Spectral Techniques for High-Resolution Thermal Characterization with Limited Sensor Data

Ryan Cochran
Division of Engineering
Brown University
Providence, RI 02912
Email: ryan_cochran@brown.edu

Sherief Reda
Division of Engineering
Brown University
Providence, RI 02912
Email: sherief_reda@brown.edu

ABSTRACT

Elevated chip temperatures are true limiters to the scalability of computing systems. Excessive runtime thermal variations compromise the performance and reliability of integrated circuits. To address these thermal issues, state-of-the-art chips have integrated thermal sensors that monitor temperatures at a few selected die locations. These temperature measurements are then used by thermal management techniques to appropriately manage chip performance. Thermal sensors and their support circuitry incur design overheads, die area, and manufacturing costs. In this paper, we propose a new direction for full thermal characterization of integrated circuits based on spectral Fourier analysis techniques. Application of these techniques to temperature sensing is based on the observation that die temperature is simply a space-varying signal, and that space-varying signals are treated identically to time-varying signals in signal analysis. We utilize Nyquist-Shannon sampling theory to devise methods that can almost fully reconstruct the thermal status of an integrated circuit during runtime using a minimal number of thermal sensors. We propose methods that can handle uniform and non-uniform thermal sensor placements. We develop an extensive experimental setup and demonstrate the effectiveness of our methods by thermally characterizing a 16-core processor. Our method produces full thermal characterization with an average absolute error of 0.6% using a limited number of sensors.

ACM Categories & Subject Descriptors
B.7.1 [Integrated Circuits]: Types and Design Styles.
General Terms: Design, Performance, Algorithms.
Keywords: Thermal management, spectral methods, spatial estimation, thermal sensors.

1. INTRODUCTION

Runtime thermal variations arise due to the spatially and temporally non-uniform power density distributions that occur during chip operation. These variations lead to thermal gradients and hot spots that degrade performance and shorten the mean time to circuit failure. State-of-the-art chips typically employ thermal sensor(s) in order to capture a number of on-chip temperature measurements during runtime. Thermal management techniques use these measurements to control cooling fan speeds, adjust workload scheduling, and assign appropriate voltages and frequencies [12, 16, 19, 4]. Thermal sensors (e.g., thermal diodes) and their support circuitry utilize silicon real estate and increase design complexity. These sensors are typically placed during design time at locations where thermal hot spots are expected to occur. However, unpredictable workloads and within-die process variations alter the locations of runtime hot spots from their expected locations.

By using Fourier analysis techniques, this paper proposes a new direction for full thermal characterization of integrated circuits using a limited number of thermal sensors. We formally ground our ideas in the classical foundations of Fourier signal analysis, which leads to superior results when compared to existing methods. Our contributions are as follows.

- We discuss the implications of using Nyquist-Shannon sampling theory to determine the fewest number of thermal sensors that can fully characterize the runtime thermal status of a processor.
- Given a limited number of thermal sensors, we propose signal reconstruction techniques that generate a full-resolution thermal characterization of a processor. We propose methods that handle uniform and non-uniform sensor placements, the latter of which may arise due to design and layout constraints.
- Using a comprehensive tool chain of thermal simulators, power estimators, and workloads, we demonstrate that our proposed methods can provide an accurate high-resolution thermal characterization for a 16-core processor. We also demonstrate the trade-off between the number of thermal sensors and the attained accuracy of the thermal characterizations.
- In addition to full thermal characterization, we quantify the effectiveness of our methods in estimating hot spot magnitudes. We demonstrate the superiority of our methods over other methods (e.g., grid-based interpolation [8] and geostatistical Kriging estimators [7]).

The organization of this paper is as follows. Section 2 provides the necessary motivation and background for this work. In Section 3, we discuss Nyquist-Shannon theory and its implications on thermal sensing. We also discuss uniform and non-uniform sampling cases. Section 4 develops a number of signal reconstruction methods. Using a 16-core processor model, we provide experimental results in Section 5 that demonstrate the effectiveness of our methods. Finally, Section 6 summarizes the main conclusions of this work and indicates directions for future work.

