. al. considered thread migration (TM or core hopping) and DVFS techniques to reduce maximum temperature and improve performance in multicore architecture with 16 cores [5]. Coskun et. al. discussed spatial gradient and thermal cycles as important factors along with maximum temperature for multiprocessing SoCs [3]. The reactive thermal management strategies, such as thread migration, have been combined with scheduling policies to reduce the hotspot and temperature gradient in space and time [3]. However, there has been minimal effort to recognize and manage the global distribution of heat/temperature in many-core processors.

This paper presents insights on the thermal fields for many core architectures in deep nanometer nodes (to 16nm), proposes proactive policies to simultaneously control the maximum temperature and spatiotemporal uniformity of temperature, and assesses the performance of these policies. Since we aim to analyze and manage the thermal field instead of only maximum temperature, the proposed approach is referred to as **'*thermal field management*'**. This paper makes following contributions:

- We analyze the thermal field in a many-core processor in extremely scaled technology nodes and illustrate the challenges in thermal field management.
- We propose proactive *spatiotemporal power multiplexing* as a principle for thermal field management to redistribute

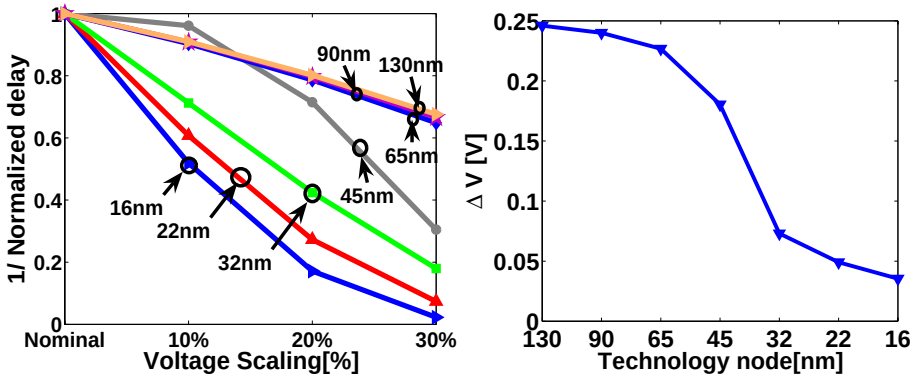

Fig. 2: Limitation of fine control for DVS (a) voltage vs frequency, and (b) voltage change for 20% frequency reduction

heat generation in both space and time.

- We present different policies for the spatiotemporal power multiplexing and analyze its impact on spatiotemporal uniformity and maximum temperature.
- We analyze the energy-efficiency of the proposed approach considering on-chip delay and leakage variations.

The effectiveness of the proposed approaches is evaluated considering a predictive 256 core system at International Technology Roadmap of Semiconductor (ITRS) predicted 16nm node [1]. Our approach can achieve uniform spatiotemporal thermal field and reduce the maximum temperature across the chip.

2. Analysis of the Thermal Field

We first analyze the power/thermal behavior of many core processors at deep nanometer technologies using the methods described in Section 5. We consider a tile-type architecture where the set of tiles (cores) are arranged in two dimensions [11, 12]. The cores are connected by a mesh network. The core complexity (simple cores with 16 million gates) and frequency (3GHz) is assumed to constant over technology generations. The technology driven core power estimation tool proposed by Sekar et. al. [13, 14] and the ITRS predicted technology parameters [1] were considered in the analysis. We arrive at the following important observations.

1. ***Thermal/power limited capacity of many-core processor:*** We find that the power and thermal budget will limit the number of cores that can be simultaneously active in a many-core processor. This is illustrated in Fig. 1 considering a 4X increase in the number of cores per generation (complexity and frequency was kept constant). Note that the reduction in the total power (active and leakage) per core is more than overshadowed by the increased number of cores. While the exact values of the power consumed and peak temperature will depend on the specifics of the device and processor architecture, we find the presence of the power/thermal limited capacity to be a fundamental architectural limitation to be explicitly

managed and overcome.

2. ***Activate-deactivate core control for power management:*** We observe that the effectiveness of DVFS will reduce with scaling. This is illustrated in Fig. 2 considering circuit simulations of delay-voltage characteristics of an 8-stage FO4 ring-oscillator at different predictive technology nodes [15]. For a target frequency reduction the opportunity for voltage reduction reduces at scaled nodes. For example, the allowable voltage reduction for a 20% frequency reduction target at 16nm node is only 5% or ~40mV (considering ITRS predictions of ~0.7V supply voltage). This suggests the utility of fine-grain DVFS will be limited at best. We predict that core level power management is better performed using activate-deactivate control. The deactivated cores employ clock and supply gating and dissipate negligible power. The recent developments of many-core system supports this need for fine-grain activate-deactivate control of cores [12]. The increased process variation with scaling will also limit the opportunity for fine-grain DVFS.

2.1. Structure of the Thermal field in Many-Core Processors

Based on the above mentioned observations, we predict that at deep nanometer nodes effective power management approach is to operate each core at nominal power and control the number of active cores. If cores are activated on a schedule such that the number of active cores is fixed and less than the total number of cores, the workload and locations of the active cores vary over time. The interaction of this execution principle with physical nature of heat flow creates a *spatiotemporally non-uniform global thermal field.* The physical properties of heat flow, such as spatial and temporal variations, do not scale with area of the cores. Thus with the reduction of the physical area of a core across technology generations the temperature variation across the physical area of a scaled core continue to diminish until from a thermal perspective, each active core behaves as a thermal hotspot creating 'distributed hotspots' across the on-chip thermal variations. This increases the on-chip spatial gradients in thermal field (Fig. 3a). Both the location and temperature of the hotspots can change over time as the location of the active cores change. This creates a spatiotemporally non-uniform and non-deterministic global thermal field. In contrast, in modern low core count architectures hotspot locations are known, e.g., the register files. On the other hand, as each core experiences sequences of active-inactive states, there will be temporal variations in the core temperature (thermal cycles). The activation-deactivation sequence as well as temporal variation in workload of the core will have a strong impact on this temporal thermal gradient (Fig. 3b). Unlike the case of single or few core systems where maximum temperature mainly depends on the total power dissipation, in many-core processors the maximum on-chip temperature will depend on both the number (i.e. total chip-power) and location of the active cores. For

