Hotspot Monitoring and Temperature Estimation with Miniature On-Chip Temperature Sensors

Pavan Kumar Chundi, Yini Zhou, Martha Kim, Eren Kursun, Mingoo Seok
Columbia University, New York, NY

Abstract

This paper presents analysis and evaluation of the impact of size and voltage scalability of on-chip temperature sensor on the accuracy of hotspot monitoring and temperature estimation in dynamic thermal management of high performance microprocessors. The analysis is based on both the layout level and the system level across state-of-the-art sensors in terms of accuracy, voltage-scalability, and silicon footprint. Our analysis shows that a sensor having compact footprint and good voltage scalability can be placed on exact hotspot locations, typically among digital cells, significantly improving accuracy in tracking hotspots and estimating temperature of microarchitecture blocks, as compared to two other sensors that have higher sensor-circuit accuracy, large footprint and little voltage scalability limiting flexible placement.

1. INTRODUCTION

Transistor scaling has led to a significant increase in the power density of high performance microprocessors, which makes them thermally limited. This mandates most of the high-performance microprocessors to employ dynamic thermal management (DTM) to maximize the performance, energy-efficiency, and reliability of the system [1,2,3]. DTM monitors the temperature at multiple points on a microprocessor chip, using the readings to trigger temperature reduction techniques when at the thermal limit.

DTM uses a temperature sensor network (TSN) for monitoring chip temperatures. A TSN consists of multiple temperature sensors, read-out circuitry such as analog-to-digital converters, and possibly post measurement processing framework.

Today's TSNs, however, have low accuracy. To avoid aging effects such as temperature bias instability (BTI) [4] or worse, burning a chip, it is typical to design a TSN to overestimate temperature. Prior work shows that the resulting margins including the margin to ensure overestimation can cause excessive throttling reducing the performance across different workloads [5-10].

The low accuracy of on-chip hotspot monitoring and temperature estimation stems from two sources: (i) process and voltage variations in the sensor circuit itself, which we define *sensor error;* (ii) from the distance between a hotspot and the nearest sensor, which we define *distance error*.

In estimating hotspot temperature, a significant portion of the total error is attributable to the distance error. One reason is that the advances in sensor circuit designs have been reducing sensor error [11,12]. The other reason is that highly-scaled transistors increase power density and therefore local thermal gradients. Our simulation confirms the dominance of distance error (Fig. 1). When we place nine sensors per core uniformly for a quad-core microprocessor running 12 different workloads, the distance error contributes more than 90% of total error in average. Note that the sensor error has different values across benchmarks because each sensor is calibrated at 50°C, and the sensor error grows as it measures temperature away from the calibrated point.

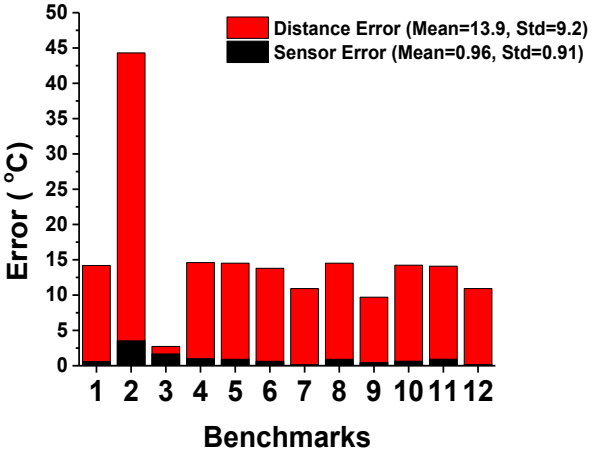

Figure 1. Error breakdown across benchmarks. In this experiment nine sensors per core are placed uniformly

Large errors and the significant contribution of distance error have motivated various studies on optimal sensor placement. However, most such prior studies assume the sensor is point-sized and can be placed anywhere on a chip [5-10], which is not very practical. To place a sensor inside of digital circuits implemented by the standard cell design flow, for example, the sensor needs to be very compact so as not to perturb digital cell placement and routing and timing closure. Furthermore, the sensor needs to operate from a digital power grid that can be scaled down to near-threshold regime (e.g., 0.5-0.7V) for supporting Dynamic Voltage Frequency Scaling (DVFS). But many of the existing sensors cannot operate at such low supply voltage (V_{DD}). Without good voltage scalability, a separate power grid or local regulation circuit is necessary to provide V_{DD} to sensors, incurring large area overhead.

