Power Mapping of Computing Devices: Fundamentals and Applications

by
Abdullah Nazma Nowroz

M.Sc., University of Southern California; Los Angeles, CA, 2009
B.Sc., Boston University; Boston, MA, 2006

A dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in School of Engineering at Brown University

PROVIDENCE, RHODE ISLAND

May 2014

This dissertation by Abdullah Nazma Nowroz is accepted in its present form
by School of Engineering as satisfying the
dissertation requirement for the degree of Doctor of Philosophy.

Recommended to the Graduate Council

Date_____ _____

Sherief Reda, Advisor

Date_____ _____

Ruth Iris Bahar, Reader

Date_____ _____

Jacob Rosenstein, Reader

Approved by the Graduate Council

Date_____ _____

Peter M. Weber, Dean of the Graduate School

Vitae

Abdullah Nazma Nowroz was born in Newcastle Upon Tyne, UK. She received her B.Sc. in Electrical and Electronics Engineering with Summa Cum Laude from Boston University in 2006 and M.Sc. in Electrical Engineering from University of Southern California in 2009. Her research focus is post-silicon power characterization using thermal emissions from the backside of silicon die using state-of-the-art infrared camera. She worked on devising frameworks for various applications of post-silicon power mapping, such as, power mapping of real processors, power mapping using sparse thermal sensors measurements, and power mapping for hardware Trojan detections. She also worked in thermal sensor allocation and signal reconstruction techniques that fully characterize the thermal status of the processor using limited number of measurements from the thermal sensors.

abdullah_nowroz@brown.edu
Brown University, RI, USA

Publications:

1. A. N. Nowroz, K Hu, F. Koushanfar and S. Reda, "Novel Techniques for High-Sensitivity Hardware Trojan Detection using Thermal and Power Maps", under review for *IEEE Transactions on Computer-aided Design for Integrated Circuits and Systems*, 2013.

2. K. Dev, A. N. Nowroz and S. Reda, "Power Mapping and Modeling of Multi-core Processors," to appear in International Symposium on Low-Power Electronics and Design, 2013.

3. K. Hu, A. Nowroz, S. Reda and F. Koushanfar, "High-Sensitivity Hardware Trojan Detection Using Multimodal Characterization," *Design Automation and Test in Europe*, pp 1271-1276, 2013. Acceptance rate 25%.

4. A. N. Nowroz, G. Woods and S. Reda, "Power Mapping of Integrated Circuits Using AC-based Thermography," *IEEE Transactions on Very Large Scale Integration Systems*, Vol. PP(99), pp 1-1, 2012.

5. S. Reda and A. N. Nowroz, "Power Modeling and Characterization of Computing Devices: A Survey," in *Foundations and Trends in Electronics Design Automation Journal*, Vol 6(2), pp. 121-216, 2012.

6. S. Reda, A. N. Nowroz, R. Cochran, S. Angelevski, "Post-Silicon Power Mapping Techniques for Integrated Circuits," *ElSevier VLSI Integration Journal*, Vol. 46(1), pp 69-79, 2013.

7. A. N. Nowroz, G. Woods and S. Reda, "Improved Post-Silicon Power Modeling Using AC Lock-in Techniques," *ACM/IEEE Design Automation Conference*, pp. 101-106, 2011. Acceptance rate 24%.

8. A. N. Nowroz and S. Reda, "Thermal and Power Characterization of Field-Programmable Gate Arrays," *ACM International Symposium on Field Programmable Gate Arrays*, pp. 111-114, 2011. Acceptance rate 24%.

9. S. Reda, R. Cochran, and A. N. Nowroz, "Improved Thermal Tracking for Processors Using Hard and Soft Sensor Allocation Techniques," *IEEE Transactions on Computers*, Vol. 60(6), pp. 841-861, 2011. Acceptance rate 23%.

10. R. Cochran, A. N. Nowroz and S. Reda, "Post-Silicon Power Characterization Using Thermal Infrared Emissions," *Proceedings of the International Symposium on Low-Power Electronics and Design*, pp. 331-336, 2010. **Best Paper Award.** Acceptance rate 23%.

11. A. N. Nowroz, R. Cochran and S. Reda, "Thermal Monitoring of Real Processors: Techniques for Sensor Allocation and Full Characterization," *Proceedings of the Design Automation Conference*, pp. 56 - 61, 2010. Acceptance rate 24%.

Acknowledgements

I would like to express my sincerest gratitude to my advisor, Prof. Sherief Reda for his invaluable guidance and support during my graduate study at Brown University. I would like to thank him for all his insights, thought provoking questions and encouragement throughout my Ph.D. experience. I am also thankful to Prof. Iris Bahar and Prof. Jacob Rosenstein to agreeing to serve as members in my dissertation committee even at hardship. Their comments and questions are invaluable to this work.

I would like to thank all of my co-authors, including Prof. Gary Wood, Prof. Farinaz Koushanfar and Kangqiao Hu from Rice University, Ryan Cochran, Kapil Dev and Stefan Angelevski, and of course Prof. Sherief Reda. I owe much of my output over the past four years to their productivity and hard effort. Thank you to Patrick Temple and Sriram Jayakumar for the contribution to the lab infrastructure. I would also like to thank all my colleagues at SCALE laboratory. My stay at Brown was made pleasant and enjoyable because of their friendly support.

Mostly, this work would not be possible without love, encouragement and support of my family. I am grateful to my parents Dr. Abdullahel Bari and Afruja Bari who had been a constant source of love and inspiration in all my pursuits. Their prayer for me is what sustained me thus far. I would like to thank my loving husband Kamrul for his love, understanding and unwavering support throughout this experience. I also would like to thank my dearest siblings, Zaffreen and Sagar for always standing by my side. Everything I have achieved to this point began with what I learned from them.

Abstract of "Power Mapping of Computing Devices:
Fundamentals and Applications" by Abdullah Nazma Nowroz, Ph.D., Brown University, May 2014

Power consumption limits the maximum achievable performance of modern processors. Accurate power modeling is an essential step in both mobile and high-end processors. Large-scale transistor-level power modeling is computationally very challenging. As a result high-level power modeling is performed at the expense of accuracy. Pre-silicon power tools need to be complemented with post-silicon characterization to determine the true power consumption of circuits. Post-silicon power maps are developed during debugging and characterization phases of the first-silicon, and then applied to improve the design during re-spins and for future designs. Post-silicon power results can also improve the accuracy of power and thermal models.