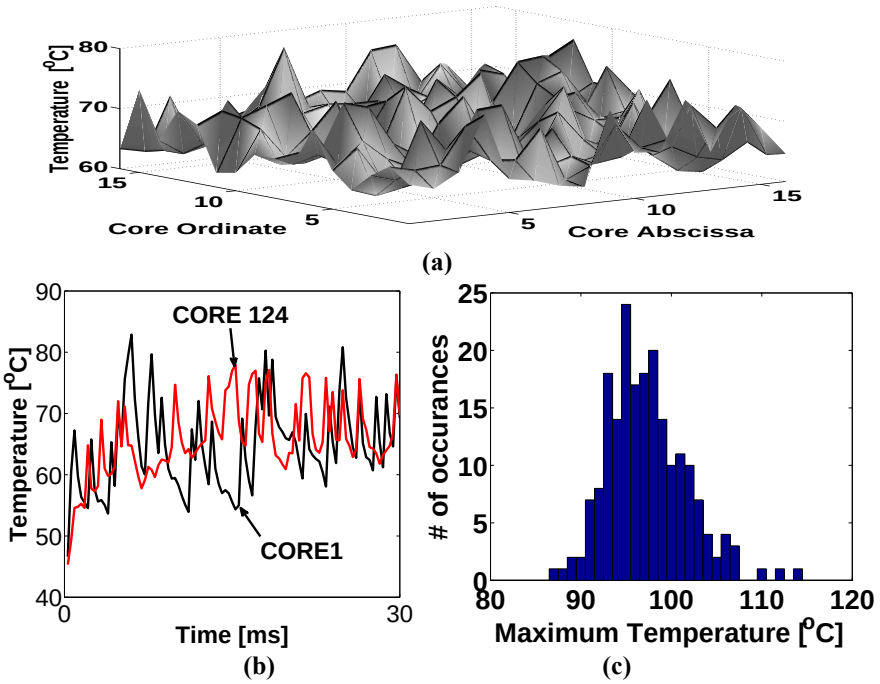

Fig. 3: Thermal field in many-core: (a) spatial, (b) temporal, (c) maximum temperature

example, same number of active cores but their different locations can result in 20°C change in temperature (Fig. 3c). If active cores are concentrated in a physical region of the chip (a physically compact set), the maximum on-chip temperature can rise faster and rise to a higher value. Hence, the power dissipation and thermal field in many-core processor has a complex time-varying relationship. Note if we assume all cores are always active we approach a special case of the generic thermal field discussed above that has spatial uniformity [16]. However, such an analysis neglects the technological trends in power-per-core and does not consider the challenges to arbitrary reduction of computation power under process variation while maintaining accuracy of the computation.

3. Basic Principles for Thermal Field Management

The primary goal of thermal field management is to reduce the spatial variation in temperature across cores and temporal variation at a core, i.e., we seek to maintain a uniform spatiotemporal field while minimizing the maximum temperature (Fig. 4). The thermal field is characterized using the following metrics (Fig. 4).

- ***Spatial gradient***: The temperature difference between two adjacent cores at a given time.
- ***Spatial difference***: The difference between maximum and minimum on-chip temperature at a given time.
- ***Temporal gradient***: The temporal temperature variation a core experiences within a short time interval (~10-100μs).

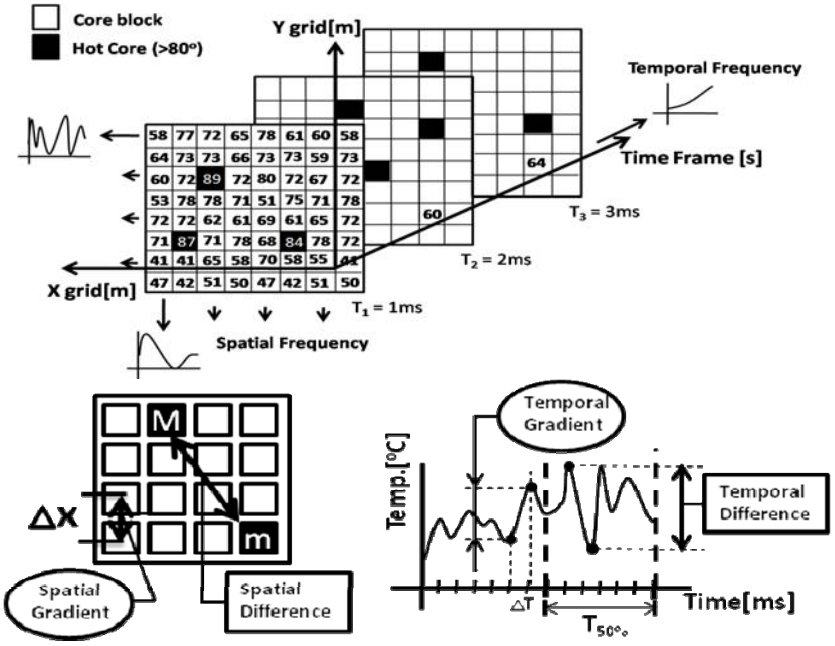

Fig. 4: Spatial and temporal frequency concept diagram

- ***Temporal difference***: The difference between the maximum and minimum temperature for a core over a long time period (~10ms).