There is no prior work that takes into account the effect of sensor size and voltage scalability on optimal sensor

978-1-5090-6023-8/17/$31.00 ©2017 IEEE

placement. Therefore, in this paper we address these issues. First, we have studied sensor impact on layout, to see if and how state-of-the-art sensors [11,12] alter the temperature map and the critical path delay of the design. Based on these analyses, we create constraints on sensor placement that is more realistic with respect to the impact on layout. Then, we employ those constraints at the system level with a suitable set of assumptions for different kinds of sensors. It is observed that a small sensor has a very small error (3.5°C) in the worst case in estimating the hottest temperature in a core compared to medium (24.5°C) and large sensor (44.3°C). It is also better at tracking individual hotspots with an error of 4.8°C compared to 6.3°C for medium sized sensor and 12.3°C for large sensor in case of the worst benchmark.

The remaining paper is organized as follows. In Sec. 2 we will discuss three state of the art temperature sensors on which we base our study. In Sec. 3, we study the effect of a miniature temperature sensor on the temperature map and the impact of sensor size on the critical path delay of the design. In Sec. 4, we draw up assumptions for the placement of the sensor at system level and compare the accuracy of hotspot monitoring for three different representative sensors. Finally, we conclude in Sec 5.

2. TEMPERATURE SENSOR CIRCUITS

Temperature sensors are based on the threshold voltage (Vth) [11,13,14], sub-threshold leakage [12], or frequencies of ring oscillators (RO) [15,16] in CMOS or on the junction voltage of bipolar transistors (BJT) [17,18,19]. Vth-based sensors generally achieve compact footprints and have better voltage scalability. However, they are less robust against process variations and device aging effects since Vth is sensitive to process variation and aging. On the other hand, the designs using BJT junction voltage are generally more accurate and reliable but exhibit footprints larger than 1000-2,000 μm^2. The BJT based sensors also have limited voltage scalability, which makes it difficult to use supply voltage below 1V in those sensor designs.

Fig.2 compares the area and accuracy of the recent temperature sensors, with each point annotated with the minimum operational supply voltage. From these, we choose three that are Pareto optimal (starred points in Fig.2) in terms of area and accuracy.

The first chosen sensor is called Sensor-Small (Sensor-S), which is proposed by Kim et al. and uses a simple front-end circuitry which includes just two PFETs. This makes the area of this sensor front end as small as 30.1μm^2 [11]. This is about 1 to 2 orders of magnitude smaller than BJT based design. In addition, the sensor can operate at 0.4 V.

The second one is Sensor-Medium (Sensor-M), designed by Saneyoshi, et al., which exploits the temperature dependency of PFET off-leakage current (I_{OUT} in Table 1) to measure temperature [12]. The circuit consists of several PFETs and switches controlled by a 3-bit select signal. This sensor design achieves a smaller error than Sensor-S (of ±1.55°C) but consumes nearly ~40X larger area (1,255μm^2 per front end).

Finally, the third one is called Sensor-Large (Sensor-L), designed by Souri, et al. It is based on BJT junction voltage which is proportional to absolute temperature (PTAT) [19]. It also employs a chopping technique to reduce low-frequency noise. Each of these front ends are large (~10,000μm^2) but achieves the smallest error of the three (±0.4°C).

