Post-Silicon Power Characterization Using Thermal Infrared Emissions

Ryan Cochran Abdullah Nazma Nowroz Sherief Reda
School of Engineering
Brown University
Providence, RI
Emails: {ryan_cochran, abdullah_nowroz, sherief_reda}@brown.edu

ABSTRACT

Design-time power analysis is one of the most critical tasks conducted by chip architects and circuit designers. While computer-aided power analysis tools can provide power consumption estimates for various circuit blocks, these estimates can substantially deviate from the actual power consumption of working silicon chips. We propose a novel methodology that provides accurate, detailed post-silicon spatial power estimates using the thermal infrared emissions from the backside of silicon die. We theoretically and empirically demonstrate the inherent difficulties in thermal to power inversion. These difficulties arise from measurement errors and from the inherent spatial low-pass filtering associated with heat diffusion. To address these difficulties we propose new techniques from regularization theory to invert temperature to power. Furthermore, we propose new techniques to compute the emissivities and conductances required for any infrared to power inversion method. To verify our results, a programmable circuit of micro heaters is implemented to create any desired power pattern. The thermal emissions of different known injected power patterns are captured using a state-of-the-art infrared camera, and then our characterization techniques are applied to invert the thermal emissions to power. The estimated power patterns are validated against the injected power patterns to demonstrate the accuracy of our methodology.

ACM Categories & Subject Descriptors
B.7.1 [Integrated Circuits]: Types and Design Styles.
General Terms: Design, Performance, Algorithms.
Keywords: Power characterization, thermal infrared emissions.

1. INTRODUCTION

It is possible to distinguish modern chips into two main categories: *power limited* and *hot spot limited* [5]. Power-limited chips are used in battery-operated mobile devices where it is necessary to minimize the total power consumption to prolong the device's lifetime between battery recharges. The wide range of applications (e.g., 3D games, web browsing, video decoding) that are possible to execute on mobile devices complicate the problem of power estimation. On the other hand, designers of hot-spot limited chips seek to minimize the peak junction temperatures which arise from the spatial and temporal distribution of power densities. Excessive chip temperatures impact reliability and constrain performance.

Design of power-limited and hot-spot limited chips requires the computation of accurate power estimates for the different blocks of a circuit. Design-time estimates are typically inaccurate in comparison to the true post-silicon power consumption for the following reasons: modeling and computational inaccuracies during design time, impact of runtime workloads, and manufacturing process variations [14, 7, 10, 11]. Our objective is to provide accurate post-silicon spatial power estimates for the various blocks of a circuit using the runtime thermal infrared emissions from the back side of the chip. The major contributions of this paper are as follows.

- We elucidate the technical challenges in thermal to power inversion. Spatial discretization and measurement errors in thermal imaging limit the ability to resolve power spatially. In addition, heat conduction has a low-pass filtering effect that attenuates the thermal impact of high-frequency spatial power variations. We investigate these phenomena theoretically and empirically and demonstrate their impact on power estimation accuracy.
- To address the outlined technical challenges we propose techniques from regularization theory. We propose quadratic formulations that are augmented with Tikhonov regularization methods to improve the accuracy of the spatial power estimation. Our techniques limit the propagation of measurement errors and lead to unique solutions.
- A major challenge in thermal to power inversion is to be able to validate the power estimations. Accordingly we design and implement a circuit with programmable micro heaters that can generate any desired power pattern. We exercise the circuit with a number of different spatial power patterns and then capture the thermal emissions in the 2.5 μm – 5 μm infrared range from the back-side of the chip using an infrared camera. We are able resolve down to 187 mW in a block area of size 1.24×10^{-6} m^2. Our work is the first to validate its post-silicon power characterization estimates.
- To address the difficulties encountered in our realistic experimental setup, we provide novel techniques (1) to compensate for the thermal emissivities of different parts of the chip, and (2) to learn the heat conductance values of the thermal-power model.

The organization of this paper is as follows. Section 2 overviews some of the recent relevant work in post-silicon power estimation. In Section 3 we describe our new methodology for thermal to power inversion. In Section 4 we put together an elaborate experimental setup to validate our techniques and test their limitations. Finally, Section 5 summarizes the main conclusions of this work.

ISLPED'10, August 18–20, 2010, Austin, Texas, USA.

2. MOTIVATION AND PREVIOUS WORK

Design time power analysis is a critical task conducted by chip architects and designers. In many cases this analysis can lead to estimates that substantially deviate from the power consumption of working silicon for the following reasons.