3.1 Spatiotemporal Power Multiplexing

Consider a 256 core processor at the 16nm technology node with a target throughout equivalent to 64 active cores (each active core dissipates ~1.5W as per our model from Section 5). We now consider the methods for distributing the heat generation in space and time. First, we consider a spatial approach with 64 active cores distributed spatially in the entire die (i.e. die-area for a 256 core processor, low power cache memory between neighboring cores). We call this approach "*Selected Core On"* (Fig. 5). Next we consider distribution of heat over time by activating all cores and turning 'all' of them periodically (Fig. 5). This condition is called "*Full core on and off*" since all cores are active during same time interval and off periodically. Our approach is to distribute the heat in both space and time –i.e. we select a set of 64 cores and keep them on for a given period of time (time-slice). The power is migrated to a different set of 64 cores in the next time-slice. As over time the total power is multiplexed between different spatially distributed core sets, we refer to this as *spatiotemporal power multiplexing*. Note that a 64 core die (all cores are physically close and active) would have much higher power density and temperature compared to all the above cases (e.g. 110^0C in our simulation at predictive 16nm node).

3.2 Thermal equivalent RC network

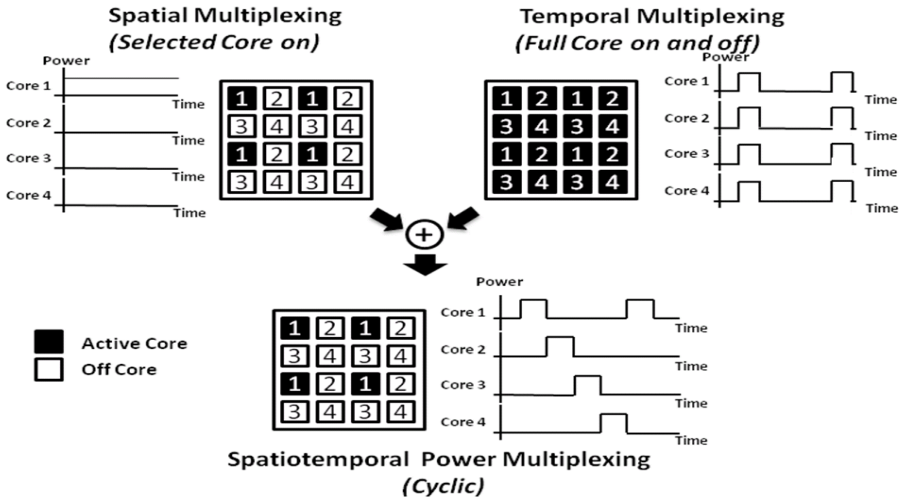

Fig. 5: Spatiotemporal Power multiplexing

First, we physically explain the advantage of the spatiotemporal power multiplexing using a thermal RC circuit (Fig. 6, represents a 4 core system). The heat generation is represented by the current sources and the voltage at each node represents the temperature. For the spatial ('*Selected Core on'*) one current source remains always active. The temporal ('*Full Core on and off'*) and spatiotemporal ('*Cyclic'*) cases consider pulsed current source (Fig. 6). For the temporal case, all sources are activated and deactivated simultaneously. In spatiotemporal mode, each source is active at the successive time-intervals and only one source is active at a time. For all these approaches average current (i.e. heat generation) over time is constant (equivalent to that of one current source).

Since a single core (i.e. current source) is always active with the s*elected core on* case it allows maximum temperature to rise at the active location and create high temperature difference. However, as the cores are not deactivated thermal cycles are low. The temperature behavior for the temporal and spatiotemporal case depends on the time-slice interval. First we explain the effect of time-slice interval on the spatiotemporal multiplexing based approach. If the current source is activated (deactivated) for a smaller time (i.e. smaller time-slice) the thermal capacitor at that location gets less time to charge (discharge). Thus a faster time-slicing reduces the voltage change across the thermal capacitors between two activations leading to a lower thermal cycles. Further, a faster switching allows less time to develop voltage difference across the lateral resistors resulting in less spatial variation in temperature (i.e. voltage) across the different nodes. Finally, a lower rate of increase in temperature in all locations implies a lower maximum temperature. Hence, for spatiotemporal power multiplexing a faster time-slicing can better manage the physical heat flow and leads to a reduced maximum temperature and lower spatial and temporal variations. However, the advantage diminishes if the time-slice is reduced beyond a certain point. The effect of