DAC'09, July 26-31, 2009, San Francisco, California, USA

2. BACKGROUND

High-end chips, such as multi-core processors and graphical processor units, are equipped with integrated thermal sensors that monitor the on-chip temperatures during runtime. Thermal diodes, which translate temperature variations into voltage variations, are a popular choice for temperature sensing. The diode voltage signal is routed to and measured by an analog to digital (A/D) converter. In older processor technologies, the A/D conversion takes place on the board, while in newer processors (e.g., Intel's Core i7), the A/D conversion takes place on-chip. State-of-the-art chips are typically equipped with more than one thermal sensor. There are a number of reasons for this: (1) complex chips with large die area require more thermal sensors to capture temperatures at a wide range of locations; (2) the unpredictability of a processor's workload can lead to continuous migration of hot spots; and (3) within-die manufacturing variations lead to leakage variability that can further conceal the locations of the thermal hot spots. The thermal sensors, together with their support circuitry and wiring, complicate the design process and increase the total die area and manufacturing costs. Thus, there is an inherent tradeoff in thermal monitoring. On the one hand, designers would like to reduce costs by using the fewest number of thermal sensors, while on the other hand, the gravity of runtime thermal problems require higher thermal resolution.

Given the limitations on the number of thermal sensors, it is necessary to optimally place them near potential hot spot locations. Mukherjee and Memik [10, 9] describe a clustering algorithm that computes the thermal sensor positions that best serve clusters of potential hot spot locations. The locations of these hot spots are identified via extensive workload thermal simulation. Even with optimized placement, it is likely that sensors will fail to detect hot spots, especially when there are a large number of cores. Thus, Long *et al.* [8] advocate using a grid-based interpolation scheme that identifies the hot spot around each sensor by interpolating the measurements at its immediate neighbors. Another entirely different potential method for full thermal characterization is based on geostatistical techniques. Liu [7] proposes using *Kriging* estimation as a general framework for estimating variability (whether manufacturing, thermal or IR drop) at various chip locations during design time. A Kriging temperature estimator computes the unknown temperature at a particular location using a weighted combination of the known measurements at other locations. In choosing the optimal weights, temperature is modeled as a random field. The Kriging estimator then chooses a set of weights such that the variance of the error values is minimized. In order to do this, Kriging estimators utilize *variograms*, which capture the spatial correlations between the temperatures at various locations on the die [7]. While Liu [7] does not specifically advocate the use of Kriging estimators for runtime thermal characterization, we implemented a Kriging adaptation for the sake of comparison.

Our objective is to achieve full characterization of on-chip runtime temperature using the measurements of a few thermal sensors. Our approach is formally grounded in spectral Fourier analysis techniques. Application of these techniques to temperature sensing is based on recognition that die temperature is simply a space-varying signal, and that space-varying signals are treated identically to time-varying signals in Fourier signal analysis.

While temperature is a continuous variable, any representation in computer memory must be discretized. Thus, the temperature $t(m,n)$ is a discrete function that is defined over a finite region $0 \leq m \leq M-1$ and $0 \leq n \leq N-1$, where M and N are the *resolutions* required for thermal characterization. For example, if a die has dimensions $1\ cm \times 1\ cm$, then with resolutions $M = N = 128$, temperatures are evaluated for every $78\ \mu m \times 78\ \mu m$ square. The two-dimensional Discrete Fourier Transform (DFT) is given by

$$T(p,q) = \sum_{m=0}^{M-1}\sum_{n=0}^{N-1} t(m,n)e^{-j2\pi pm/M}e^{-j2\pi qn/N} \quad (1)$$

for all $p = 0, 1, \ldots M-1$ and $q = 0, 1, \ldots N-1$, and the inverse DFT is given by

$$t(m,n) = \frac{1}{MN}\sum_{p=0}^{M-1}\sum_{q=0}^{N-1} T(p,q)e^{j2\pi pm/M}e^{j2\pi qn/N} \quad (2)$$

for all $m = 0, 1, \ldots M-1$ and $n = 0, 1, \ldots N-1$. To compute the DFT and the inverse DFT efficiently, the Fast Fourier Transform (FFT) and the inverse FFT are used, achieving an $O(MN \log MN)$ runtime as opposed to $O((MN)^2)$ [13].

There are two main aspects to full thermal characterization:

1. **Thermal Sensing or Sampling.** Thermal sensors are treated as discrete temperature samples at various die locations. With a handful of thermal sensors in even state-of-the-art chips, thermal measurements are very sparse relative to the objective of full thermal resolution. An interesting aspect of Fourier analysis is that it is capable of determining the exact number of thermal sensors required for full thermal reconstruction.
2. **Signal Reconstruction.** Given thermal sensor values, signal reconstruction techniques based on Fourier analysis can be used to reconstruct the temperature at all die locations, thus achieving full-resolution thermal characterization. Given appropriately spaced samples, thermal reconstruction can theoretically be performed with no information loss. In reality, a number of constraints hinder this theoretical possibility, and one can only hope for minimal information loss. An attractive feature of thermal characterization is that the continuous nature of heat flow gives relatively smooth on-chip temperature gradients that do not change abruptly. This feature limits the bandwidth of the thermal signal in the spectral domain, which permits near perfect signal reconstruction.