We provide fundamentals for post-silicon power validation using captured thermal infrared emissions from back-side of integrated circuits. We identify challenges associated with thermal-to-power inversion [8]: (1) spatial heat diffusion which blurs underlying power maps, and (2) measurement noise in thermal imaging systems. We address these challenges by devising optimization formulation that incorporates Tikhonov filtering. To further increase the accuracy of power mapping, we propose to use AC-based thermography which reduces the impact of flicker noise and spatial heat diffusion. The average power mapping error reduced from 40% using DC-based method to about 8.5% using the proposed AC-based method. A programmable circuit of micro heaters is implemented, and used to validate our methods and to quantify the improvements in power mapping attained from regularization techniques and AC-based methodology.

Using our novel power mapping framework, we present versatile applications of post-silicon power mapping. First, we power map a soft processor embedded in a FPGA chip, and a real multi-core processor using DC-based techniques, where we decompose the

power-per-block into dynamic and leakage power. We demonstrate methods for utilizing AC framework on real processors. Second, we propose alternate way of post-silicon power mapping with reconstructed thermal maps using measurements from on-chip thermal sensors in-place of infrared imaging, which reduces the cost of post-silicon power mapping drastically. To reconstruct the thermal maps, we utilize full thermal characterization methods based on various frequency-domain techniques. Third, we utilize the high resolution thermal and power maps to formulate high-sensitivity hardware Trojan detection techniques, which are scalable to large circuits.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Problem Characterization

In the past decade power has emerged as a major challenge to computing advancement. A recent report by the National Research Council (NRC) of the National Academies highlights power as the number one challenge to sustain historical improvements in computing performance [28]. Power is limiting the performance of both mobile and server computing devices. At one extreme, embedded and portable computing devices operate within power constraints to prolong battery operation. The power budgets of these devices are about tens of milli-Watts for some embedded systems (e.g., sensor nodes), 1–2 Watts for mobile smart phones and tablets, and 15–30 Watts for laptop computers. At another extreme, high-end server processors, where performance is the main objective, are increasingly becoming hot-spot limited [33], where increases in performance are constrained by a maximum junction temperature (typically 85 °C). Economic air-based cooling techniques limit the total power consumption of server processors to about 100–150 Watts, and it is the

spatial and temporal allocation of the power distribution that leads to hot spots in the die that can comprise the reliability of the device. Because server-based systems are typically deployed in data centers, their aggregate performance becomes power limited [7], where energy costs represent the major portion of total cost of ownership.

The emergence of power as a major constraint has forced designers to carefully evaluate every architectural and design feature with respect to its performance and power trade-offs. This evaluation requires pre-silicon power modeling tools that can navigate the rich design landscape. Every architectural feature has to be judged in terms of its performance, area, power and reliability. A typical design space has an exponential number of possible combination of settings for the various features. Thus, there is a strong need for power modeling methods that enable designers to efficiently explore the design space and to evaluate the impact of various high-level system architectural choices and optimizations on power consumption. These architectural features and choices vary by the medium of the computing substrate. For multi-core processors, the choices include, for example, pipeline depth, instruction issue width, and cache sizes. For SoC-based embedded systems, the choices include the functionality of the custom blocks and the on-chip communication architecture (e.g., network topology, buffer sizes and transfer modes). In some embedded systems, the boundary between hardware (HW) and software (SW) is fluid, where the choice of the implementation (SW or HW) of every component could be decided based on its impact on performance, power, and area. Field-programmable gate array (FPGA) power modeling is also challenging as the user's design is not known during the design and fabrication of the FPGA.

Various techniques have emerged for pre-silicon power modeling [25, 13, 105, 77, 63, 12, 49, 30, 46, 54, 101, 92, 62]. Low-level power estimation tools which rely on transistor-level circuit characteristics, gate-level switching activities and detailed RTL netlists can be very accurate, but is highly infeasible for large designs. To improve computation effi-

ciency, designers rely on high-level power estimation tools at the early stage. Some notable techniques for high-level power estimation are macro-modeling, analytical modeling and regression modeling [57, 30, 16, 101, 92, 49, 77]. Macro-modeling is a bottom-up approach where black-box models (macro-models) are built for circuit blocks by a process of characterization that models the block power as a function of the input/output signal statistics (probabilities) of the block. Other details can also be used to query the power models, such as the bus width, average capacitance, etc. This approach of power modeling is suitable for blocks with irregular structure, such as, arithmetic logic unit (ALU), control unit etc. Analytical power modeling techniques attempt to correlate power consumption of circuits block with mathematical equations, which are suitable for functional blocks which are organized in a regular way, such as, cache memory, translation lookaside buffers (TLBs) and register file. Regression modeling is a statistical inference method, where the relationship between a dependent variable power and independent variable architectural parameters is established. While it is highly desirable, these pre-silicon power estimation tools are inevitably inaccurate, or approximate [65] because these tools trade-off accuracy for speed.

Many of the high-level power modeling tools estimates power per block. Figure 1.1 gives one such example, the data is taken from power estimation work by Natarajn *et. al* [68]. Here high-level per-block power has been estimated, but validating the accuracy of the estimation still remains a fundamental challenge. In fact, such a pie chart is extremely difficult to measure directly, since there are no on-chip ammeters to determine how current or power divides between different CPU units. McPAT [35], which is a recently proposed power modeling tools for multi-core processor, has compared the estimated total peak dynamic power with post-silicon total measured chip power. But no validation has been done for per-block power, or the estimated leakage power. The accuracy of these tools could be 30% from the actual power consumption [13]. For all computing substrates

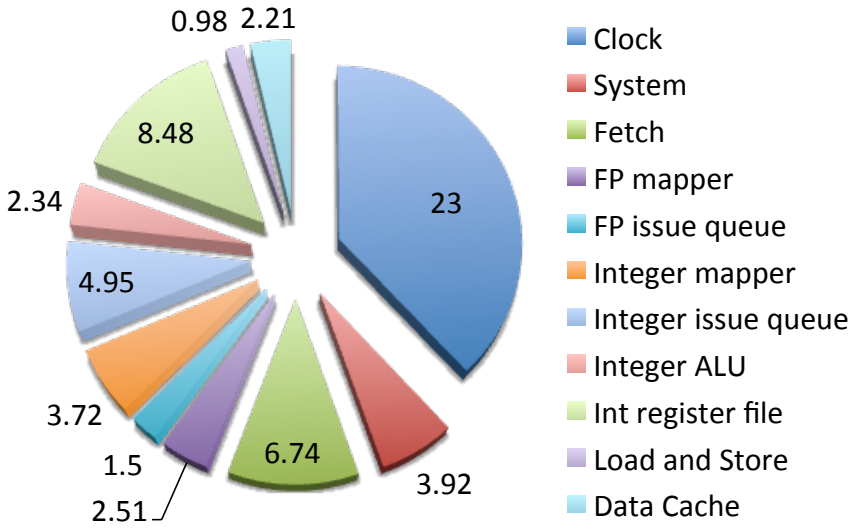

Figure 1.1: Average estimated per-block power distribution (W) of Alpha 21264 pipeline design.

estimating power consumption of each circuit block in the chip accurately is very crucial in making good design decisions at an early stage. So, to gain trust and improve accuracy of pre-silicon power estimation tools, these tools need to be validated with post silicon measurements.