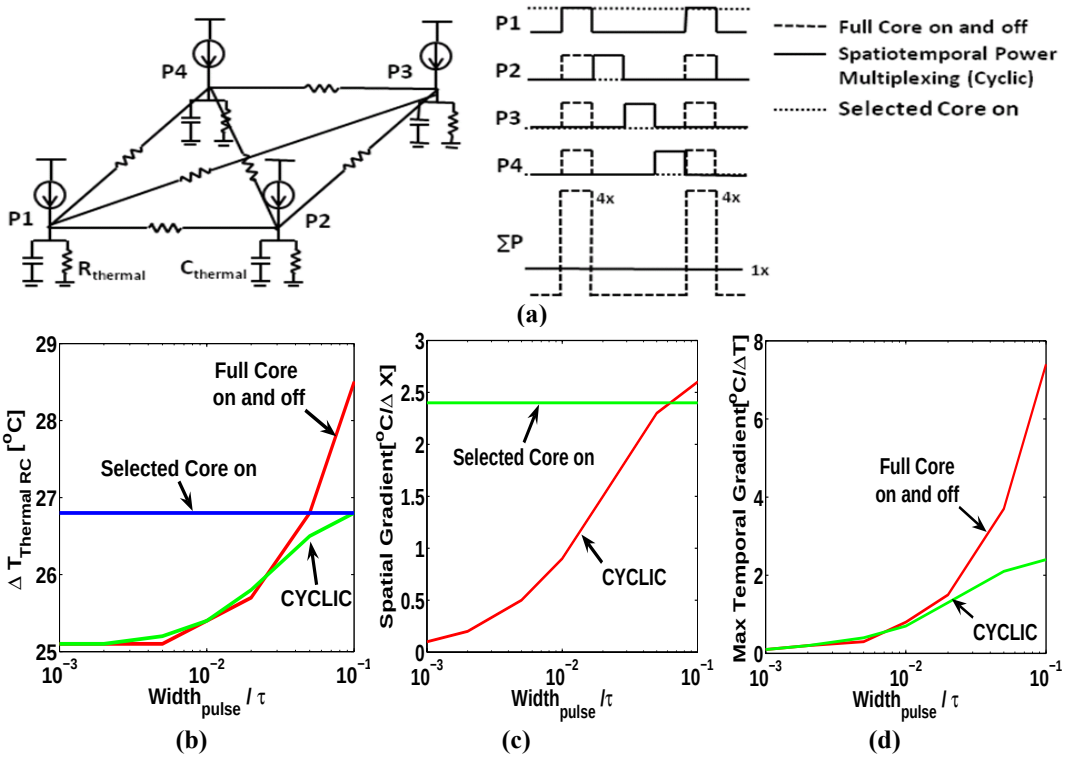

Fig. 6: RC Equivalent Circuit Analysis: (a) circuit, (b) maximum temperature, (c) spatial gradient, (d) temporal gradients.

time-slicing for the temporal (i.e. *full core on and off*) case is similar to the spatiotemporal case explained earlier. However, as the overall heat generation (i.e. total current) during the active period is much higher, the temperature rise at a much faster rate and to a higher value. Similarly, the temperature reduction is also higher. For a given timeslice, *full core on and off* case results in a higher maximum temperature and temporal gradients, particularly, for larger time-slice intervals (Fig. 6).

The above observation is verified using full-chip thermal simulations of a 256 core system with a targeted throughput equivalent to 64 active cores (Fig. 7). Fig. 7 shows that the *selected core on* case has low temporal gradient, but has high spatial gradient. The *full chip core on and off* case has higher temporal gradient but has lower spatial gradient. The spatiotemporal power multiplexing provides better temporal uniformity compared to the temporal method; better spatial uniformity compared to the spatial method; and lower maximum temperature compared to both approaches. We also observe that to maintain a target maximum temperature and temporal gradients, a much faster timeslicing is required for when heat is distributed only temporally.

4. Spatiotemporal Power Multiplexing Policies

We discuss three policies to determine the next set of active cores at the end of each time-slice interval.

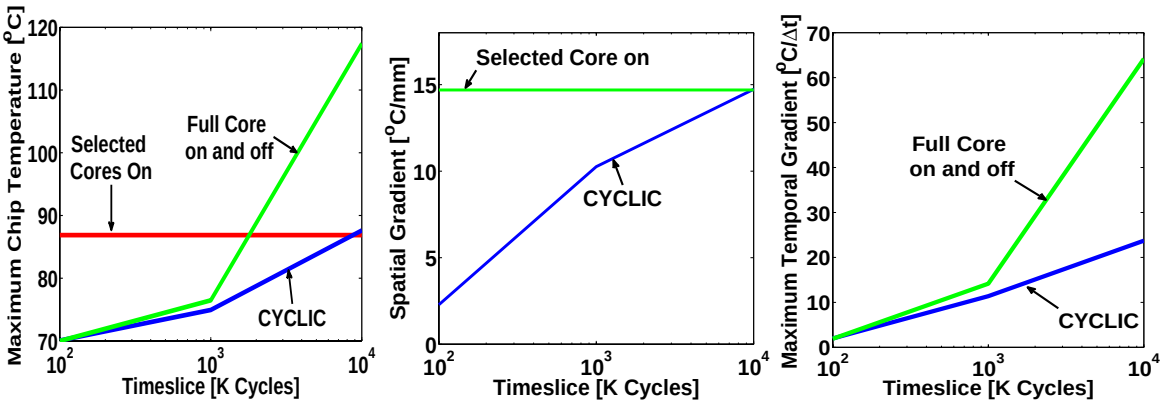

Fig. 7: Spatiotemporal Power multiplexing: (a) maximum temperature, (b) spatial gradient, and (c) temporal

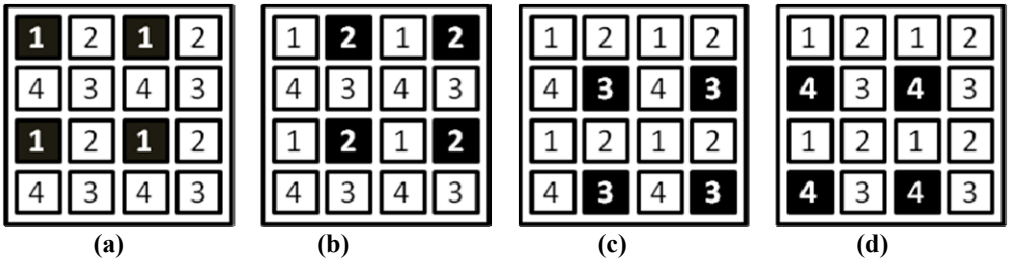

Fig. 8: Spatiotemporal power multiplexing with the cyclic policy.

4.1. Random power multiplexing

In this policy a random set of cores are activated at every timeslice, while currently active cores are deactivated. This can be achieved by using an on-chip random number generator that generates a new set of random number in every time slice. This is a proactive approach as no real-time temperature information is used. But the activation sequences of the cores are random resulting in increased complexity of state migration (i.e. transferring state information from one location to another). However, this policy allows a fine-grain control of the number of 'on' cores for highly time-varying parallel workload.