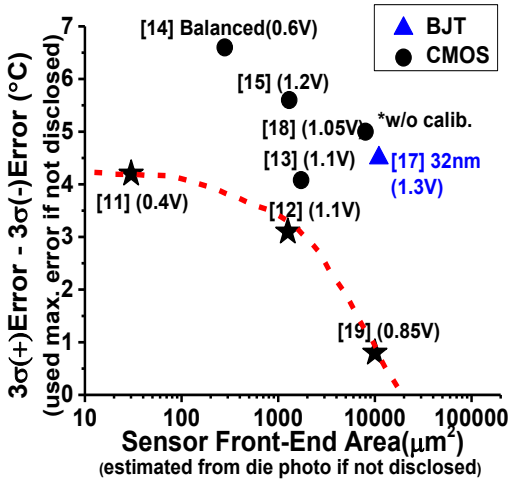

Figure 2: Recent thermal sensors and their trade-offs between ±3-σ error and sensor front-end area. The minimum supply voltage for the sensor is indicated in parenthesis. The star symbols represent the designs used in this paper.

3. LAYOUT LEVEL ANALYSIS

In many prior studies, the sensor is assumed to be point-sized and can be placed anywhere on a chip [5-10]. However, such assumption is not very practical. The placement of the sensor among densely-placed digital cells can affect signal routing and therefore timing closure. Sensors that cannot operate under sub-1V V_{DD} cannot use digital power grids and require additional power routing and local regulation, incurring large overhead. Therefore, in this section, we examine how temperature sensor's placements impact the size and location of hotspots and digital critical path delay.

3.1. Effect of sensor size on the critical path delay

The effect of sensor size on the critical path delay of the design is very important, as one would not want to place a sensor where it would degrade microprocessor performance. A study on the effect of sensor size on the critical path delay therefore is necessary to understand the requirement on the size of the sensor so that performance of the system is not affected.

We perform this study on two 32-bit multipliers: both have the same netlists but one is placed and routed at 75% standard cell area utilization and takes 110x115 μm^2 area and the other at 60% utilization and 110x144 μm^2. A sensor

is placed at the center of the design core and then the standard cells are placed and routed. We optimize the placement and routing to meet a timing constraint of 1.2ns for the case of the multiplier having 75% utilization and that of 1ns for the multiplier having 60 % utilization in a 65nm.

Fig. 3 plots the slack as a function of the sensor size as a percentage of the core area. As the sensor grows, so does the critical path delay, because large sensors occupy significant design area, separating blocks on the critical path and causing significant increases in wire delay. We find that the notable increase appears when the sensor is larger than 5-6% of core area or larger than 700 μm^2.

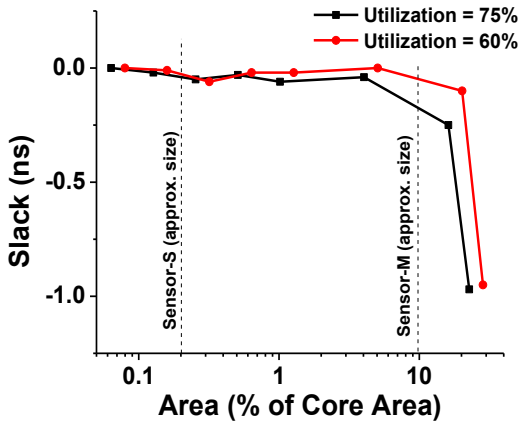

Figure 3. The slack as a function of sensor size as a percentage of the area of a 32-bit multiplier

3.2. Effect on temperature map

We also perform a case study for understanding the effect of sensor placement on thermal characteristics of digital circuits. In this study, we use a 32-bit multiplier, whose spatial power dissipations and floorplan from post APR simulations are provided as input to the thermal analysis software Hotspot [27] for generating temperature maps.