To validate pre-silicon design estimates, to calibrate power-modeling CAD tools, and to estimate the impact of variabilities introduced during fabrication, there is need for post-silicon characterization. Furthermore, post-silicon methods can guide design changes

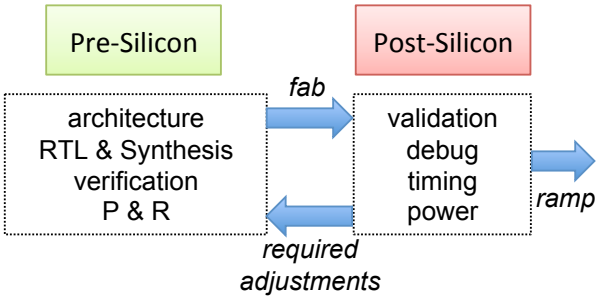

Figure 1.2: Post silicon characterization.

(layout, sensor-locations, etc.) during the period between the first-silicon and the product-silicon as shown in Figure 1.2. In a state-of-the-art custom chip, it is expected to spend about half of the time to market in the design phase, starting with architecture and culminating with placement, routing, timing closure, and meeting the power specifications. The initial design could take more than a year, depending on complexity and IP re-use [19]. Once completed, the initial design is shipped for fabrication and first silicon is received about three months later. At this point, the silicon still has to be debugged, nearly always requiring at least one major re-spin of the design [93]. The debug phase, including working out yield problems as volume is ramped up, can easily take as long as the initial design where in many cases the mismatch between pre-silicon and post-silicon requires major changes in the design and implementation of the chip. Popular examples of large mismatches between pre- and post-silicon estimates include IBM's Cell processor, where the post-silicon power and thermal measurements led to large changes in its specifications and implementation [31]. Once a design is implemented and a physical prototype is available for direct measurements, post-silicon characterization methods should be utilized to perform required design adjustments.

The most versatile approach for post-silicon power characterization is to capture the thermal emissions from the back side of the die and to use these captures to estimate the power [33, 61, 17]. Post-silicon power mapping involves many challenges at both the experimental and modeling fronts as follows.

1. At the experimental front, it is necessary to use a infrared transparent heat sink. For oil-based sink, it is required to control the speed and temperature of the oil flow on top of the chip to remove the generated heat, while maintaing good optical transparency to the infrared imaging system.

2. It is essential to accurately synchronize all the measurements of the system, includ-

ing thermal maps, fluid state measurements and total power consumption.

3. There are two major challenges in infrared imaging; spatial heat diffusion and measurement noise in the thermal imaging system. Heat diffusion blurs the underlying power map and reduces the accuracy of post-silicon power maps as it filters out the spatial high-frequency power patterns. Small noise in temperature measurement can be amplified into significant error in power estimation.

4. At the processing front, challenges include the need to model the relationship between power consumption and temperature.

5. Decomposing the total power into leakage and dynamic is a challenging task due to the dependency of leakage on process variability and temperature.

1.2 Major Contributions

There has been a diversification in possible computing substrates that offer different trade-offs in performance, power, and cost for different applications. These substrates include application-specific custom-fabricated circuits, application-specific circuits implemented in field-programmable logic arrays (FPGAs), general-purpose processors whose functionality is determined by software, general-purpose graphical processing units (GP-GPUs), digital signal processors (DSPs), and system-on-chip (SoC) substrates that combine general-purpose cores with heterogeneous application-specific custom circuits. None of these substrates necessarily dominate the other, but they rather offer certain advantages that depend on the target application and the deployment setting of the computing device. Enormous amount of computing devices have emerged utilizing these substrates in last few decades. Portable computing devices include smartphones, laptops, notebooks, tablets and stationary computing devices include desktop computers and high performance

server computers. While power modeling and characterization for these devices share common concepts, each of these devices has its own peculiarities and bring new challenges to pre-silicon power modeling and post-silicon power characterization. In addressing the challenges inherent to power modeling and estimation, we provide fundamentals of post-silicon characterization where we propose an accurate, detailed framework for post-silicon power mapping which is applicable across all the fabrics. The focus of this work is to devise algorithmic and experimental methods for post-silicon power mapping by using infrared imaging techniques, where the captured thermal emissions from the backside of the die are inverted to yield the underlying spatial power maps. The major contributions of this thesis are as follows.

1. **Fundamentals of Post-silicon Power Mapping using DC-based methods:** We propose a novel DC-based methodology that provides accurate, detailed post-silicon spatial power estimates using the thermal infrared emissions from the backside of silicon die [17, 84]. 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 using quadratic optimization formulation that incorporates Tikhonov filtering techniques to find the most accurate power maps. 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 which shows 30% improvement in power mapping accuracy over previous approaches.

2. **Fundamentals of Post-silicon Power Mapping using AC-based methods:** Two major challenges in thermal to power inversion is the spatial diffusion of heat and measurement noise. Heat diffusion blurs the underlying power map and reduces the accuracy of post-silicon power models as it filters out the high spatial frequency power patterns. AC excitation reduces the amount of spatial heat diffusion as the AC excitation frequency increases [55, 10]. AC excitation also has the benefit of reducing flicker noise making the inversion more stable. In contrast to previous usages of AC-based thermography, we use AC-based thermography to improve post-silicon power mapping, where AC excitation signals are applied to the chip instead of DC excitation signals [71, 72]. We prove and demonstrate that using AC excitation sources reduces the impact of flicker noise and spatial heat diffusion, which translates to significant improvements in power mapping accuracy by reducing the average error to 8.5% compare to 40% in the DC case. We devise a lock-in based thermal to power inversion methodology that maps spatial power consumption on a real chip. By using a custom test chip that can be programmed to control the spatial and temporal power consumption, we perform a number of experiments. We use the test chip to analyze the noise in our thermal imaging system, and to quantify the improvements in power mapping attained from the proposed AC-based methodology. We elucidate the impact of the AC excitation frequency on both the signal-to-noise ratio and power mapping accuracy.