4.2. Cyclic Multiplexing:

A very simple but effective policy is the cyclic multiplexing policy (Fig.8). The entire chip is divided into smaller blocks called local areas. In our simulations, we divide the chip into 64 blocks of 2x2 cores each. To reduce thermal interaction between neighboring units, we initially placed active cores in checker board fashion. The active cores are then shifted after each timeslice, in a circular fashion as illustrated in Fig.8 [sequence: (a)→(b)→(c)→(d))]. This policy requires very simple hardware implementation and has a predefined core-activation sequence. As the cores are physically located in the vicinity of each other, cache sharing may also be possible which can significantly decrease the state transition time and energy (i.e. the energy associated with transferring data from one location to another). The primary disadvantage is that fine-grain control of

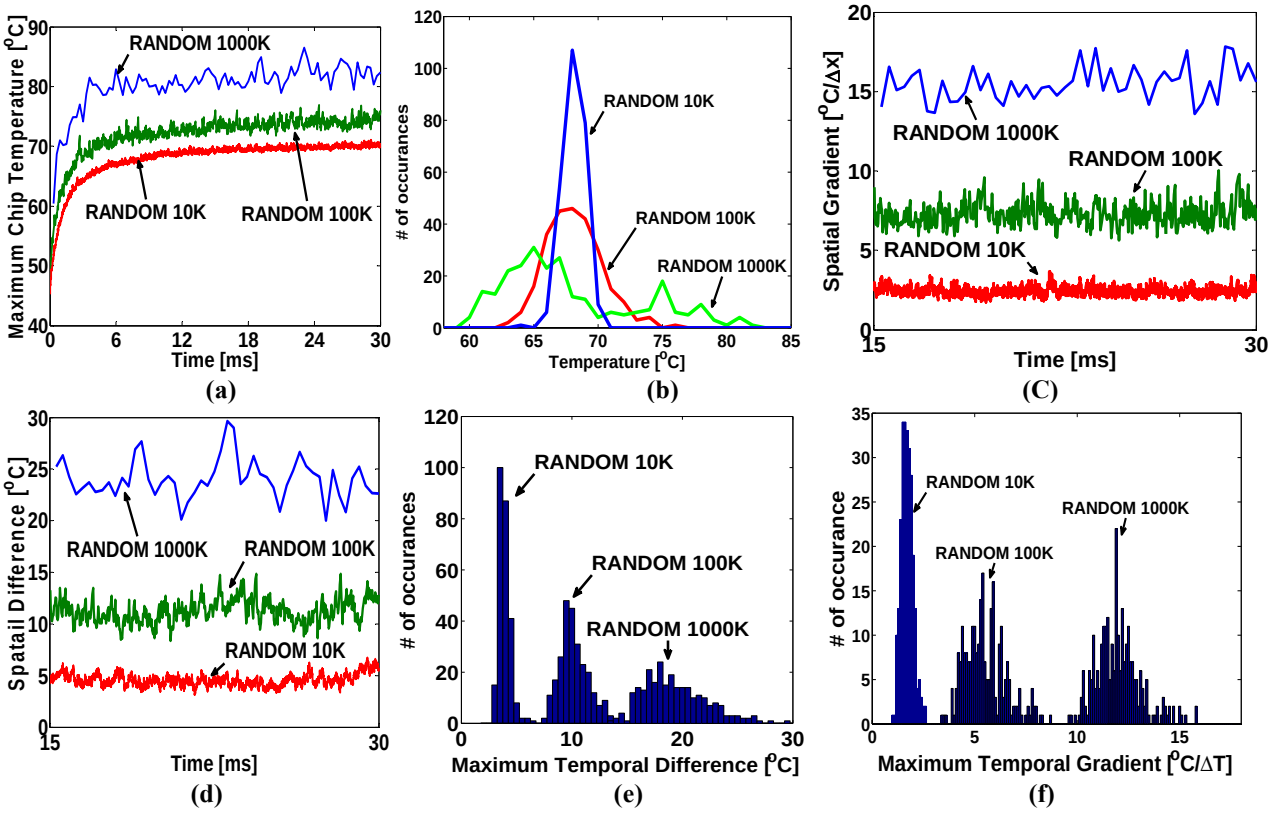

Fig. 9: Effect of timeslice on random migration (a)Max temperature (b) core-to-core temperature distribution, (c) spatial gradient (d) spatial difference, (e) temporal difference , and (f) temporal gradient. The 100K refers to 100,000 clock cycles i.e. 33μs for a 3GHz clock.

the number of 'on' cores is more difficult.

4.3. Global Coolest Replace

A partially reactive control approach using on-chip real-time temperature data is to replace the hottest 64 cores with the coolest 64 cores. This can help decrease the maximum temperature of the chip, as the coolest cores in the chip are turned on at the end of each timeslice. However, it requires complex implementation involving sampling of temperatures of all the cores and maintaining a sorted list of all the cores according to the instantaneous temperatures at the end of each timeslice. The core-activation sequence is also not predefined and the state migration can be costly (active cores may migrate to physically distant locations).

5. Modeling and Simulation Methodology

We demonstrate the effectiveness of spatiotemporal multiplexing on a tile-type homogeneous 256-core processor in predictive 16nm technology [11, 12]. A complete circuit and/or architectural level simulation of this predictive many-core processor is currently not feasible. Hence, we have used the technology driven power estimation tool (Intsim) presented by