- It is computationally infeasible to carry-out an accurate transistor-level power simulation for the entire execution time of a workload. High-level power estimation tools address this problem at the expense of power estimation accuracy [3, 11]. These tools rely on parameterizable power models that are indexed by circuit blocks and their access counts. The inaccuracy of high-level power estimates could be up to 30% [3].
- The dependency of leakage current on temperature complicate design-time power modeling. Self-consistent power analysis approaches are computationally demanding and can only provide accurate estimations for small-size circuits [10, 8]. Furthermore, process variabilities create within-die leakage variations that lead to further inaccuracies in the power models [14].

The aforementioned reasons imply that it is necessary to validate and calibrate the accuracy of design-time power models and estimators. Hence, it is necessary to conduct post-silicon power characterization. Measuring the total power consumption does not yield detailed spatial power maps that give the power consumption of different circuit blocks. To conduct fine-grained post-silicon power analysis, there have been two main research directions:

1. *Power Characterization Using Infrared Emissions*: The first direction uses thermal infrared emissions to estimate the power consumption of different circuit blocks [5, 9]. In one approach by an industrial team, the spatial steady-state power consumption is estimated from temperature using least-square estimation [5]. A second approach by an academic team used a combinatorial optimization formulation based on genetic algorithms to find a power solution that leads to the measured emissions [9]. Both approaches did not the address the inherent limitedness of resolving spatial power from temperature measurements, and more importantly, they did not provide any validations on the accuracy of their spatial power estimates.
2. *Dynamic Power Characterization Using Performance Counters*: The second direction uses the embedded frequency counters in most modern processors to estimate the switching activity of various circuit blocks [7, 11]. These switching activities are used with design-time estimates of the capacitances of various circuit blocks to compute the blocks' dynamic power consumption. These methods need their own calibration and do not take into account leakage and its variability.

Comparing the two directions, we find that the second direction relies on design time estimates which are not necessarily accurate and requires performance-counter infrastructure that is specific to processors. Furthermore, the second approach cannot estimate leakage power and its spatial variability. The first direction which we follow in this paper is more generic and powerful.

3. PROPOSED METHODOLOGY

The physical relationship between power and temperature is governed by the heat diffusion equation. In steady state, the heat equation is given by

$$\bigtriangledown \cdot (k(x,y) \bigtriangledown t(x,y)) - ht(x,y) = -p(x,y), \qquad (1)$$

where $k(x,y)$ is the thermal conductivity at location (x,y), $t(x,y)$ is temperature at location (x,y) relative to the ambient temperature, $p(x,y)$ is the power density at location (x,y), and h is the heat transfer coefficient. The explicit use of location in the thermal conductance $k(x,y)$ is important as different parts of the chip have different material composition which impact lateral heat diffusion.

In any practical implementation, the heat equation has to be discretized. This discretization comes from the finite memory size of any computer and more importantly from the use of thermal imaging equipment with limited spatial resolution. In a discretized form, the continuous temperature signal $t(x,y)$ is represented by a vector $\mathbf{t}$ that gives the measured temperatures at every pixel of the imaging system, and the continuous power signal $p(x,y)$ is represented by a vector $\mathbf{p}$ that gives the power density at a set of discrete die locations. In such case, Equation (1) is approximated by the following linear matrix formulation

$$\mathbf{Rp} + \mathbf{e} = \mathbf{t}, \qquad (2)$$

where the matrix $\mathbf{R}$ is determined by the thermal resistivities between different locations, and the vector $\mathbf{e}$ denotes the errors in measurements of the infrared imaging system. Previous approaches used least-square estimation to find $\mathbf{p}$ that gives temperatures as closest as possible (in L_2 norm) to $\mathbf{t}$ [5]. We argue that previous approaches ignored the ill-posed and ill-conditioned nature of thermal to power inversion where high-frequency spatial thermal gradients and measurements error could lead to significant power estimation errors. We describe our arguments in the next subsection.

3.1 Limits of Thermal to Power Inversion

We describe two limitations that impact the accuracy of inverting measured temperatures to spatial power estimates. The first limit arises from the physics of heat diffusion and the second limit arises from the technology used to perform thermal imaging.