3. **Applications of Post-silicon Power mapping:** The proposed DC and AC-based post-silicon power mapping overcome many challenges that results into a very accurate post-silicon power estimation framework. These methods can be applied to various computing devices to get access to detailed accurate spatial post-silicon power maps. We apply our proposed methods to two different computing substrates, an embedded soft processor in a FPGA and a real multi-core processor. We propose an alternative way to perform the proposed post-silicon power mapping by using

measurements from available on-chip thermal sensors. Furthermore, we formulate various techniques for hardware Trojan detection using the detailed spatial power maps. Following are the three different applications of the post-silicon power mapping.

(a) **Power Mapping of Real Processors:** We apply our DC-based post-silicon power mapping to two different real processors, firstly to an embedded soft processor [70] and secondly to a multi-core processor [21]. For the embedded processor, we invert the thermal emissions of the soft processor while running a standard application, into spatial power estimates. We consider different configurations of the soft processor to evaluate the accuracy of the power mapping. For the multi-core processor, we propose various new techniques to tackle challenges that accompanies power mapping and modeling of multi-core processors. To decompose the estimated power maps of the multi-core processor to dynamic and leakage power, we utilize thermal conditioning techniques to build leakage power models for the die and use these models to analyze within-die spatial leakage variations. In our experiments, we capture thermal images from a 45 nm quad-core processor under different workload conditions, and then we reconstruct the dynamic and leakage power maps for different blocks. Our results show great accuracy in mapping and modeling, revealing good insights into the trends of power consumption in multi-core processors. We also demonstrate the basic applicability of our AC-based power mapping technique on a dual-core processor [72].

(b) **Power Mapping from Sparse Thermal Sensor Measurements:** We propose to apply post-silicon power mapping on thermal maps reconstructed from thermal sensor measurements [69] instead of the thermal maps obtained by infrared imaging. We characterize temperature signals of real processors and demonstrate that on-chip thermal gradients lead to sparse signals in the fre-

quency domain. We exploit this observation to devise signal reconstruction techniques that fully characterize the thermal maps of the processor using the limited number of measurements from the thermal sensors. Using the reconstructed thermal maps, we perform the post-silicon power mapping procedure to obtain spatial power estimates. We present results using a custom FPGA chip and a real processor.

(c) **Power Mapping for Hardware Trojan Detection:** Vulnerability of modern integrated circuits (ICs) to hardware Trojans has been increased considerably due to the globalization of semiconductor design and fabrication. In this work, we present fast and easy-to-implement methods for Trojan detection that is based on post-silicon thermal and power characterization technique and 2-dimensional principal components analysis (2DPCA) [37]. Our approach first estimates the detailed post-silicon spatial power consumptions using thermal emissions of the IC, then applies 2DPCA to extract features of the spatial power consumptions, and finally uses statistical tests against the features of authentic ICs to detect the Trojan. To characterize real-world ICs accurately, we add 20% - 40% process variation to gates' length, width and oxide thickness of ICs in our experiments. We have designed Trojans modules with varying power consumption. The results reveal that our experiments can detect Trojans with power consumptions as small as 0.05% to 0.2% of the total power consumption of the chip.

The remainder of this thesis is structured as follows. Chapter 2 presents the required background for power modeling in ICs, challenges in pre-silicon power modeling and motivation for post-silicon characterization, related works and basics of infrared imaging including detail techniques for thermal calibration. We provide the fundamentals of post-silicon power mapping in Chapter 3 and 4. Chapter 3 presents our framework for post-

silicon power mapping using infrared emissions using DC-based techniques, and Chapter 4 describes techniques for post-silicon power mapping using AC-based thermography. In Chapter 5, we give detail description of three different applications for post-silicon power mapping. Finally, in Chapter 6, we summarize our findings and outline directions for possible research extensions.

Chapter 2

Background

2.1 Power Consumption Mechanisms

Power consumption of digital Complementary Metal Oxide Semiconductor (CMOS) circuits is caused by two mechanisms. The first mechanism is *dynamic power*, which arises when signals transition their values, and the second mechanism is *static power*, which causes circuits to dissipate power when no switching activity is occurring. One of the main advantages of using CMOS technology over earlier bipolar technology was that CMOS circuits consumed power only during circuit switching. However, aggressive technology scaling in the past decade led to a situation where static power is no longer negligible, but rather a significant contributor to total power consumption.

Dynamic Power: Logic gates implemented in CMOS chips use two complementary transistor types, NMOS and PMOS, to build the functionality of each gate. One terminal of PMOS transistors is typically connected to the voltage supply, V_{DD}, while one terminal of NMOS transistors is connected to the ground voltage V_{GND}. Figure 2.1 gives the

schematic of an inverter gate that consists of one NMOS transistor and one PMOS transistor. To understand the operation of the gate, assume that the input voltage is first at logic 0 (i.e., V_{GND}). In this case the PMOS transistor is in `on` state with a very low resistance (ideally 0), while the NMOS transistor is in `off` state with a very high resistance (ideally ∞), and a path exists to charge the load capacitance C_L until the output voltage reaches V_{DD}. The *load capacitance*, C_L, represents the total capacitance arising from the output diffusion capacitances of the two transistors, the input capacitances of fan-out gates, wiring capacitance, and parasitics. When the input voltage switches to a logic 1 (i.e., V_{DD}), the PMOS transistor is in `off` state with a very high resistance (ideally ∞), while the NMOS transistor is `on` state with a very low resistance (ideally 0), and a path exists to discharge the charges on the load capacitance to the ground until the output voltage reaches 0. The sum of energy consumed during the charging and discharging, i.e., the energy per cycle, is equal to $C_L V_{DD}^2$. The *dynamic power* consumed, which is the switching energy per second, by the gate is equal to

$$P_{dynamic-gate} = sC_L V_{DD}^2, \tag{2.1}$$

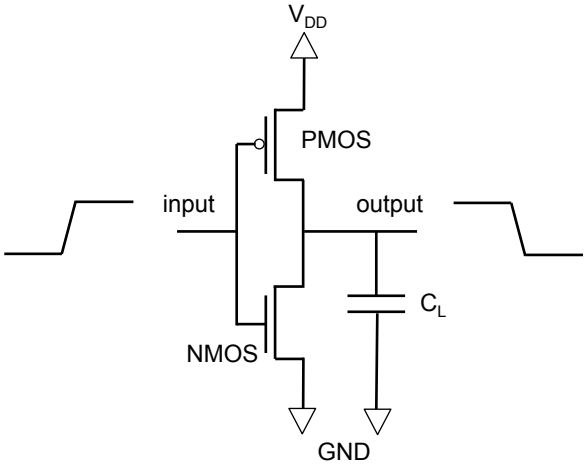

Figure 2.1: CMOS inverter.

where s is an activity factor that denotes the number of switching cycles per second. If a circuit has N gates, then the total dynamic power is equal to

$$P_{dynamic} = \sum_{i}^{N} s_i C_{L_i} V_{DD}^2, \tag{2.2}$$

Another power component that is incurred during switching is *short circuit power*. If the transition edge (from 1 to 0 or 0 to 1) of the input signal is not sharp, there will exist a brief moment of time where the NMOS and the PMOS transistors are both turned on and current will flow from the supply terminal to the ground. Short circuit power is incurred only whenever a switching activity occurs, which makes it proportional to dynamic power consumption. Its exact value is determined by the slopes, or transition times, of the input and output signals. With proper circuit design, short circuit power is usually about 10% of the dynamic power [29].