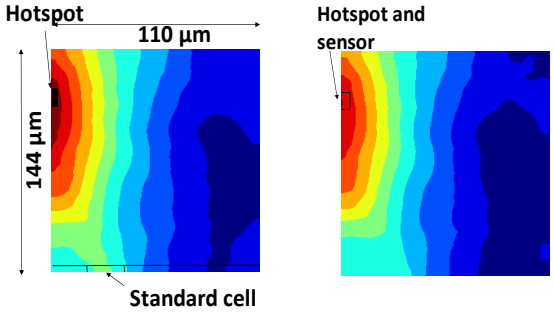

Figure 4. Temperature maps of a 32-bit multiplier (a) before and (b) after the placement of a sensor

Fig. 4(a) shows the temperature map of the 32 bit-multiplier before the placement of Sensor-S. The hotspot in this temperature map is at the left edge of the design. Fig. 4(b) shows the location of the placed sensor on the hotspot. The placement location is a small white space that would accommodate a decoupling capacitor. Since the Sensor-S is very small, we can place it in the space without any disturbance to the digital cells in the design. Since the Sensor-S can share the digital power grid, it requires no additional power lines or local voltage regulation.

By contrast, it is impossible to place a large sensor (>1000 μm^2) without disturbing the remaining standard cells, because the area of the core itself is 110x144 μm^2. The poor voltage scalability of Sensor-M and Sensor-L does not allow one to connect the sensors to the existing power network. Therefore, a large sensor would have to be placed outside the 32-bit multiplier, while the smaller ones with sufficient voltage scalability allow easy placement and in the exact location of the hotspot. This is particularly critical to attain high accuracy if there is a strong temperature gradient.

4. SYSTEM LEVEL ANALYSIS OF PLACEMENT OF SENSORS

We now take three different-sized sensors, and evaluate them when incorporated into a full temperature sensor network for a quad-core multiprocessor. The three sensors are: extremely small with high voltage scalability (Sensor-S [11], 30.1 μm^2), medium (Sensor-M [12], 1255 μm^2), and large (Sensor-L [19], 10000 μm^2). We will evaluate the sensors, in their network context, by comparing the sensors on two different error metrics and arrive at the best choice of the sensor from this study.

Table 1: Microarchitecture parameters

Parameter	Value
Technology node	32 nm HP; aggr. interconnect
Supply voltage	1.35 V
ISA	x86-64, Gainestown
Number of cores	4 @ 3.6 GHz
pipeline	Out-of-Order
L1-I cache per core	32 kB, 4-way assoc., private
L1-D cache per core	32kB, 8-way assoc., private
L2 cache per core	256 kB, 8-way assoc., private
L3 cache	8 MB, 16-way assoc., shared
Area	102.074mm^2

Table 2: Benchmarks

1	blackscholes	PARSEC
2	cholesky	SPLASH
3	fmm	SPLASH
4	lu.cont	SPLASH
5	lu.ncont	SPLASH
6	water.nsq	SPLASH
7	radix	SPLASH
8	water.sp	SPLASH
9	streamcluster	PARSEC
10	swaptions	PARSEC
11	w1 (perlbench,bzip,gcc,games)	SPEC
12	w1 (perlbench, sjeng, libquantum, calculix)	SPEC

4.1. Creating a Temperature Map

We create temperature maps of a microprocessor across several benchmark software. For the microprocessor, we use the Sniper simulator [20] with Table 1 summarizing the microarchitecture of the microprocessor. We also consider both multi-thread and multi-program workloads. For multi-thread, we use seven SPLASH-2 [21] and three PARSEC [22] benchmarks. For multi-program, we use two workloads, named w1 and w2, each of which consists of four randomly-chosen benchmarks out of the twelve in SPEC CPU2006 [23]. All the benchmarks used are listed in Table 2. Using the McPat power/area model [24], we collect area and power information of the microprocessor. From such information, we construct a processor floorplan. Finally, with a floorplan and power trace, we use Hotspot [25] to generate temperature maps.

In the thermal map generation, we perform the iterations for modeling temperature-aware leakage dissipation. In the first iteration, we have McPat produce the initial power traces based on a default temperature (330K). After Hotspot simulates temperatures using the initial power traces, we have McPat to iteratively simulate power traces with the new temperature information. This process continues until the temperature map from Hotspot match the temperatures used in McPat. This is performed for all the 12 benchmarks and the obtained temperature maps represent the ground truth for chip temperatures in our experiments. The final temperature maps have a resolution of 256x256 points and a spatial resolution of 40x40 μm^2 for the targeted 100-mm^2 quad-core microprocessor.