In the next two sections we will describe our methods for signal sampling and signal reconstruction. Methods for signal sampling are given in Section 3, and methods for signal reconstruction are given in Section 4.

3. THERMAL SAMPLING

With thermal sensors treated as discrete samples of a continuous phenomenon, the following important question arises: *what is the minimum number of sensors required to fully characterize the thermal status of a chip?* According to the Nyquist-Shannon sampling theorem, a continuous signal can be perfectly reconstructed from its samples if the sampling frequency is at least twice the highest frequency contained in the signal. This result is remarkable because it shows that sampling entails no loss of information for an appropriately band-limited signal. Another classical result in signal processing is that time-limited signals are not band-limited, and that band-limited signals are not time-limited; i.e., a signal cannot be simultaneously time-limited and band-limited [14]. Temperature, a space-varying signal, is space-limited by the edges of the chip. Thus, the spectral representation of the temperature is not band-limited. As a result, perfect reconstruction is not possible in the case of on-chip temperatures; however, near-perfect reconstruction with negligible loss of information is possible if these *edge effects* are appropriately handled during reconstruction such that the higher frequency magnitudes in the spectral domain are minimized.

Techniques for handling edge effects are mentioned in Section 4. We now consider two possible sampling schemes.

Uniform Sampling. We first consider the case in which the locations of the sensors are aligned on a lattice such that they are equally spaced from each other. There are a number of lattices for which such uniformity can be achieved, including hexagonal lattices, diamond lattices, and rectangular grids. We focus on rectangular grids as they are more suitable for the tiled layouts typically encountered in multi-core processors and GPUs. To sample the temperature $t(m,n)$, we assume that the thermal sensors have been placed with a horizontal/vertical spacing of $P \in \mathcal{Z}^+$. Thus, the sampling signal $s(m,n)$ can be described with

$$s(m,n) = \sum_{u=0}^{\lfloor \frac{M-1}{P} \rfloor} \sum_{v=0}^{\lfloor \frac{N-1}{P} \rfloor} \delta(m-uP, n-vP), \qquad (3)$$

where $\delta(\cdot,\cdot)$ is the Dirac delta function. To obtain the sampled temperature signal $t_s(m,n)$, we multiply the temperature with the sampling function to get

$$\begin{aligned} t_s(m,n) &= t(m,n)s(m,n) \\ &= t(m,n) \sum_{u=0}^{\lfloor \frac{M-1}{P} \rfloor} \sum_{v=0}^{\lfloor \frac{N-1}{P} \rfloor} \delta(m-uP, n-vP) \\ &= \sum_{u=0}^{\lfloor \frac{M-1}{P} \rfloor} \sum_{v=0}^{\lfloor \frac{N-1}{P} \rfloor} t(uP,vP)\delta(m-uP,n-vP) \end{aligned} \qquad (4)$$

Consider the DFT of $t_s(m,n)$. Since $t_s(m,n)$ is the product of $t(m,n)$ and $s(m,n)$, then the DFT $T_s(p,q)$ of the sampled signal $t_s(m,n)$ is equal to the convolution of the $T(p,q)$ and $S(p,q)$, which are the DFTs of $t(m,n)$ and $s(m,n)$ respectively. The Fourier transform of a periodic impulse train is a periodic impulse train as well; i.e.,

$$\sum_{u=0}^{\lfloor \frac{M-1}{P} \rfloor} \sum_{v=0}^{\lfloor \frac{N-1}{P} \rfloor} \delta(m-uP,n-vP) \leftrightarrow \frac{1}{P^2} \sum_{u=0}^{P-1} \sum_{v=0}^{P-1} \delta(p-u\frac{M}{P}, q-v\frac{N}{P}) \qquad (5)$$

Thus, the convolution of $T(p,q)$ and $S(p,q)$ gives

$$T_s(p,q) = \frac{1}{P^2} \sum_{u=0}^{P-1} \sum_{v=0}^{P-1} T(p-u\frac{M}{P}, q-v\frac{N}{P}). \qquad (6)$$

Temperature sampling by the thermal sensors in the space-domain has led to periodic copies of the Fourier transform of the temperature signal in the spectral domain. Figure 1 visually illustrates the impact of sampling, where Figure 1.a gives a thermal map of a 16-core processor in both the space and spectral-domain. After sampling, the spectral-domain representation of Figure 1.b shows the repetition of the spectral map of Figure 1.a. According to the Nyquist-Shannon theorem, if copies of the spectral-domain temperature signals are spread far "enough" apart, then overlap or *aliasing* between the copies will not occur. The necessary spreading, which is controlled by the sampling frequency, depends on the highest frequency seen in the temperature signal. If the temperature signal is band-limited with frequency B, then sampling at a rate higher than $2B$ guarantees full reconstruction of the original signal. Sampling at a higher rate is achieved by decreasing the spacing P between the thermal sensors. Thus, thermal sensors should be spaced such that $1/P \geq 2B$, or equivalently $P \leq 1/(2B)$. As mentioned earlier, the temperature signal is space-limited by the size of the chip,