Static Power: Static power is the power consumed by transistors when they are not switching. When CMOS logic gates do not switch, they have no electrical path between the supply terminal, V_{DD}, and the ground terminals. Thus, static or *leakage current* was historically negligible; however, aggressive scaling in sub-100 nm technologies has led to a substantial increase in its magnitude. Note that modern computing devices include analog components (e.g., phase-locked loops, sense amplifiers) that incur static DC power consumption. For digital CMOS, there are three main components for leakage current: (1) subthreshold leakage current between the transistor's source and drain; (2) gate leakage between the transistor's channel and gate; and (3) reverse bias current between the transistor's drain and well [88, 2]. Figure 2.2 illustrates these three leakage components. With the introduction of high-k dielectrics, the main source of leakage current is the subthreshold leakage.

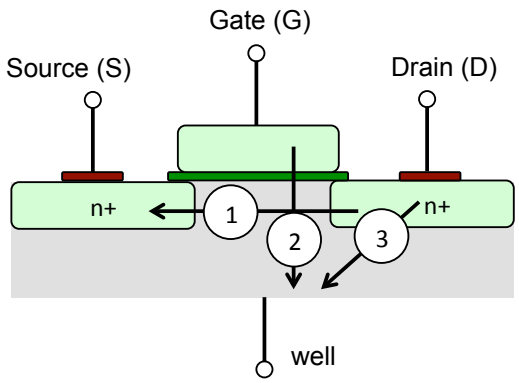

Figure 2.2: Three components of leakage current in a MOS transistor.

MOS transistors operate by modulating an energy barrier between the source and drain of the transistor. The height of the energy barrier is called the *threshold voltage* (V_{th}). By increasing the potential difference between the gate and the source, $|V_{GS}|$, it is possible to reduce the energy barrier enabling more flow of electrical carriers between the source and the drain of the transistor. During the off state when $V_{GS} = 0$, the average carrier's energy is lower than the barrier's energy; however, the carriers do not have a uniform energy distribution, and there is a probability that some carriers will have higher energy than the height of the barrier, enabling them to leak from the source to the drain [28]. The probability of a carrier having an energy higher than the average energy drops exponentially with a factor that is proportional to temperature. The subthreshold leakage current can be mathematically expressed as

$$I_{leak} = I_o e^{\frac{-qV_{th}}{\alpha kT}}, \tag{2.3}$$

where I_o is a constant that depends on the transistor's geometrical dimensions and fabrication technology, α is a number greater than one, T is the temperature of the transistor, q is the charge of the electrical carrier, and k is the Boltzmann constant [28]. Equation (2.3) reveals a number of sensitivity factors that impact leakage [1], such as, process sensitivity

and temperature sensitivity. While the leakage current of an individual logic gate also depends on its input vector, the sum of these vector-dependent variations typically average out for large circuits; furthermore, they are dwarfed by the impact of temperature and V_{th} variations [1]. The sensitivities of leakage power (i.e., $V_{DD} \times I_{leak}$) to the supply voltage, threshold voltage, and temperature introduce variabilities in leakage modeling.

2.2 Power Modeling and Estimation

2.2.1 Challenges in Pre-silicon Power Modeling

While computer-aided power analysis tools can provide power consumption estimates for various circuit blocks, these estimates can deviate from the actual power consumption of working silicon chips. There are a number of reasons these pre-silicon estimates can deviate from the actual post-silicon power measurements.

- *Large input vector space:* The most significant obstacle in trying to estimate power dissipation is that the power is pattern dependent. In other words, it strongly depends on the input patterns being applied to the circuit. The exponentially vast number of possible input vector sequences and the billions of transistors implemented in current designs make it impossible to simulate the power consumption incurred from every possible input vector sequence or software application.

- *Spatiotemporal correlation between signals:* Probabilistic approaches are commonly used to solve the pattern-dependence problem [29]. But the problem becomes more difficult when re-convergent fanout introduces correlations between input signals. In order to achieve good accuracy, one must model the correlations between internal

node values, which can be very expensive [66]. As a result, the tools typically apply heuristics that trade off accuracy for speed in order to obtain the necessary transition probability estimates [65].

- *Errors in coupling capacitance estimation:* Coupling capacitance between neighboring wires is determined by the exact waveform activity experienced by the wires. The computational in-feasibility of simulating every input vector automatically makes estimating the coupling capacitance during design time a difficult process. Wrong estimates for the coupling capacitance impact switching power the following ways [29]. Wrong capacitance values incorrectly estimate incomplete voltage transitions arising from crosstalk noise. They impact signal slew times which determine short circuit power. Furthermore, they impact the signal timing delays which determine the occurrence of glitches.

- *Process variations:* Intra-die and inter-die process variations are unique for each die. Process variations impact leakage power, and they also impact the signal delays along the circuit paths, which determine glitches [9, 103, 73, 83]. Process variations can introduce serious deviations in power consumption compared to pre-silicon design time estimates.

- *Glitch power:* It has been observed that power dissipation due to glitch signals is typically 20% of the total power, but can be as high as 70% of the total power in some cases such as combinational adders [91]. This component of the power dissipation is computationally expensive to estimate, because it depends on the timing relationships between signals inside the circuit. Consequently, many power estimation techniques ignore this issue to improve computation time.

- *Spatiotemporal thermal variations:* Incorrect estimates for dynamic power implies incorrect estimation of thermal variations. The inaccuracy in thermal variations

estimation in turn leads to further deviations in leakage power estimates because leakage power is strongly dependent on thermal status of the chip.