Sekar et. al. to estimate the power of each core at 16nm [13-14]. The total power of a logic core is computed by taking the dynamic (assuming activity of 0.1) and leakage power of the logic gates including the wiring capacitance into account, the register and clock power, and the power associated with the repeaters in the signal interconnects [13]. The core power is determined by the system level parameters such as number of logic gates, the total silicon area of a core, operating frequency, etc. and the technology parameters such as feature size, device on/off current, threshold and supply voltage, interconnect dielectric thickness etc. The technology parameters are obtained from ITRS high-frequency logic process roadmap [1]. Each core is assumed to have 16 million gates, a local cache, and running at 3GHz frequency. Note the number of gates is much lower than that in the current microprocessor cores (to reflect simple cores) and the frequency is not scaled over the generations. For simplicity, we have neglected the cache power. Under these assumptions, we predict ~1.5W power per core. This is a reasonable estimate if we extrapolate the each core of 45nm microprocessors running at 3GHz which consumes ~100W of power [18]. Assuming a 2X reduction in switching capacitance in successive generations due to scaling and ~70% reduction of V_{DD} as predicted by ITRS from 45nm to 16nm node, we obtain ~6W of power per-core. If we consider 3-4X reduction in core complexity, this reduces to ~1.5W-2W. We consider different numbers and locations of the active cores in a chip to emulate the time-varying power map. While the different number of active cores modifies the total power, the different locations only modifies the power map. The generated time-varying power map is coupled to the full-chip thermal simulator HotSpot to obtain the on-chip temperature distribution [17]. The spatiotemporal power multiplexing policies are used to dynamically vary the on-chip power map and estimate their effect of temperature distribution. Although this approach disregards the exact architecture details of the core, it brings important insights into the thermal behavior of the many-core systems.

6. Impact of Spatiotemporal Power Multiplexing

We analyze the impact of spatiotemporal power multiplexing based proactive control mechanism on the thermal field. We show our results considering a throughput target of 64 active cores. Similar trends were observed for other throughput values as well.

6.1. The effect of the Timeslice Interval

As the thermal behavior of the chip is similar to a distributed RC network, it is clear that a larger timeslice will lead to a higher maximum temperature and higher spatiotemporal non-uniformity (as explained in section 3). This is captured in Fig. 9 which shows the effect of time slice on maximum temperature, spatial gradient, temporal gradient, spatial difference, and

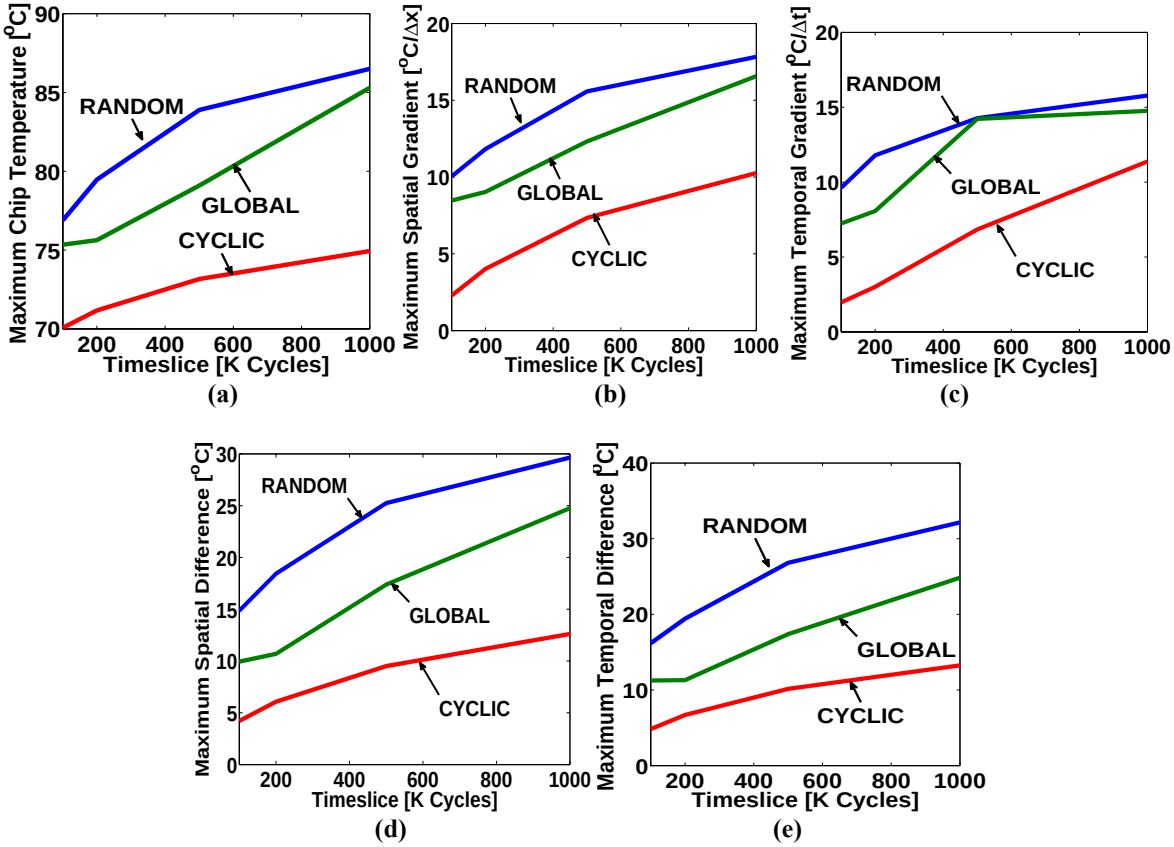

Fig. 10: Effect of policies: (a) Max temperature, (b) spatial gradient, (c) temporal gradient, (d) spatial difference, and (e) temporal difference.

temporal difference. A smaller timeslice shows reduced maximum temperature (Fig. 9a). However, the improvement from 1000K to 100K is more pronounced than that from 100K to 10K clock cycles. A faster migration implies that core-to-core temperature variation is much less (Fig. 9b). Both the spatial difference and spatial gradient reduce with a faster timeslice (Fig. 9c,d). This implies that on-chip thermal variations are more uniform at all time instants which imply better functional reliability, cooling efficiency, and physical reliability. Finally, both the average value and core-to-core variations of temporal gradients and temporal difference of all cores reduce significantly at a faster timeslice (Fig. 9e,f). This implies all the cores have a lower thermal cycle (better reliability). Further, the reliability degradations across cores are similar.