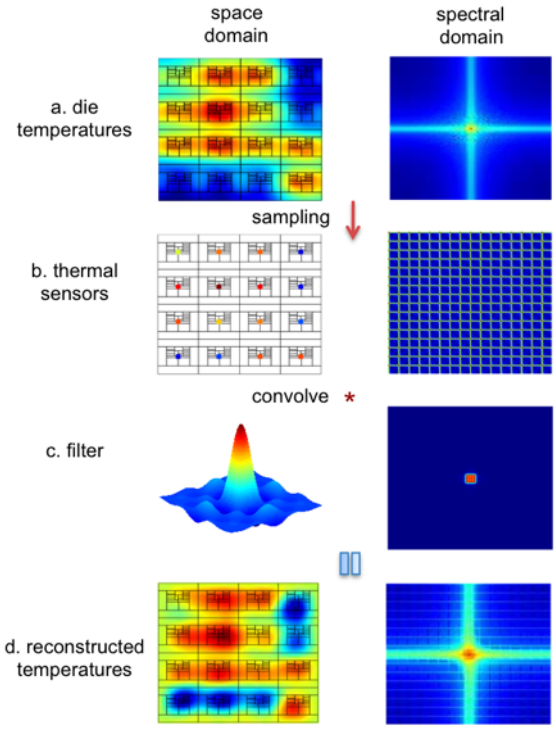

Figure 1: Main steps used for signal reconstruction. The log of the magnitude of the 2D DFT is plotted.

and consequently, the spectral-domain representation of the temperature is not band-limited. To minimize information loss due to sampling, one must pick a bandwidth B below which most of the signal's energy is concentrated. The impact of sensor spacing on the thermal characterization accuracy will be quantified in our experimental results section (Section 5).

Non-uniform Sampling. In many cases, it may be necessary to place sensors in a non-uniform pattern for the following reasons:

1. There could be regions on the die that have greater utilization and higher thermal gradients than other regions with uniform utilization. Thus, it may be advantageous to allocate more sensors to locations that have higher gradients or unpredictability.
2. Design constraints could force designers to place the sensors in a non-uniform pattern. Even in the case of multi-core processors where one can expect regular layouts, on-chip switches and cache structures can create irregular design layouts. Automatic layout and placement tools can also move blocks from their intended locations if it leads to timing or die area improvements.

Full thermal characterization consists of two problems: optimal sensor placement, and maximal use of thermal information given a set of temperature samples. The first problem is not our focus in this work. In the next section we provide solutions for the second problem for both uniform and non-uniform sample patterns.

4. THERMAL RECONSTRUCTION

We have seen in the previous section that sampling a signal in the space-domain leads to periodic copies of the Fourier transform, $T(p,q)$, of the original signal $t(m,n)$, in the spectral-domain. If the temperature signal is band-limited (or has negligible energy beyond a certain frequency) then these periodic copies are well separated from each other and aliasing is minimal. The Whittaker-Shannon-Kotelnikov (WSK) classical theorem states that to recover

the original signal from the samples, it is sufficient to extract only one copy of the signal in the spectral-domain [14]. This extraction can be achieved using a low-pass box filter as shown in Figure 1.c. In the frequency domain, this box can be expressed as

$$\begin{aligned} F(p,q) &= 1 \text{ if } |p| \leq B \text{ and } |q| \leq B \\ &= 0 \text{ otherwise.} \end{aligned} \quad (7)$$

Taking the inverse DFT of the box filter gives the spatial-domain representation of the filter $f(m,n)$ which is equal to

$$f(m,n) = \text{sinc}(\frac{m}{B})\text{sinc}(\frac{n}{B}). \quad (8)$$

Reconstruction is achieved by *convolving* the space-domain samples with the space-domain filter representation. That is, the reconstructed temperature of a chip $t_r(m,n)$ can be found using

$$t_r(m,n) = \sum_{u=0}^{M-1}\sum_{v=0}^{N-1} t_s(u,v)\text{sinc}(\frac{m}{B}-u)\text{sinc}(\frac{n}{B}-v). \quad (9)$$