Because of the aforementioned complexities, pre-silicon power estimates might not be accurate when compared to the real post-silicon characteristics. In many cases the mismatch between pre-silicon and post-silicon characteristics forces major changes in the design and implementation of integrated circuits. In a study from 2005, it was shown that 70% of new designs require at least one design re-spin to fix post-silicon problems and that 20% of these re-spins are due to power and thermal issues [93]. It is likely that these figures are much higher now as chips are more complex with a larger number of transistors.

2.2.2 Post-silicon Characterization

Once a design is fabricated, the manufactured devices can yield a wealth of power characterization data that can improve various design choices and runtime applications. The produced devices can be directly characterized for their true power consumption under various loading conditions with no need for approximate models or simulations as in the case with the pre-silicon design. One of the main research directions for post-silicon characterization is detailed power mapping for validation. Design-time power estimates are approximate. By measuring total electrical current and thermal status of a fabricated device, the true power estimates of the circuit structures are revealed. This is possible because of two reasons, (1) the inherent relationship between power and temperature can be quantized into a modeling matrix, (2) the thermal emission which is a representative of the thermal status of the chip can be obtained from the back side of the silicon die. By combining the thermal emissions with the modeling matrix, the detailed power consumption of

the chip can be estimated. These post-silicon power estimates can be used to validate design time power estimates, calibrate various empirical models used by CAD tools, or make re-adjustments during design re-spin if necessary. We describe the relationship between chip power and temperature and basics of thermal imaging as follows.

a. Relationship between Power and Temperature

The relationship between power and temperature is described by the physics of heat transfer. Mainly, heat diffusion governs the relationship between power and temperature in the bulk of the die and the associated metal heat spreader, while heat convection governs the transfer of heat from the boundary of the integrated metal spreader/sink to the surrounding fluid medium which is either air or liquid. The heat diffusion equation is given by

$$\bigtriangledown \cdot (k \bigtriangledown T(x,y,z)) + p(x,y,z) = \rho c \frac{\partial T(x,y,z)}{\partial t}, \tag{2.4}$$

where $T(x,y,z)$ is the temperature at location (x,y,z), $p(x,y,z)$ is the power density at location (x,y,z), ρ is the material density, c is the specific heat density, and k is the thermal conductivity [74]. k, ρ, and c are all functions of location [74]. In the steady-state DC case, the heat diffusion equation can be written as follows,

$$\bigtriangledown \cdot (k \bigtriangledown T(x,y,z)) = -p(x,y,z), \tag{2.5}$$

where $T(x,y,z)$, $p(x,y,z)$, and k are the temperature, power density, and the thermal conductivity at location (x,y,z) respectively [74]. The power transferred at the boundary by convection is described by Fourier's law for heat transfer, and it is proportional to the

temperature difference between the boundary of the heat sink and the ambient temperature. The constant of proportionality is the heat transfer coefficient which depends on the geometry of the heat sink, the fluid used for heat removal and its convection characteristics (e.g., speed and laminarity / turbulence) [40].

In any practical implementation, the heat equation must 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 is represented by a vector $\mathbf{T}$ that gives the *true* temperatures at a discrete set of back-side die locations. The length of the vector $\mathbf{T}$ is determined by the spatial resolution of the infrared camera and the dimensions of the die. The continuous power signal is represented by a vector $\mathbf{p}$ that gives the power consumption of each of the circuit's units. In such case, Equation (2.5) is approximated by the following linear matrix formulation,

$$\mathbf{Rp} = \mathbf{T}, \tag{2.6}$$

where the matrix $\mathbf{R}$ is a linear operator that captures the impact of all modes of heat transfer, which is called the *modeling matrix* for the system.

b. Basics of Infrared Imaging

Any body above absolute zero Kelvin emits infrared thermal radiation with an intensity that depends on its temperature, its emissivity, and the radiation wavelength [86]. Transistors and interconnects in integrated circuits operate at elevated temperatures due to resistive heating by charge carriers [75]. Modern computing integrated circuits use flip-chip

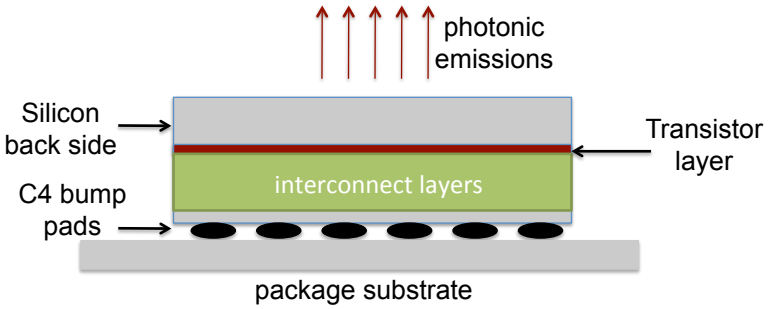

Figure 2.3: Infrared Imaging of silicon back side.

packaging, where the die is flipped over and soldered to the package substrate. By removing the package heat spreader, one can obtain optical access to every device on the die through the silicon backside. It is possible to use infrared-transparent heat removal techniques that have similar thermal characteristics to the original heat spreader [81, 59] or to numerically translate the thermal maps to the original thermal maps.

Figure 2.3 shows one such setup, where the emission from the transistor layer can be captured from the silicon back side. Silicon is transparent to infrared emissions with photon energies that are less than its bandgap energy (1.12 eV), which corresponds to wavelengths larger than 1.1 μm. This transparency is ideal from an infrared imaging perspective as it enables the capture of photonic emissions from the devices, which provides valuable information for thermal and power characterization of computing devices operating under realistic loading conditions.

One of the key specifications of an infrared imaging system is its *spectral response*, which determines the part of the infrared spectrum that the imaging equipment can detect. For the range of temperatures encountered during chip operation, the mid-wave infrared (MWIR) range, which stretches from 3 μm to 5 μm shown in Figure 2.4, yields the most sensitive and accurate characterization of thermal emissions. Detecting emissions in the MWIR range requires the use of InSb quantum detectors which have to be cooled to

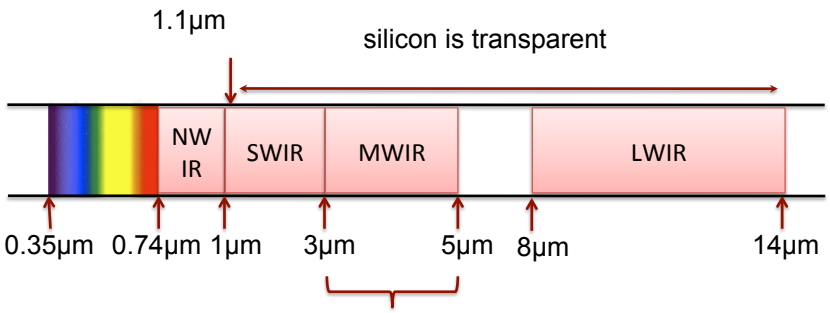

Figure 2.4: Infrared wavelength range.

cryogenic temperatures to ensure sensitivity. As a consequence, high-resolution MWIR imaging systems tend to be fairly expensive.