The first limit in thermal to power inversion is inherent in the physics of heat transfer. It is well-established in the literature that a chip's temperature map is a low-pass filtered form of its power dissipation map [4, 6]. While the power dissipation can spatially vary abruptly according to the chip's layout and application behavior, the temperature will always vary smoothly in space. This low-pass filtering effect is inherent in the physical behavior of heat conduction as governed by Equation (1). While we do explicitly model and measure the conductances $k(x,y)$ of different parts of a chip, in this section we will consider a die with homogenous material, i.e., $k(x,y) = k$, to simplify the discussions and to focus on the key concepts. In this case, Equation (1) simplifies to

$$k(\frac{\partial^2 t(x,y)}{\partial x^2} + \frac{\partial^2 t(x,y)}{\partial y^2}) - ht(x,y) = -p(x,y). \qquad (3)$$

The 2D Fourier transform of some function $t(x,y)$ is defined by

$$F_t(u,v) = \int_0^{-\infty} \int_0^{-\infty} t(x,y) e^{-j2\pi(ux+vy)} dxdy, \qquad (4)$$

where u and v are variables that represent the spatial frequencies in the x and y directions respectively. Applying the 2-D Fourier transform to both sides of Equation (1) gives

$$k(-u^2 F_t(u,v) - v^2 F_t(u,v)) - hF_t(u,v) = -F_p(u,v). \qquad (5)$$

Thus, the power-temperature transfer function $G(u,v) = \frac{F_t(x,y)}{F_p(x,y)}$ is equal to

$$G(u,v) = \frac{1}{k(u^2+v^2)+h} \qquad (6)$$

It is clear from Equation (6) that the higher frequency components of spatial power patterns will be subjected to greater attenuation. The equation also shows that k, which determines the degree of lateral conductance, has an impact on the filtering effect. If there is no lateral conductance, i.e. $k = 0$, then the frequency content of the power signal is preserved at all spatial frequencies. As the lateral conductance increases, attenuation at higher frequencies further increases. This inherent low-pass filtering effect has profound impact on the problem of inverting temperature to power. If high-frequency components of power patterns are sufficiently attenuated, no spatial temperature variations will result leading to an irreversible loss of information. High-frequency power components can easily occur when neighboring chip blocks are independently triggered at different times. This low-pass filtering effect makes the temperature to power inversion problem *ill-posed* as many power patterns could lead to the same thermal image.

The second limit arises from discretization and measurement errors introduced during infrared imaging. Consider a small variation $\delta\mathbf{t}$ in the discrete temperature image as given by Equation (2). Then by linearity the change in the power estimation is equal to $\delta\mathbf{p} = \mathbf{R}^{-1}\delta\mathbf{t}$. It can be shown [2] that

$$||\delta\mathbf{p}|| \leq \frac{1}{\sigma_{\min}}||\delta\mathbf{t}||, \tag{7}$$

where $\sigma_{\min}$ is the minimum singular value of $\mathbf{R}$. On the other hand,

$$||\mathbf{t}|| \leq \sigma_{\max}||\mathbf{p}||, \tag{8}$$

where $\sigma_{\max}$ is the maximum singular value of $\mathbf{R}$. By combining Equations (7) and (8), we get that

$$\frac{||\delta\mathbf{p}||}{||\mathbf{p}||} \leq \frac{\sigma_{\max}}{\sigma_{\min}}\frac{||\delta\mathbf{t}||}{||\mathbf{t}||}. \tag{9}$$

The ratio $\sigma_{\max}/\sigma_{\min}$ is the *condition number* of $\mathbf{R}$ and it controls error propagation from the measurements to the solution. Large condition numbers make the inversion problem *ill-conditioned* further complicating temperature to power inversion. For example, the condition number of our chip's $\mathbf{R}$ matrix is equal to 4.035×10^3 which could transform a milli-Kelvin error – a typical number in cryogenic thermal sensors – in temperature measurement into a significant error in power estimation.

3.2 Methods for Thermal to Power Inversion

Given the thermal measurements $\mathbf{t}$, the objective is to find the best power pattern $\mathbf{p}$ that minimizes the total squared error between the temperatures as computed using Equation (2) and the measured temperatures; that is,

$$\min \parallel \mathbf{Rp} - \mathbf{t} \parallel^2,$$

under the constraints that

$$\begin{aligned} [1 \cdots 1]\mathbf{p} &= p_{total} \text{ and} \\ \mathbf{p} &\geq \mathbf{0}, \end{aligned}$$

where p_{total} is the total power consumption of the chip which could be measured externally through an ammeter. As discussed earlier in Subsection 3.1, the ill-posed nature of the continuos problem and the ill-conditioned nature of the discrete problem could lead to power estimates with significant errors. To address these issues, we propose techniques from *regularization theory* [2]. Regularization techniques consider a family of approximate solutions depending on a positive parameter called the *regularization parameter*. When the measured thermal data is noise-free, the solution converges to the true power solution as the regularization parameter goes to zero. When the data is noisy, we can get an optimal approximation of the true solution with a non-zero value of the regularization parameter.