This result is illustrated in Figure 1.d. One of the practical problems that arises when using the sinc function is that it is not space-limited. In any implementation, the sinc function must be truncated, which has the effect of smearing its spectral-domain representation, leading to a less than sharp box filter edge. This smearing effect can be minimized by *windowing* the sinc function. In our implementations, we multiply the sinc function by a *Hamming window* of the same size. The severity of the edge effects depends on the size of the sinc function used. Edge effects can be minimized by extending the temperature data by a distance larger than half the size of the filter function. The values in the extended region can be copies of the edge values, periodic repetitions of the temperature signal, or even a mirror image of the temperature signal. It is also useful to investigate other filter functions that approximate a low-pass box filter in the spectral domain without windowing [11]. We investigate three other functions.

- **Nearest neighbor:** The simplest interpolation function is nearest neighbor, in which each location is given a temperature equal to the value measured by the sensor closest to it. This is achieved by convolving the sampled temperature signal with a rectangular function expressed as follows:

$$\begin{aligned} f(x,y) &= 1, \text{ for } x \in [-0.5, 0.5] \text{ and } y \in [-0.5, 0.5] \\ f(x,y) &= 0, \text{ otherwise} \end{aligned}$$

- **Linear function:** Linear interpolation amounts to convolving the temperature samples with a round cone function. This function corresponds to a modestly good low-pass filter in the spectral-domain. However, it attenuates frequencies near the cut-off frequency, resulting in smoothing of the thermal characterization results. It also passes a good amount of energy above the cut-off frequency. The linear function is expressed as follows:

$$\begin{aligned} f(x,y) &= g(x)g(y), \text{ where} \\ g(u) &= (1-u) \text{ for } u \in [0,1] \end{aligned}$$

- **Cubic B-spline function:** Cubic B-spline functions are reasonably good low-pass filters. They are positive in the whole interval from 0 to 2, so they smooth somewhat more than necessary below the cut-off frequency. These filters are symmetric, so they only need to be expressed on the interval $[0,2]$. The cubic B-spline function we use is expressed as follows:

$$\begin{aligned} f(x,y) &= g(x)g(y), \text{ where} \\ g(u) &= \frac{u^3}{2} - u^2 + \frac{4}{6} \text{ for } u \in [0,1] \\ g(u) &= \frac{-u^3}{6} + u^2 - 2u + \frac{8}{6} \text{ for } u \in [1,2] \end{aligned}$$

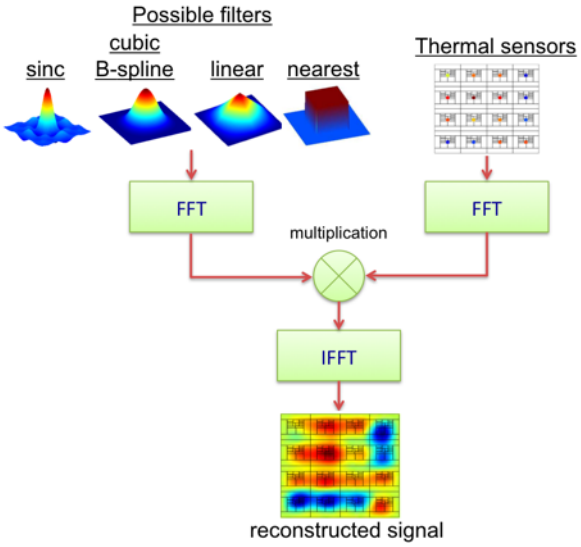

Figure 2: Flow of the proposed runtime thermal characterization technique.

All three interpolation functions are contrasted with the sinc function in Figure 2. To achieve a computationally efficient thermal characterization, we propose the flow in Figure 2. Instead of directly convolving in the space-domain, we first take the Fast Fourier Transform (FFT) of the temperature samples and the space-domain representation of the interpolation functions. We then multiply the resultant spectral-domain representations to get the FFT of the reconstructed 2D thermal signal. We then apply the inverse FFT (IFFT) to get the full-resolution thermal characterization in the space-domain. If the required thermal resolution is $R = MN$, then application of the proposed flow of Figure 2 has a runtime of $O(R \log R)$, which gives a significant advantage over the $O(R^2)$ runtime achieved by straightforward convolution. Such speedup boost is necessary for runtime thermal characterization.

Reconstruction from Non-Uniform Samples. If thermal sensors are placed non-uniformly, then directly applying the proposed flow of Figure 2 could lead to large errors. Instead, it is necessary to apply signal reconstruction algorithms that are devised to handle non-uniform samples. One such algorithm is proposed by Sauer and Allebach [15]. This iterative algorithm consists of two steps:

1. Given the locations of the thermal sensors, construct a Voronoi diagram for the die. All die locations that belong to the Voronoi cell of a thermal sensor are assigned the same temperature as the thermal sensor at center of the cell.
2. The resultant 2D temperature map obtained from the first step is low-pass filtered B (by convolving the temperature with the sinc function) using the flow of Figure 2.