4.2. Sensor Placement Constraints

As we have found in Sec. 3, the size and the voltage scalability of sensors impose different constraints in their placement in a microprocessor. Sensor-S, because of its small size and high voltage scalability we can assume that it can be placed anywhere on the chip. This is supported by the fact that Sensor-S had little impact on critical path delay and did not change the temperature map (see Figs 3 and 4).

On the other hand, we assume that Sensor-M can be placed at the edge of blocks such as Execution Unit, Memory Management Unit, Renaming Unit, Instruction Fetch Unit and Load Store Unit (Please refer Fig. 5 for the relative sizes and locations of these blocks). This is because Sensor-M requires an additional power grid due to the limited voltage scalability, and it makes the addition of the power grid simple and little invasive to place the sensors outside of a block.

Sensor-L would be even more restricted with respect to placement. So, we assume that the Sensor-L will be placed uniformly in each core. This helps in comparison in case one wants to avoid time-consuming thermal simulations and instead utilize post-processing frameworks from spatially uniformly-placed sensors for tracking hotspots and estimating temperature [26]. Note that for the purpose of system level placement we model every sensor to be placed in one 40x40-μm^2 pixel in the temperature maps, even though Sensor-L would occupy more than five pixels on the temperature map.

The sensor placed at different locations sample the temperature values from the ground truth. We take inherent sensor-circuit error numbers from the measurements of each circuit reported in the literature. Note that the sensor error is dependent on the difference between the temperature measured and the temperature at which the sensor was calibrated. In this paper, we assume that all the sensors were calibrated at 50°C.

4.3. Uniform v. Targeted Sensor placement

In this section, we first compare Sensor-L that is uniformly placed and Sensor-S that is placed at potential hotspots. Specifically, we find all the potential hotspots from the thermal simulations across benchmarks, and then use the k-means clustering method [16] on the union of local maxima of the temperature maps of all the workloads, separately on each core. We use several 'k' values that determine the number of sensors per core. A Sensor-S is placed at the hottest local maxima in each cluster. Fig. 5 shows the hotspots and the locations of Sensor-S placed based on the above method when k is 5 (i.e., 5 sensors per core; 20 sensors for 4 cores).

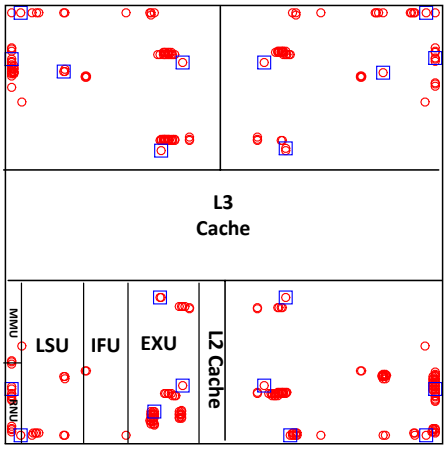

Figure 5. Floorplan of the microprocessor with the locations of sensors (blue squares) and the local maxima (red circles). We place 5 sensors in each core.

We also define the core hotspot error in the estimated hotspot temperature of core 'i' as:

$$\boldsymbol{Core\ HS\ Error_i} = |\boldsymbol{Max(T_i)} - \boldsymbol{Max(S_i)}|$$

where T_i is the set of temperature values for i-th core and S_i the set of sensor data for that core. Then, the error for a particular benchmark 'b' across all the cores is defined as:

$$\boldsymbol{HS\ Error[b]} = \boldsymbol{Max}(\{\boldsymbol{Core\ HS\ Error_i} | \boldsymbol{i} \in \{\boldsymbol{1\ to\ 4}\}\})$$

Fig. 6 shows the scaling of the HS error averaged across the 12 workloads as a function of the number of sensors embedded in each core. We see that 15 Sensor-L's would be

needed to reduce the average HS error below 10°C. This can impose a non-negligible amount of area overhead as each sensor takes about 10,000 μm².