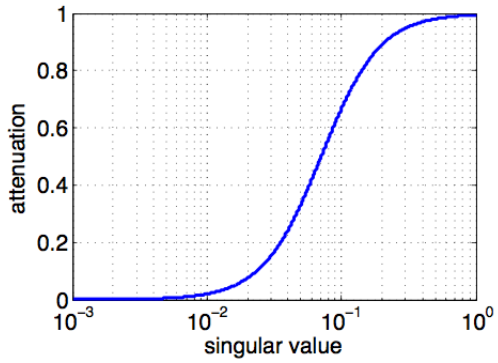

Figure 1: Attenuation of the singular values as a function of σ. The regularization parameter $\alpha = 0.05$.

Tikhonov regularization finds the power solution that gives the least total squared error while penalizing solutions with large L_2 norm; that is,

$$\begin{aligned} \mathbf{p}_\alpha &= \arg\min_{\mathbf{p}} \parallel \mathbf{Rp} - \mathbf{t} \parallel^2 + \alpha \parallel \mathbf{p} \parallel^2 \\ &= (\mathbf{R}'\mathbf{R} + \alpha\mathbf{I})^{-1}\mathbf{R}'\mathbf{t}, \end{aligned} \tag{10}$$

where $\alpha > 0$ is the regularization parameter and $\mathbf{R}'$ is the transpose of matrix $\mathbf{R}$. Penalizing the power estimates with large L_2 norm guarantees the uniqueness of the solution. Using the Singular Value Decomposition (SVD) $\mathbf{R} = \mathbf{USV}'$, where $\mathbf{U}$ and $\mathbf{V}$ are unitary matrices and $\mathbf{S}$ is a diagonal matrix with the diagonal elements are the non-zero singular values. Using the SVD, Equation (10) can be further analyzed as follows

$$\begin{aligned} \mathbf{p}_\alpha &= (\mathbf{VS}'\mathbf{U}'\mathbf{USV}' + \alpha\mathbf{VIV}')^{-1}\mathbf{VS}'\mathbf{U}'\mathbf{t} \\ &= \mathbf{V}(\mathbf{S}'\mathbf{S} + \alpha\mathbf{I})^{-1}\mathbf{S}'\mathbf{U}'\mathbf{t} \\ &= \mathbf{V}\text{diag}(\frac{\sigma_i^2}{\sigma_i^2 + \alpha} \times \frac{1}{\sigma_i})\mathbf{U}'\mathbf{t} \end{aligned} \tag{11}$$

Thus the impact of Tikhonov regularization parameter α is that it filters out singular components that are small relative to α [15]. Figure 1 illustrates this effect by giving the value of the *attenuation factor* $\frac{\sigma_i^2}{\sigma_i^2+\alpha}$ as a function of the singular value σ_i for $\alpha = 0.05$. The removal of these singular components makes the problem more well conditioned and reduces the error propagation as given earlier by Equation (9).

3.3 Emissivity and Temperature Computation

Measuring the temperature is slightly complicated by that fact that an infrared camera is really a photon detector that measures the infrared radiation intensity at different parts of the chip. Thus, it is necessary to convert photon counts recorded by the camera to temperatures. One problem is that radiation intensity is not constant among different materials even if they are at the same temperature. Perfect radiation emitters are *black bodies* with an *emissivity* of 1. The emissions of real materials are a fraction of the black-body level, and each material is characterized with an emissivity value (ε), which is defined as the ratio of that material's thermal emission to that of a perfect black-body at the same temperature [13]. As integrated circuits are comprised of a mixture of different materials (e.g., copper, silicon and dielectrics) with varying spatial material densities, the radiation intensities of different parts of the

chip could be different even if the chip is isothermally kept at homogenous temperature by external means. In order to recover the true thermal status of a chip from infrared emissions, the emissivities of different parts of the chip have to be measured.

In addition to varying emissivity values, different materials reflect radiation from the surrounding environment with varying intensity. These reflections obscure the true thermal status of the material under observation and are particularly problematic for the highly reflective metals in integrated circuits. The emissivity and reflection at each chip location can be learned by taking many calibration thermal images of a chip after forcing the chip to a known isothermal status through external means. Let $N_m(x, y, t)$ denote the photon count measured over some integration period at pixel location (x, y) with temperature t and let the black-body photon count at the same temperature be denoted by $N_b(t)$. Then from basic rules of physics [1], $N_m(x, y, t)$ and $N_b(t)$ are related by

$$N_m(x, y, t) = \varepsilon(x, y)N_b(t) + (1 - \varepsilon(x, y))N_a(x, y), \quad (12)$$

where $N_a(x, y)$ is photon reflection count from the ambient environment at location (x, y). If the entire chip is kept at an isothermal status – isothermal status can be achieved by heating or cooling the chip externally while the chip is turned off – then $N_b(t)$ can be determined by measuring the radiations levels for a black body (or near-perfect black body) at the same isothermal status. With $N_m(x, y, t)$ and $N_b(t)$ measured at a number of isothermal temperature settings, the remaining two unknowns $\varepsilon(x, y)$ and $N_a(x, y)$ in Equation (12) can be computed for each chip's location using least-square regression.