6.2. The effect of Migration Policies

For a fixed timeslice, the effectiveness of the spatiotemporal power multiplexing depends on the policy. Our analysis shows that the cyclic policy realizes a much lower maximum temperature, spatial difference/gradients, and temporal difference/gradients (Fig. 10), i.e., for a fixed timeslice cyclic migration provides better uniformity. On the other hand, for a given maximum temperature or spatiotemporal uniformity target, the cyclic migration can be performed much less

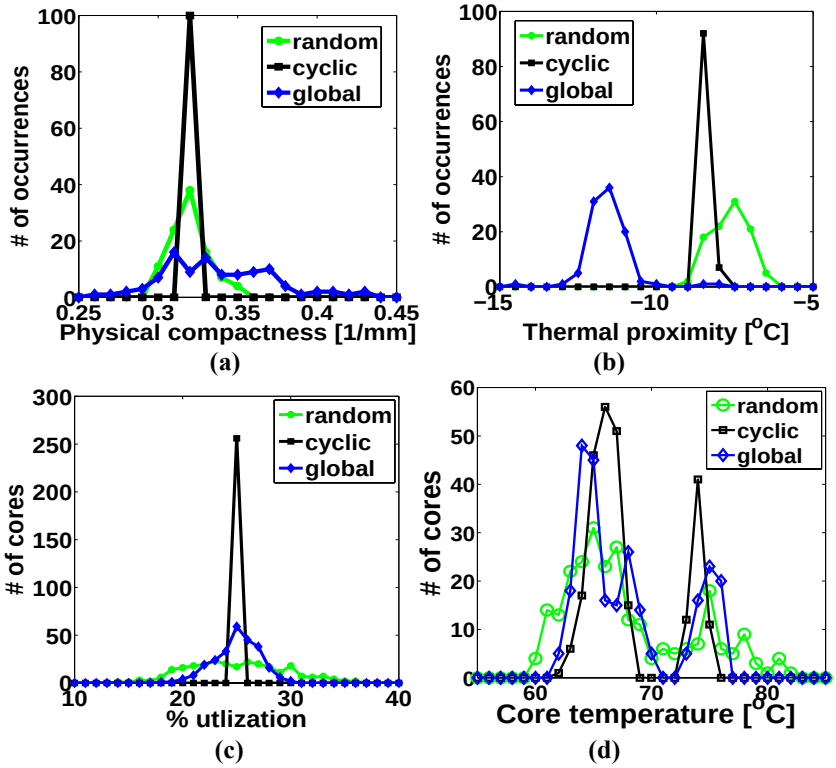

Fig. 11: Effect of migration policies: (a) Compactness (b) Proximity, (c) Utilization, and (d) Core-to-core temperature variation

frequently. Hence, a co-design of the migration policies and timeslice is necessary to maximize the advantage. To understand the reason for the above mentioned differences, we estimate the physical compactness of the set of active cores over the entire time duration. The physical compactness is measured by computing the core-to-core distance between all pairs of cores in an active core set and estimating the inverse of the average of all the distances. A higher value of compactness implies the active cores are physically close which results in a faster rise in temperature. Fig. 11a shows the physical compactness distribution for different policies. The active core sets in the cyclic policy always has the same compactness (because of the checkerboard configuration). Note that both global coolest and random policy generates several instances of very high compactness leading to higher maximum temperature and spatiotemporal non-uniformity. This explains why cyclic policy performs better. However, to explain the difference between global coolest and random policies we estimate the thermal proximity, defined as the difference between the average temperature of current set of active cores and the next set of cores to be activated at each timeslice intervals. A more negative difference implies that migration is performed to a set of more cool cores. As expected, the global coolest policy always migrates to a much cooler set of cores compared to the random policy resulting in a lower maximum temperature compared to the random migration (Fig. 11b).

One point to note here is that by executing power multiplexing, the utilization of any core per unit time is reduced. The cyclic policy ensures all cores have equal utilization (25%). However, the core utilization can vary between 15-35% for random and 20-30% for global coolest replace policies (Fig. 11c). This implies that while cores age at a similar rate for cyclic policy, the global coolest replace and random policy can result in different aging rates. The core-to-core temperature distribution is much tighter for the cyclic policy for a fixed timeslice (~1000K for Fig. 10d) condition.

6.3. Effect on Delay and Leakage Variations

By providing lower temperature and more uniformity, spatiotemporal power multiplexing helps to reduce the core-to-core delay and leakage variations. The delay variations are estimated considering an 8 stage FO4 ring-oscillator using 16nm predictive models [15]. Fig. 12 shows the histogram of delay and leakage increase (estimated for an inverter) compared to 45^{o}C temperature, considering random migration policy for different timeslice. A lower on-chip delay and leakage variation implies a better functional reliability. From Fig. 11d we expect that cyclic policy will provide even lower delay/leakage spread.

6.4. The Power overhead and overall Power Saving

We estimate the power overhead of activation-deactivation of all possible cores at the same time (i.e. 64 cores are turned off and 64 cores are turned on). We assume during transition all internal nodes make transition (i.e. entire core capacitance switches). The switched capacitance is estimated from the power-per-core. Further, we consider all registers need to make transition for state migration. Finally, the transition of the supply network capacitance (estimated based on ITRS roadmap) and 10% additional core level decoupling capacitance is considered. However, as the per core switch capacitance is reduced at deep nanometer nodes and time-slicing is performed at very low frequency, the power overhead is very low (Table. 1). This is overshadowed by the dynamic and leakage power saving achievable by operating the system at a lower voltage. This is permissible as lower temperature implies that same target frequency can be achieved at a lower voltage (Table. 1, estimated based on the ring-oscillator simulation in predictive 16nm models [15]). Compared to ‘Selected Core on’ (i.e. with no transition case), even with random migration policy, the overall power saving (i.e. after accounting for the additional transition power) can be ~5-6%.