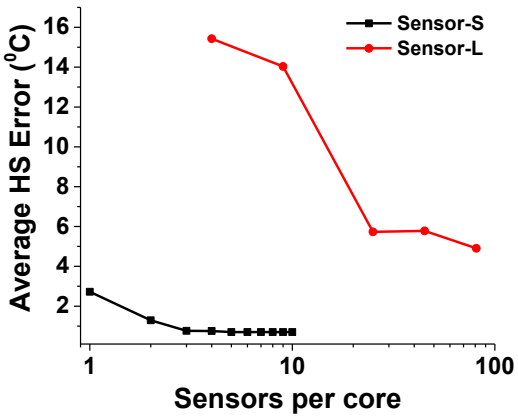

Figure 6. HS error averaged across all the benchmarks for uniformly placed Sensor-L and targeted-placement of Sensor-S

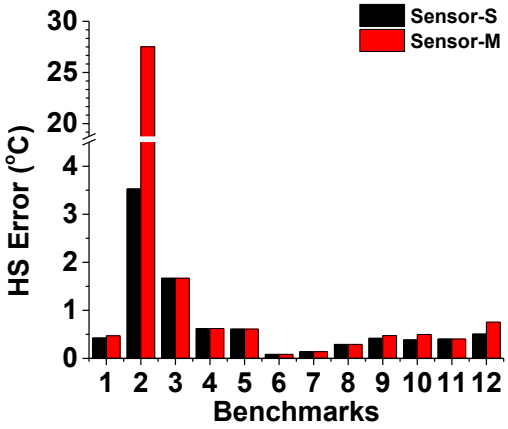

Figure 7. HS Errors of the placement of Sensor-S's and Sensor-M's

In contrast, as shown in Fig. 6, we find that having five Sensor-S's per core can achieve <1°C HS error averaged across the 12 benchmarks, if we place them at potential hotspots based on the k-means clustering method. The small error is the result of nearly entirely eliminating the distance error: i.e., the remaining error is mostly attributed to circuit error. Note that having more than five sensors improves little as all the hottest spots in each core in all the benchmarks are being directly tracked by the sensors on them.

4.4. Targeted Placement of Sensor-S and Sensor-M

We now compare Sensor-S and Sensor-M, both of which are placed on hotspots. The significant differences in size and voltage scalability allow only Sensor-S to be placed inside of digital blocks, which can have impact on the error in hotspot monitoring and fine-grained temperature estimation. We place five sensors in each core. Sensor-M is placed on the closest edge of the block in which the hottest spot for that cluster is located while Sensor-S is placed exactly at the hottest spot.

Fig. 7 shows that Sensor-S and Sensor-M are almost equally accurate in most of the benchmarks except the benchmark-2 (cholesky) where the Sensor-S largely outperforms Sensor-M. This is because the temperature gradient in this benchmark is very high and therefore having the sensor at the edge of the block closest to the hotspot still causes large HS error.

4.5. Fine-Grained Temperature Estimation

In Secs. 4.3 and 4.4, we focus on hotspot tracking with complete avoidance of temperature underestimation. This is because it is typical to avoid catastrophic effects such as burning a chip. Although it may track hotspot temperature, this practice can overestimate temperatures of microarchitecture blocks and a chip. If one wants to perform fine-grained thermal management, such as redirecting workloads, scheduling instructions, etc. there is a need for more fine-grained thermal sensing.

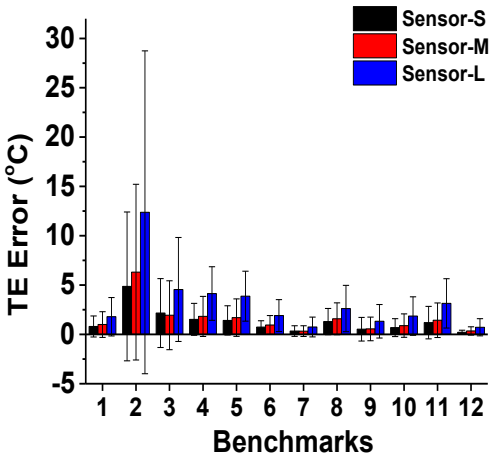

Figure 8. TE Error across benchmarks, together with the ±σ errors across local temperature maxima.