With the emissivity values for the different parts of the chip in hand, it is readily possible to estimate the temperature from the radiation levels. The relationship between black body radiation and temperature is given by

$$N_b(t) = \alpha e^{\beta t}, \quad (13)$$

where α and β are constants that depend on the camera's detectors and the specified integration time. These constants are typically provided as part of the imaging system calibration process. Combining Equation (12) and Equation (13), it is possible to compute the temperatures for each location of the chip by using the following equation

$$t(x, y) = \frac{1}{\beta} \ln \frac{N_m(x, y, t) - (1 - \varepsilon(x, y))N_a(x, y)}{\alpha\varepsilon(x, y)}. \quad (14)$$

Finally, we filter the temperatures computed from Equation (14) by using a *median filter* to reduce "salt and pepper" noise which arises from the thermal noise fluctuations associated with the quantum detectors of the infrared camera. The resultant temperature maps are used as an input to our thermal to power inversion method which was discussed Subsection 3.2.

4. EXPERIMENTAL SETUP AND RESULTS

One of our main research goals is to assess the accuracy of thermal to power inversion. Previous studies [5, 9] presented spatial power estimates for the different processor blocks "as is".[1] None of the previous studies provided a thorough validation for the accuracy of their spatial power estimations. This lack of validation is a consequence of the experimental setup that previous studies chose which offered no alternative methods to accurately estimate

[1][5] used a dual core PowerPC 970MP and [9] used an AMD Athlon 64 processor.

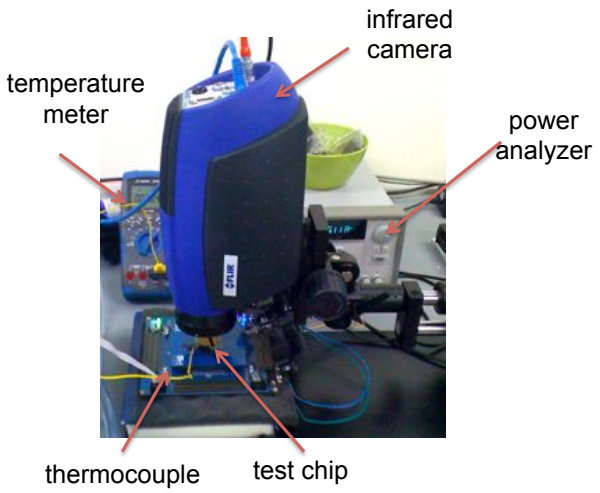

Figure 2: Experimental setup illustrating the different equipment that make up our measurement system.

the spatial power consumption. Thus, we designed our experimental platform with special emphasis on the ability to estimate the spatial power consumption through two different means, and hence we are able to provide estimations for the accuracy of our proposed power characterization methodology.

To test our post-silicon power characterization methodology, we put together the following sophisticated experimental setup which is shown in Figure 2.

- To capture the thermal emissions from the back-side of our die, we use a FLIR SC5600 infrared camera with a mid-infrared spectral range of 2.5 μm – 5.1 μm. The camera is capable of operating at 380 Hz with a spatial resolution of 30 μm with a 0.5$\times$ microscopy kit.
- For our test chip we utilize a 90 nm Altera Stratix II (EP2S180) FPGA with 180,000 logic elements with total die dimensions of 23 mm $\times$ 24 mm. Figure 3.a shows the die image of the FPGA. For our experiments, we use a relatively homogenous section of the die that spans 6.4 $mm \times 7\ mm$ as shown in Figure 3.b. This section has 30$\times$ 30 logic array blocks where each logic array block contains 16 logic elements.
- We use an Agilent E3634A power supply to supply and measure the power consumption of FPGA chip.

Test Chip Configuration. We use our experimental FPGA chip to create a grid of *micro heaters*, where each micro heater consists of

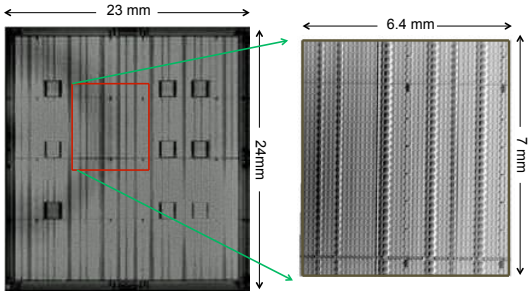

a. The FPGA test chip

b. Die section where power patterns are generated.