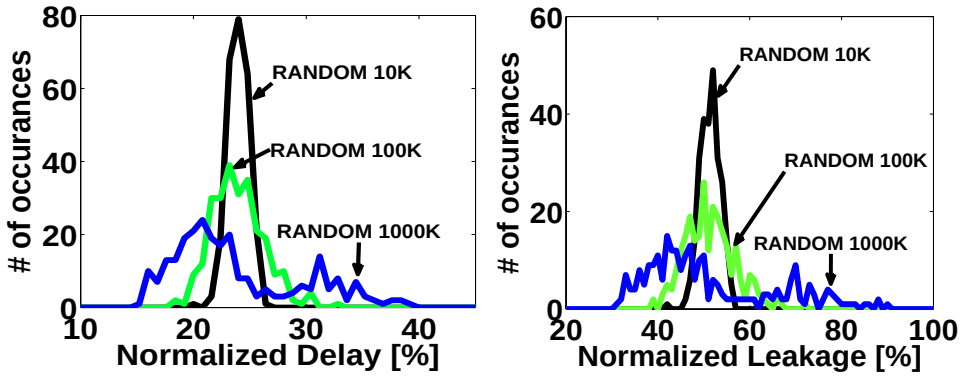

Fig. 12: Effect on delay (left) and leakage (right) distribution.

Table. 1: Power overhead and saving [% of total chip power]

Timeslice Interval	100K	200K	500K	1000K
Power overhead	0.028	0.014	0.006	0.003
Power Savings	6.1	6.1	6.1	6.1

Table. 2: Results with 4X lower Power Density and 10000K Timeslice

	Max Temp. [°C]	Spatial Difference [°C]	Spatial Gradient [°C/Δx]	Temporal Difference [°C]	Temporal Gradient [°C/Δt]
Random	64.63	17.32	11.53	18.19	14.95
Cyclic	59.52	11.42	8.61	11.92	10.38

6.5. Discussion

We analyze the impact of variation in the active power of a core (Normal distribution with standard deviation of 20% of nominal value) and non-zero power in the deactivated mode (finite leakage in case of state-preserving supply gating i.e. voltage is not reduced to '0'). Our analysis shows that these conditions will increase the maximum on-chip temperature requiring a smaller time-slice interval. Also, we study the effect of a lower power density for each core. Since lower power density implies reduced heat generation in each active core, a much larger time-slice interval can be used to control the maximum temperature and uniformity for all policies. As in case of the higher power density, the cyclic policy shows better performance compared to the random approach (Table. 2)

7. Summary

We have analyzed the unique thermal behavior in many-core system to show that the thermal field in many-core systems can have significant spatiotemporal non-uniformity. Hence, global management of the thermal field to reduce such non-uniformity along with the maximum temperature is necessary in many-core systems. We have presented spatiotemporal power multiplexing at fixed time intervals (timeslice) as a feasible approach for thermal field management. The key idea is to distribute the generated heat in space and time to reduce temperature as well as maintain uniformity. Several migration

policies are discussed. The effects of migration policies and timeslice on the thermal management are estimated. Our analysis shows that timeslicing provides effective proactive management of the thermal field to reduce maximum temperature, spatial and temporal gradients.

8. References

[1]. International Technology Roadmap of Semiconductors, 2008.

[2]. S. Borkar, Thousand Core Chips - A Technology Perspective, Design Automation Conference, 2007. DAC '07, pp. 746 – 749

[3]. A. Coskun, et. al, "Static and Dynamic Temperature-Aware Scheduling for Multiprocessor SoCs", IEEE TVLSI 2008

[4]. D. Brooks, et. al, "Dynamic Thermal Management for High-Performance Microprocessors", HPCA 2002

[5]. P. Chaparro, et. al, "Understanding the Thermal Implications of Multicore Architectures", IEEE Transaction on Parallel and Distributed Systems, 2007.

[6]. R. McGowen, et. al, Power and temperature control on a 90-nm Itanium family processor, IEEE JSSC, January 2006. pp. 29-237.

[7]. J. Tschanz, et. al., "Adaptive Frequency and Biasing Techniques for Tolerance to Dynamic Temperature-Voltage Variations and Aging," ISSCC, 2007.

[8]. J Srinivasan, et. al, "The Case for Lifetime Reliability-Aware Microprocessors," ISCA, 2004, pp. 276- 287.

[9]. V. Gektin, et. al, "Substantiation of Numerical Analysis Methodology for CPU Package with Non-uniform Heat Dissipation and Heat Sink with Simplified Fin Modeling", Thermal and Thermomechanical Phenomena in Electronic Systems, 2004.

[10]. K. Sankaranarayanan, et., al, "A Case for Thermal-Aware Floorplanning at the Microarchitectrual Lelvel", Journal of Instruction-Level Parallelism, 2005.

[11]. S. Bell, et. al., "TILE64 Processor: A 64-Core SoC with Mesh Interconnect," IEEE ISSCC, 2008, pp. 88-598.

[12]. S. R. Vangal, et. al. An 80-Tile Sub-100-W TeraFLOPS Processor in 65-nm CMOS, *IEEE Journal of Solid-State Circuits*, vol. 43, no. 1, Jan. 2008, pp. 29-41

[13]. D. C. Sekar, et. al, Intsim: a CAD tool for optimization of multilevel interconnect networks, IEEE/ACM ICCAD, 2007.

[14]. D. C. Sekar, Optimal Signal, Power, Clock and Thermal Interconnect Networks for High-Performance 2D and 3D Integrated Circuits, PhD Dissertation, Georgia Institute of Technology, 2008.

[15]. Predictive Technology Model: http://www.eas.asu.edu/~ptm/

[16]. W. Huan, et. al, "Many-core design from a thermal perspective," DAC 2008.

[17]. W. Huang, et. al, "Hotspot: A Compact Thermal Modeling Method for CMOS VLSI systems," IEEE TVLSI, pp. 501–513, May 2006,

[18]. http://www.intel.com/products/processor_number/chart/xeon.htm