Sauer and Allebach [15] prove that iterating these two steps leads to convergence and reconstruction of the original signal if it is band-limited.

5. EXPERIMENTAL RESULTS

Simulation Infrastructure. To evaluate the effectiveness of our methodology, we set up a *tool chain* that simulates temperatures for a 32 nm 16-core processor. Our tool chain takes as inputs the processor's floor-plan and the workload that will run on each core and produces as output the steady-state temperatures at various

grid locations. Using workload instruction traces, dynamic power traces for each micro-architectural unit are calculated and then fed together with the floor-plan into the thermal simulator. Once the steady-state temperatures are calculated, they are fed into a leakage calculator which outputs the corresponding leakage power of each unit. The leakage power values are then added to the original dynamic power values, and a new thermal simulation is performed. This process is iterated until the temperatures converge to stable values. Power and thermal simulations are performed using the following tools.

- For dynamic power estimation, we use a Wattch-like [1] power simulator, with the power consumption appropriately scaled to 32 nm technology based on ITRS predictions [3]. For the leakage consumption of the processor core units, we construct a leakage model using the expressions for leakage power from PTScalar [6]. To accurately model cache leakage power, we use CACTI 5.0 [18], which has accurate cache leakage values at current and future technology nodes.
- We utilize HotSpot (version 4.0) [17] for thermal simulation. HotSpot takes as inputs the processor floor-plan and workload power traces and produces as output the steady-state temperatures for a set of grid locations.
- We use the Alpha 21264 processor as our baseline core [5]. The 21264 is an out-of-order speculative execution core that is commonly used as a test-bench core in thermal management research [6, 8]. We create a 16-core processor based on the Alpha processor. The die area of the processor is $1.1\ cm \times 1.1\ cm$, and we discretize the temperature by defining a grid with resolution $M = 64 \times N = 64$, where each grid location represents the temperature for an area of size $172\ \mu m \times 172\ \mu m$.
- For workloads, we use eight benchmarks from the SPEC2000 suite [2]. We use four integer benchmarks: gcc, bzip, mcf, and twolf, and four floating point benchmarks: ammp, equake, lucas, and mesa.

In each simulation, we assign each core in the 16-core processor a workload from our SPEC2000 benchmark selection such that each core gets a random workload from the available eight. We then assign each core a random frequency in the range $1.5 - 3$ GHz. We then execute the tool chain to find the true temperatures at all grid locations. Thermal sampling is accomplished by zeroing temperatures at all grid locations except those corresponding to thermal sensors. The number of samples and the sample locations are varied, and each proposed method is evaluated for each sensor configuration using the *average* and *maximum* absolute post-reconstruction error across a set of 32 simulations.

Experiments. In our first set of experiments, we determine the average and maximum absolute error for each reconstruction method for uniformly spaced samples while varying the total number of sensors on a 16-core processor. We present results for the following cases: 1 sensor per core, 4 sensors per core arranged on a 2×2 grid, and 9 sensors per core arranged on a 3×3 grid. The first case corresponds to 16 total sensors, while the second and third cases correspond to 64 and 144 total sensors respectively. The barplot of Figure 3 summarizes the average absolute error calculated for each proposed reconstruction method. We include error values for a geostatistical-based Kriging estimator for sake of comparison. The results show that as the number of sensors increases, the thermal characterization error decreases. Furthermore, the results show that our proposed reconstruction methods are capable of achieving full thermal characterization with minimum average absolute error

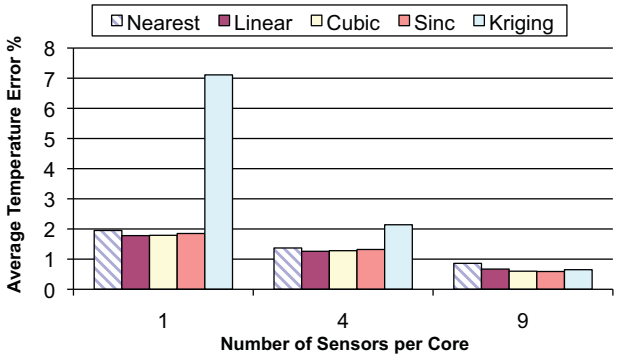

Figure 3: Results of full thermal characterization. We report the average error percentage for temperature estimation at all grid locations.

ranging from 1.8% to 0.6%, depending on the number of the sensors. In this experimental setting, our convolution filters deliver near identical performance, and they significantly outperform the Kriging estimator, especially when the number of sensors are few. Our second set of experiments will reveal a differentiation in the performance of our proposed convolution filters. Using a single representative workload and frequency assignment, Figure 4 compares the thermal resolution attained by using 1, 4, and 9 sensors per core to the true thermal characterization. We include the latter case, despite a possibly unrealistic number of total sensors for a 16-core processor, to illustrate that increasing the number of sensors eventually reconstructs the original signal to near perfection.