Figure 3: Die image of Altera Stratix II EP2S180 device.

15-stage ring oscillator together with one flip-flop that controls its operational status. Each micro heater fits exactly in a logic array block with a spatial size of about 223 $\mu m \times 124\ \mu m$. Each micro heater consumes an average of 7.5 mW and runs at about 140 MHz as measured by an external power supply and frequency meter respectively. The created grid consists of 30 $\times$ 30 micro heaters that are precisely placed at the grid locations using Altera's Quartus II placement assignment editor. To create any desired power pattern, we inject control bits into the flip-flops of the micro heaters to selectively turn on the micro heaters that correspond to the required power pattern. Our experimental novelty of using a chip of programmable micro heaters enables us to achieve the following two technical goals which have not been attained in previous works:

1. The grid structure of the micro heaters, where every micro heater can be selectively controlled, enables us to create any desired spatial power pattern on a real chip. Previous works used processors in which independent control of various processor blocks is infeasible.
2. The regular and homogenous structure of the micro heater grid enables us to estimate the power consumption of each micro heater by simply measuring the *incremental* total FPGA power consumption and dividing it by the number of enabled micro heaters. Hence, we are able to construct the spatial power map through an alternative path that is different from using the thermal infrared emissions. Previous works lacked this ability to validate their spatial power estimations through different means.

In order to capture the chip's thermal emissions it was necessary to remove the heat spreader. While removing the heat spreader is going to change the spatial thermal behavior, the change in spatial thermal emissions does not change the underlying dynamic power consumption [12]. The spatial power consumption remains relatively intact, and the new interactions between temperature and power are captured in the learned $\mathbf{R}$ matrix as described next.

Estimating the Thermal Model Matrix ($\mathbf{R}$). It is first necessary to estimate the thermal to power inversion matrix, ($\mathbf{R}$ is defined in Section 3.2), which is necessary for our model computations. Previous approaches in the literature [5] used a scanning laser beam with known power intensity. For each discretized location of a test chip, the laser beam is focused on the location to transfer power to the die and then the temperature measurements everywhere are recorded. The applied power values at different locations together with the associated measured temperatures are used to learn the values of the matrix $\mathbf{R}$, where the temperature response to a unit power pulse at node i provides the values of column i of $\mathbf{R}$. The modular structure of our micro heater array enables us to measure $\mathbf{R}$ in a much easier and simpler fashion. We utilize a sliding window of a power pulse that consists of a 5 $\times$ 5 micro heaters and for each location of the power pulse, we record the emitted temperatures in our field of measurements as shown in Figure 4. We use the ammeter to measure the exact power consumption of the power pulses at each location and together with the measured temperatures from the camera, we learn the values of the matrix $\mathbf{R}$.

Experiment 1. In the first experiment we assess the accuracy of our thermal to power inversion methodology by testing its power estimates using a number of arbitrary spatial power patterns. The locations of the activated micro heaters in these patterns are illustrated in Figure 5.a; the measured temperature maps after emissivity calibration are illustrated in Figure 5.b; and the estimated power maps from our methodology are illustrated in Figure 5.c. We report the *estimation error* in percentage which is equal to the sum

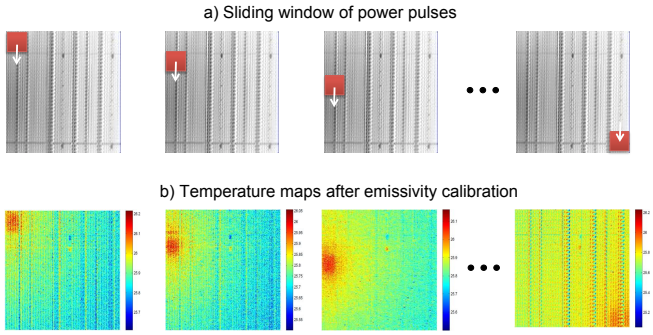

Figure 4: Measuring the power to thermal inversion matrix $\mathbf{R}$ using a sliding window of power pulses.

of the absolute differences between the power estimates and their true values divided by the total power. The average error for these spatial estimations are: 0.0%, 0.0%, 16.7%, 2.8%, 0.0%, 8.3%, 11.1%, and 8.3% with an average estimation error of about 5.9%. By visually comparing the inject power patterns and the estimated ones we notice that our technique recovers the injected pattern to a relatively good extent. Our results illustrate the accuracy of the proposed thermal to power inversion methodology as none of the previous works validated the accuracy of their power estimations.