To capture the thermal emissions from an operational die, it is necessary to remove the optical obstruction introduced by traditional heat removal mechanisms (e.g., integrated heat spreader, metal heat sink and fan), and to substitute them with mechanisms that can remove heat while being transparent to infrared emissions. The standard technique to achieve such a system is through the use of infrared transparent oil-based heat sink [33, 60, 81, 59, 69]. For our experiments with real multi-core processors, we machined a special sapphire oil-based infrared-transparent heat sink that precisely controls the oil flow over the computing device shown in Figure 2.5. Chilled oil is forced into an inlet valve of the sink, which then flows on top of the die to sweep the heat and then exits through an outlet valve. The oil maintains its flow using an external pump, and the temperature of the oil is controlled using a thermoelectric cooler. By controlling the oil flow and its temperature, as well as the sapphire window dimensions, it is possible to get similar

Figure 2.5: Infrared-transparent oil-based heat removal system.

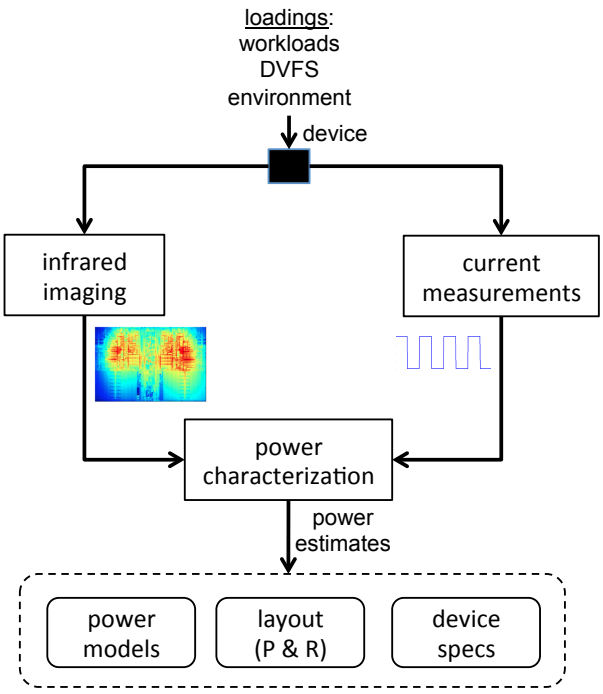

Figure 2.6: Using post-silicon power characterization for design validation and CAD tool calibration. P & R stands for placement and routing.

thermal characteristics as the original heat removal package [81, 59]. For our experiments with an FPGA test chip, we selected 10% of the chip as our test area to keep the power consumption moderate. After removing the heat spreader from the top of the FPGA a fan from the side at a fixed location was enough to maintain the chip temperature below maximum tolerance level such as 85° during operation.

c. Need for Post-silicon Power Mapping

Post-silicon power characterization of integrated computing devices provides the true spatial and temporal power characteristics under representative loading conditions, such as workloads and dynamic voltage and frequency (DVFS) settings. Figure 2.6 gives an integrated framework, where measurements from infrared imaging equipment together with electrical current measurements are used for power characterization. The electric current

measurements can provide the total power consumption. If the chip supports multiple independent power supply networks, then it is possible to increase the pool of current measurements. The thermal emissions captured from the infrared imaging system, together with the lumped electrical current measurements, can be inverted to yield high-resolution spatial power maps for the individual circuit blocks. The results of post-silicon power characterization can aid the IC design and validation process in the following ways.

1. High-level power modeling tools rely on the use of parameters that are estimated from empirical data [13, 77]. The results from post-silicon power characterization can be used to calibrate and tune high-level pre-silicon power modeling tools.

2. The power characterization results can drive heat sink design. Passive heat sinks remove heat indiscriminately from the die, and thus, their design is mainly driven by total power consumption. Active heat removal systems, such as thermoelectric Peltier coolers [89], can make use of the true post-silicon power characterization results to maximize their heat removal capabilities.

3. To evaluate "what if?" design re-spin questions, the power characterization results can substitute the power simulator estimates and directly feed the thermal simulator. For example, if the thermal characterization results are unacceptable, then the layout can be changed to reduce the spatial power densities and hot spot temperatures. The power characterization results are fed together with the new layouts to the thermal simulator to evaluate the impact of the changes.

4. The post-silicon power characterization results can also force a re-evaluation of the computing device specifications (e.g., operating frequencies and v-oltages).

5. Post-silicon power characterization can reveal valuable information for chip-wide process variability estimation. Leakage power is dependent on process variation.

Design time leakage power is estimated using statistical methods or analytical equations, whereas post-silicon power mapping of leakage power is estimated using real runtime chip temperature.

6. IC Trojans are implemented by unsought chip modifications during the third-party fabrication process. Infrared-based power mapping techniques can generate high-resolution spatiotemporal power maps, which can be used for non-intrusive Trojan detection techniques.

2.3 Related Work on Post-Silicon Power Mapping

Power related issues in modern multi-core processors have made post-silicon power analysis a necessity in IC design flow. To conduct fine-grained post-silicon power analysis, there have been two main research directions: 1) power characterization using infrared emissions, and 2) dynamic power characterization using performance counters. The first direction uses thermal infrared emissions to estimate the power consumption of different circuit blocks [33, 59, 61]. The second direction uses the embedded frequency counters in most modern processors to estimate the switching activity of various circuit blocks [42, 77]. These switching activities are combined with design-time estimates of the capacitances of various circuit blocks to compute the dynamic power consumption of the blocks. These performance counters based methods need their own calibration. Furthermore, the second approach cannot estimate leakage power and its spatial 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. The first direction, which we follow in this work, is more generally applicable and powerful.