Full thermal characterization is particularly useful for advanced thermal management techniques, examples of which include workload scheduling and per-core frequency and voltage assignment. In simpler thermal management techniques (e.g., fan speed control), only the magnitude and the location of the maximum hot spot are relevant. Thus, in our second set of experiments, we consider the hypothetical performance of our methods in such applications. We first identify the location and magnitude of the maximum hot spot in the true thermal characterization. We then evaluate the temperature at that location in our reconstructed results and report the average absolute error. We report error values for each method explored in our first set of experiments, and in addition we implement and compare to the neighborhood interpolation scheme in [8]. Our results in Figure 5 show that the sinc filter function is consistently better at interpolating the maximum hot spot (0.4% errors for the case of 9 sensors per core), and that all of our reconstruction tech-

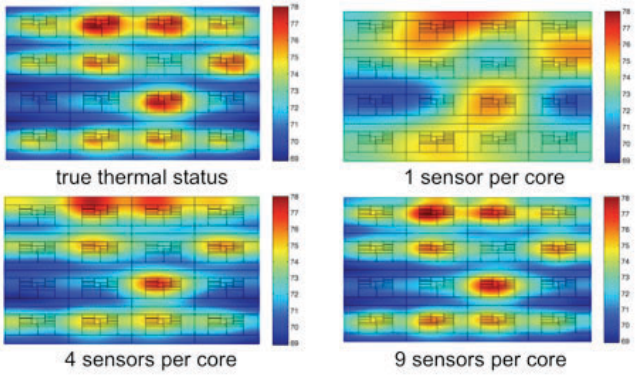

Figure 4: Impact of increasing the number of thermal sensors on the full thermal characterization.

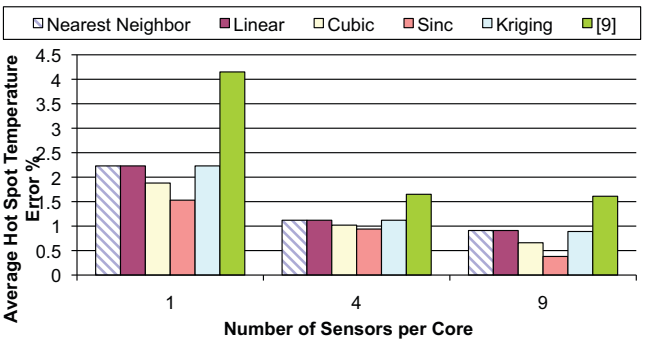

Figure 5: Hot spot estimation. We report the average error percentage for temperature estimation at the hottest die location.

niques outperform the Kriging estimator and the method proposed in [8].

In our third set of experiments, we consider non-uniform sensor placements. Figure 6 shows a representative workload and marks the locations of the thermal sensors in the processor's floor-plan with '*'. We use the same workload and frequency assignment as in Figure 4. Figure 6.a gives the output of the first step of the iterative algorithm which constructs the Voronoi diagram, and 6.b gives the results after the algorithm converges using a sinc filter function. The average absolute error for full thermal characterization using non-uniform samples is found to be 2.39%. This value confirms that our methods are capable of handling signal reconstruction for both uniform and non-uniform sensor placements.

6. CONCLUSIONS

In this paper we have presented a new direction for runtime thermal characterization in integrated circuits based on Fourier analysis techniques. The application of Fourier analysis to thermal characterization creates a theoretical foundation for this field of research. We have analyzed the implications of the Nyquist-Shannon sampling theory in determining the optimal number of thermal sensors. We have also analyzed the trade-off between thermal estimation error and the number of thermal sensors. We have proposed the use of signal reconstruction methods for full thermal characterization with limited sensor measurements. We have analyzed the case in which the thermal sensors are non-uniformly spaced, and we have proposed and implemented an algorithm that is capable of handling such a case. The effectiveness of our techniques have been evaluated on a 16-core processor model, and we have demonstrated the superiority of these techniques to existing techniques.