Therefore, in this section, we focus on such errors using similar evaluation framework. We use the same placement constraints and schemes as before for the three sensors. We also create the Temperature Estimation Error (TE Error) metric. It is the average error in estimating the temperature of all the local maxima, which can be formulated as:

$$\boldsymbol{TE\ Error[b] = Mean[|T_m - S_{f(m)}|]} \qquad (2)$$

where T_m is the temperature of the m-th local maximum, $f(m)$ is the cluster ID (from the k-means clustering method) of the local maximum, $S_{f(m)}$ is the temperature readings from the sensor associated with that cluster, and 'b' is the benchmark index. For the Sensor-Ls which are placed uniformly, we associate each sensor to one of the clusters based on the minimum distance from the centroid of that cluster.

Fig. 8 shows TE errors of the placements of the three types of sensors. Sensor-S achieves a moderate amount of

reduction in TE Error over Sensor-M. This is because Sensor-M underestimates the hottest spot while Sensor-S overestimates relatively cool local maxima in the same cluster.

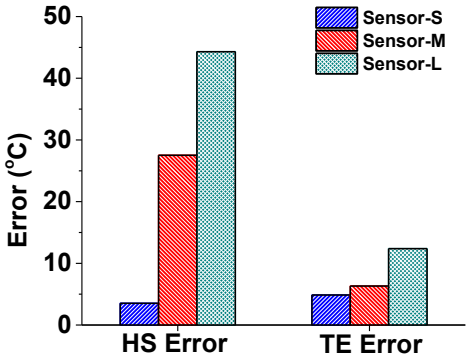

Figure 9. HS Errors and TE Errors from the benchmark 2 (the worst-case one).

Finally, Fig 9 summarizes the HS Error and TE Error of the placements of three sensors for the benchmark 2. The Sensor-S shows significant reductions in HS Error and moderate reductions in TE Error, confirming the importance of size and voltage scalability on accurate on-chip thermal monitoring.

5. CONCLUSIONS

In this paper, we have examined how the size and voltage scalability of a temperature sensor impacts their placement on a chip. Because small sensors with voltage scalability can be placed in a small white space among densely placed digital cells while sharing digital power grid, their placement is little restricted, allowing designers to place them at or very near the anticipated hotspots. Thus, despite small sensors being typically less accurate in their standalone evaluation than large ones, overall, they can achieve more accurate hotspot tracking and fine-grained temperature estimation. In addition, we find that the smaller the sensor, the smaller the impact on digital circuit timing closures, again supporting the use of small sensors.

ACKNOWLEDGEMENT

This work is partially supported by the NSF (CCF-1453142, CCF-1253772), Catalyst Foundation, and DARPA (HR0011-13-C-0003). We also appreciate Prof. Sung Woo Chung (Korea University) for valuable feedback and discussions.

REFERENCES

[1] K. Skadron et al., "Control-Theoretic Techniques and Thermal-RC Modeling for Accurate and Localized Dynamic Thermal Management," *International Symposium on High-Performance Computer Architecure (HPCA)*, 2002.

[2] D. Brooks et al., "Dynamic Thermal Management for High-Performance Microprocessors," *International Symposium on High-Performance Computer Architecture (HPCA)*, 2001.

[3] A. K. Coskun et al., "Dynamic Thermal Management in 3D Multipcore Architecture," *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2009.

[4] T. Yang et al., "In-situ Techniques for In-Field Sensing of NBTI Degradation in an SRAM register file," *IEEE ISSCC*, 2015.

[5] R. Mukherjee et al., "Systematic Temperature Sensor Allocation and Placement for Microprocessors," *ACM/DAC/IEEE DAC*, 2006.

[6] R. Mukherjee et al., "Thermal Sensor Allocation and Placement for Reconfigurable Systems," *International Conference on Computer Aided Design (ICCAD)*, 2006.