Experiment 2. To further gain insight into the behavior of thermal to power inversion, we assess the accuracy of our methodology as a function of the spatial frequency of power maps. As discussed earlier in Subsection 3.1, the nature of heat conduction on chips leads to a low-pass filtering effect. Hence, we create checker board patterns of increasing spatial frequencies as illustrated in Figure 6.a which lead to the temperature emissions illustrated in Figure 6.b. The average estimation error of the estimated spatial power estimates are given in in Figure 6.c. The average errors are: 8.3%, 0.0%, 16.7.3%, and 41.7%. The results generally agree with our earlier discussions in Subsection 3.1 that concluded that increasing the spatial frequencies of power patterns leads to a deterioration in the accuracy of the thermal to power inversion due to the impact of low pass filtering.

5. CONCLUSIONS

In this paper we have presented a new methodology for spatial post-silicon power characterization using the infrared emissions from the back of the silicon die. We have elucidated the various limitations that underly thermal to power inversion. We have demonstrated mathematically and empirically how inherent low-pass filtering, discretization, and measurement errors could all compromise the accuracy of power estimation. We have proposed new techniques from regularization theory to improve the accuracy of temperature to power inversion. Furthermore, we have provided techniques to compensate for the emissivities of different chip materials and to measure the thermal conductance matrix. We designed a highly modular, reconfigurable test chip based on an array of micro heaters that are precisely placed in a grid pattern. Our test chip and realistic infrastructure enabled us to validate our methodology by comparing its power estimates against the injected spatial power density maps. Our realistic experimental setup camera has provided deep insights on the challenges that are involved in post-silicon power characterization and has confirmed that our methodology works very well in practice providing power estimates with an average estimation error of about 6%.

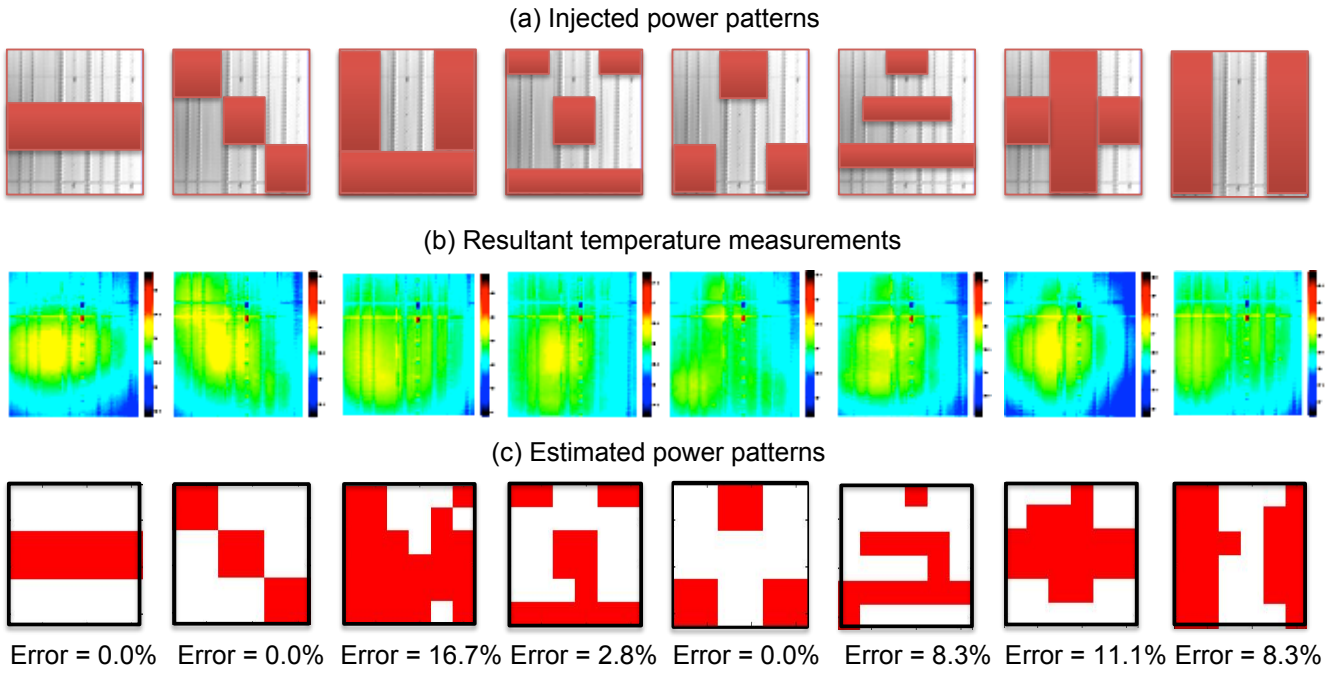

Figure 5: Accuracy of estimating arbitrary power patterns using thermal emissions.