Our proposed work is the first to consider full thermal characterization with limited sensor data. The success of our proposed methods will allow thermal management techniques to take advantage of a full-resolution thermal map for a processor. Our future work will focus on building a software prototype that utilizes actual thermal sensor measurements obtained from Intel Core i7 to achieve full thermal characterization. We will also investigate the impact of calibration errors in the thermal sensor measurements on the results of our proposed methods.

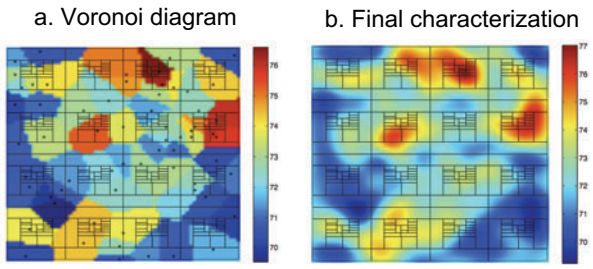

Figure 6: Results of thermal characterization using non-uniform sampling.

7. REFERENCES

[1] D. Brooks, V. Tiwari, and M. Martonosi, "Wattch: A Framework for Architectural-Level Power Analysis and Optimizations," in *International Symposium on Computer Architecture*, 2000, pp. 83–94.

[2] J. Henning, "SPEC CPU2000: Measuring CPU Performance in the New Millennium," *IEEE Computer*, vol. 33(7), pp. 28–35, 2000.

[3] ITRS, "International Technology Roadmap for Semiconductors," http://public.itrs.net, 2007.

[4] M. Kadin and S. Reda, "Frequency and Voltage Planning for Multi-Core Processors Under Thermal Constraints," in *International Conference on Computer Design*, 2008, pp. 463–470.

[5] R. Kessler, E. McLellan, and D. Webb, "The Alpha 21264 Microprocessor Architecture," in *Proc. International Conference on Computer Aided Design*, 1998, pp. 90–95.

[6] W. Liao, L. He, and K. Lepak, "Temperature and Supply Voltage Aware Performance and Power Modeling at Microarchitecture Level," *Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 24(7), pp. 1042–1053, 2005.

[7] F. Liu, "A General Framework for Spatial Correlation Modeling in VLSI Design," in *Proc. Design Automation Conference*, 2007, pp. 817–822.

[8] J. Long, S. Memik, G. Memik, and R. Mukherjee, "Thermal Monitoring Mechanisms for Chip Multiprocessors," in *ACM Transactions on Architecture and Code Optimization*, vol. 5(2), 2008, pp. 9:1–9:23.

[9] S. O. Memik, R. Mukherjee, M. Ni, and J. Long, "Optimizing Thermal Sensor Allocation for Microprocessors," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 27(3), pp. 516–527, 2008.

[10] R. Mukherjee and S. Memik, "Systematic Temperature Sensor Allocation and Placement for Microprocessors," *Design Automation Conference*, pp. 542 – 547, 2006.

[11] J. Parker, R. Kenyon, and D. Troxel, "Comparison of Interpolating Methods for Image Resampling," *IEEE Transactions on Medical Imaging*, vol. 2(1), pp. 31–39, 1983.

[12] M. D. Powell, M. Gomaa, and T. N. Vijaykumar, "Heat-and-Run: Leveraging SMT and CMP to Manage Power Density Trhough the Operating System," in *International Conference on Architectural Support for Programming Languages and Operating Systems*, 2004, pp. 260–270.

[13] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C*, second, Ed. Cambridge University Press, 1996.

[14] M. J. Roberts, *Signals and Systems*, First, Ed. McGraw Hill, 2004.

[15] K. D. Sauer and J. P. Allebach, "Iterative Reconstruction of Band-Limited Images from Nonuniformly Spaced Samples," *IEEE Trans. Circuits and Systems*, vol. 34, no. 12, pp. 1497–1506, 1987.

[16] K. Skadron, "Hybrid Architectural Dynamic Thermal Management," in *Design, Automation and Test in Europe*, 2004, pp. 10–15.

[17] K. Skadron, S. Ghosh, S. Velusamy, K. Sankaranarayanan, and M. Stan, "HotSpot: A Compact Thermal Modeling Methodology for Early-Stage VLSI Design," *Transactions on VLSI Systems*, vol. 15(5), pp. 501–513, 2006.

[18] S. Wilton and N. P. Jouppi, "CACTI: An Enhanced Cache Access and Cycle Time Model," *IEEE Journal Solid-State Circuits*, vol. 31(5), pp. 677–688, 1996.

[19] S. Zhang and K. S. Chatha, "Approximation Algorithm for the Temperature-Aware Scheduling Problem," in *Proc. International Conference on Computer Aided Design*, 2007, pp. 281–288.