[7] J. Long et al., "Thermal Monitoring Mechanisms for Chip Multiprocessors," *ACM Transactions on Architecture and Code Optimization (TACO)*, 2008.

[8] S. O. Memik et al., "Optimizing Thermal Sensor Allocation for Microprocessors," *IEEE Transactions on Computer Aided Design (TCAD)*, 2008.

[9] A. N, Nowroz et al., "Thermal Monitoring of Real Processors: Techniques for Sensor Allocation and Full Characterization," *ACM/DAC/IEEE DAC*, 2010.

[10] R. Cochran et al., "Spectral Techniques for High-Resolution Thermal Characterization with Limited Sensor Data," *ACM/DAC/IEEE DAC*, 2009.

[11] S. Kim et al., "A 30.1μm2, < ±1.1°C-3σ-Error, 0.4-to-1.0V Temperature Sensor based on Direct Threshold-Voltage Sensing for On-Chip Dense Thermal Monitoring," *IEEE CICC*, 2015.

[12] E. Saneyoshi et al., "A 1.1V 3.5μm x 3.5μm thermal sensor with supply voltage sensitivity of 2°C/10%-supply for thermal management on the SX-9 supercomputer," *IEEE Symposium on VLSI Circuits*, 2008.

[13] D. Shim et al., "A Process-Variation-Tolerant On-Chip CMOS Thermometer for Auto Temperature Compensated Self-Refresh of Low-Power Mobile DRAM," *IEEE JSSC*, 2013.

[14] T. Yang et al., "0.6-to-1.0V 279 μm^2 , 0.92μW Temperature Sensor with Less Than +3.2/-3.4°C Error for On-Chip Dense Thermal Monitoring," *IEEE ISSCC*, 2014.

[15] K. Kim et al., "A 366kS/s 400μW 0.0013mm^2 Frequency-to-Digital Converter based CMOS temprature sensor using multiphase clock," *IEEE CICC*, 2009.

[16] S. Hwang et al., "A 0.008mm2 500uW 469kS/s Frequency-to-Digital Converter based CMOS Temperature Sensor with Process Variation Compensation," *IEEE Transactions on Circuits and Systems (TCAS-I)*, 2013.

[17] J. Shor et al., "Rationmetric BJT-Based Thermal Sensor in 32nm and 22nm Technologies," *IEEE ISSCC*, 2012.

[18] Y. W. Li et al., "A 1.05V 1.6mW 0.45°C 3σ-Resolution ΔΣ-Based Temperature Sensor with Parasitic-Resistance Compesation in 32nm CMOS," *IEEE ISSCC*, 2009.

[19] K. Souri et al., "1.27 A 0.85V, 600nW All-CMOS Temperature Sensor with an Inaccuracy of of ±0.4°C (3σ) from -40°C to 125°C," *IEEE ISSCC*, 2014.

[20] T. E. Carlson et al., "Sniper: Exploring the level of abstraction for scalable and accurate parallel multi-core simulations." *International Conference for High Performance Computing, Networking, Storage and Analysis*, Nov. 2011.

[21] S.C. Woo et al., "The SPLASH-2 Programs: Characterization and Methodological Considerations", *ISCA*, 1995.

[22] C. Bienia et al., "The PARSEC benchmark suite: characterization and architectural implications", International Conference on Parallel Architectures and Compilation Techniques, 2008.

[23] John L. Henning, "SPEC CPU2006 Benchmark Descriptions", *ACM SIGARCH newsletter, Computer Architecture News*, September 2006.

[24] S. Li et al.,"McPAT: An Integrated Power, Area, and Timing Modeling Framework for Multicore and Manycore Architectures," *IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Dec. 2009

[25] K. Skadron et al., "Temperature-Aware Microarchitecture", *ISCA*, 2003.

[26] S. Paek et al., "All-Digital Hybrid Temperature Sensor Network for Dense Thermal Monitoring," *IEEE ISSCC*, 2013.