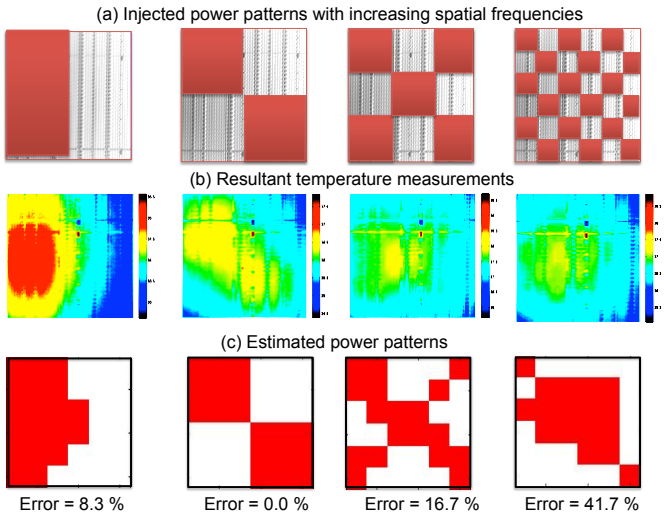

Figure 6: Impact of increasing spatial frequencies of power patterns on the accuracy of thermal to power inversion.

Acknowledgments: This work is partially supported by a DoD ARO grant number 55544-CS-RIP and a NSF CAREER grant number 0952866. S. Reda would like to thank J. Renau from UCSC for the helpful discussions on the setup of infrared imaging systems.

6. REFERENCES

[1] G. C. Albright and J. McDonald. Microthermal Imaging in the Infrared. *Electronics Cooling*, January 1997.

[2] M. Bertero and P. Boccacci. *Introduction to Inverse Problems in Imaging*. Institute of Physics Publishing, 1998.

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

[4] K. Etessam-Yazdani, M. Asheghi, and H. Hamann. Investigation of the Impact of Power Granularity on Chip Thermal Modeling Using White Noise Analysis. *IEEE Trans on Components and Packaging Technologies*, 31(1):211–215, 2008.

[5] H. Hamann, A. Weger, J. Lacey, Z. Hu, and P. Bose. Hotspot-Limited Microprocessors: Direct Temperature and Power Distribution Measurements. *IEEE Journal of Solid-State Circuits*, 42(1):56–65, 2007.

[6] W. Huan, M. R. Stan, K. Sankaranarayanan, R. J. Ribando, and K. Skadron. Many-Core Design from a Thermal Perspective. In *Design Automation Conference*, pages 746–749, 2008.

[7] C. Isci, G. Contreras, and M. Martonosi. Live, Runtime Phase Monitoring and Prediction on Real Systems with Application to Dynamic Power Management. In *International Symposium on Microarchitecture*, pages 359–370, 2006.

[8] S. Lin, G. Chrysler, R. Mahajan, V. De, and K. Banerjee. A Self-Consistent Substrate Thermal Profile Estimation Technique for Nanoscale ICs—Part I:Electrothermal Couplings and Full-Chip Package Thermal Model. *IEEE Transactions on Electron Devices*, 54(12):3342–3350, 2007.

[9] F. J. Mesa-Martinez, M. Brown, J. Nayfach-Battilana, and J. Renau. Measuring Performance, Power, and Temperature from Real Processors. In *International Symposium on Computer Architecture*, pages 1–10, 2007.

[10] M. Pedram and S. Nazarin. Thermal Modeling, Analysis, and Management in VLSI circuits: Principles and Methods. *Proceedings of the IEEE*, 94(8):1487–1501, 2006.

[11] M. Powell, A. Biswas, J. Emer, and S. Mukherjee. CAMP: A Technique to Estimate per-Structure Power at Run-Time Using a Few Simple Parameters. *International Symposium on High Performance Computer Architecture*, pages 289–300, 2009.

[12] J. Renau. Personal Communication.

[13] A. Rogalski and K. Chrzanowski. Infrared Devices and Techniques. *Opto-Electronics Review*, 10(2):111–136, 2002.

[14] A. Srivastava, R. Bai, D. Blaauw, and D. Sylvester. Modeling and Analysis of Leakage Power Considering Within-Die Process Variations. In *International Symposium on Low Power Electronics and Design*, pages 64–67, 2002.

[15] C. R. Vogel. *Computational Methods for Inverse Problems*. Society for Industrial and Applied Math, 2002.