An industrial team, Hamann *et al.* [33] introduced experimental techniques for post-silicon characterization which allows spatially resolved imaging of microprocessor power (SIMP), where infrared imaging is performed with infrared transparent heat sink to get the thermal maps of the microprocessor. The measured temperature field is then de-convolved to obtain the underlying power map. In the proposed technique, the spatial steady-state power consumption is estimated by minimizing the total squared error between the temperatures computed from the estimates and the actual temperature measurements. Firstly, one of the most important factors in estimating post-silicon power is to have an accurate modeling matrix $\mathbf{R}$ which relates power to temperature. Hamann et al. [33] constructed the modeling matrix by using a laser measurement setup that injects individual power pulses on the actual chip and measures the resultant response. Power sources are provided to the chip in form of a scannable focused laser beam, where the laser mimics individual power sources of the processor. The drawback of this setup is that it needs expensive automated scanning of the laser beam by using a pair of galvo-directing mirrors and it estimates the modeling matrix $\mathbf{R}$ not on the operating chip, but on a dummy chip, which can introduce error. Furthermore, a fairly simple least square formulation is used to invert power to temperature, which does not overcome the inherent challenges in post-silicon power mapping.

A second approach by an academic team also used an infrared transparent heat sink setup to obtain the die temperature of the processor. This group used combinatorial optimization formulation based on genetic algorithms to find a power solution that leads to the measured emissions [60, 59, 61]. The power solution is resolved into dynamic power and leakage power. The authors used analytical power equations with constant parameters that relate the power of a floorplan block to the average temperature of the corresponding block. The genetic algorithm finds various power equation parameters for each floorplan block after gathering results for several benchmarks at different ambient temperatures and

processor activities. Using only the average temperature per block can introduce inaccuracies due to the presence of hotspots in the chip. Furthermore, the drawbacks of genetic algorithms are that they are based on heuristic, require a lot of characterization data and are not known to find the best solution in many cases.

Moreover, the shortcomings of the previous works are that both approaches did not address the inherent limitations of resolving spatial power from temperature measurements, such as, measurement noise and spatial low pass filtering. Previous approaches mainly focused on minimizing the total squared error between the temperatures computed from the power estimates and the thermal measurements [32, 33, 61]. These approaches can produce suboptimal results as they ignore the ill-posed nature of the problem where measurement noise and spatial diffusion reduce the accuracy of power estimation. Compared to previous approaches, our proposed numerical techniques handle many of the challenges associated with inversion. We leverage regularization techniques to reduce the impact of noise and improve the numerical instability. We introduce AC-based technique to diminish the effect of spatial low-pass filtering and flicker noise.

More importantly, previous works did not provide any validations on the accuracy of their spatial power estimates. Previous studies presented spatial power estimates for the different processor blocks "as is" [33, 61]. This lack of validation is the result of the experimental setup that previous studies chose. Earlier experimental setups used general-purpose processors (a dual core PowerPC 970MP [33] and an AMD Athlon 64 processor [61]) in which it is impossible to fully control the underlying spatial power consumption, and there exists no alternative means to verify the spatial power maps. To address this issue, we design our experimental test chip 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.

Chapter 3

Fundamentals of Post-silicon DC-based Power Mapping

3.1 Introduction

Our objective is to provide accurate post-silicon spatial power estimates for the various circuit blocks using the runtime thermal infrared emissions emitted from the back side of semiconductor chips. Spatial power mapping from thermal emissions has emerged in the past few years through the research work of two groups (Hamann *et al.* [33] and Renau *et al.* [61]) as discussed in Chapter 2. The objective of this chapter is to provide novel accurate methods for post-silicon power mapping framework. The contributions are as follows:

1. 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 spatial power estimation accuracy.

2. To address the outlined challenges we propose techniques from regularization theory. We propose quadratic formulations that are augmented with Tikhonov regularization methods to improve the accuracy of spatial power estimation. Our techniques limit the impact of measurement errors and improve power estimation compared to previous techniques.

3. 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 power maps with various intensities and spatial constructions. We exercise the circuit with a number of different spatial power maps and then capture the thermal emissions in the mid-wave infrared range from the backside of the chip using a cryogenic-cooled InSb-based infrared camera. Our work is the first to validate its post-silicon power characterization estimates. We also trace the source of errors using statistical and frequency domain analyses techniques.

4. To address the difficulties encountered in realistic experimental setups, we provide novel techniques (1) to model the relationship between power and temperature, and (2) to calibrate thermal imaging equipment to compensate for the different material emissivities of semiconductor chips.

The organization of this chapter is organized as follows. In Section 3.2 we elucidate the challenges associated with thermal to power inversion, and in Section 3.3 we describe our proposed methodology for thermal to power inversion that handles the outlined challenges. In Section 3.4 we describe our proposed test chip design and modeling techniques.

In Section 3.5 we provide our experimental and validation results. Finally, Section 3.6 summarizes the main conclusions of this chapter.

3.2 Challenges in Thermal to Power Inversion

An inversion problem is *well-posed* if it satisfies three conditions (first outlined by Hadamard [8]): *existence*, *uniqueness*, and *stability*. In the context of thermal to power inversion, existence means that for every measured temperature map, there exists a power map that leads to the temperature map; uniqueness means that there exists one and only power map that leads to the temperature map; and stability means that small perturbations in the temperature measurements lead to small perturbations in the estimated power maps. In the post-silicon power characterization, there are two main challenges that can lead to ill-posed thermal to power inversion problems. The first challenge arises from the physics of heat diffusion and the second challenge arises from the physics of noise and the technology used to perform thermal imaging. Previous approaches used least-square estimation to find $\mathbf{p}$ that gives temperatures as closest as possible (in ℓ_2 norm) to $\mathbf{T}$ [33]. We argue that previous approaches ignore 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. In our proposed framework, we identify and address such challenges as follows.

3.2.1 First Challenge: Spatial Filtering

The first limit in thermal to power inversion is inherent to 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 density map [27, 38]. While the power density 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 governed by Equation (2.5). In this subsection we will consider a 2-D die with homogenous material, i.e., $k(x, y, z) = k$, to simplify the discussions and to focus on the key concepts. In this case, Equation (2.5) simplifies to

$$k(\frac{\partial^2 T(x,y)}{\partial x^2} + \frac{\partial^2 T(x,y)}{\partial y^2}) = -p(x,y). \tag{3.1}$$

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

$$F_T(u,v) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} T(x,y)e^{-j2\pi(ux+vy)}dxdy, \tag{3.2}$$

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 (3.1) gives

$$k(-u^2 F_T(u,v) - v^2 F_T(u,v)) = -F_p(u,v). \tag{3.3}$$

Thus, the power-temperature transfer function, $G(u, v)$, is equal to

$$G(u,v) = \frac{F_T(u,v)}{F_p(u,v)} = \frac{1}{k(u^2+v^2)} \tag{3.4}$$

It is clear from Equation (3.4) that the higher frequency components of spatial power maps will be subjected to greater attenuation. This spatial filtering phenomenon is illus-

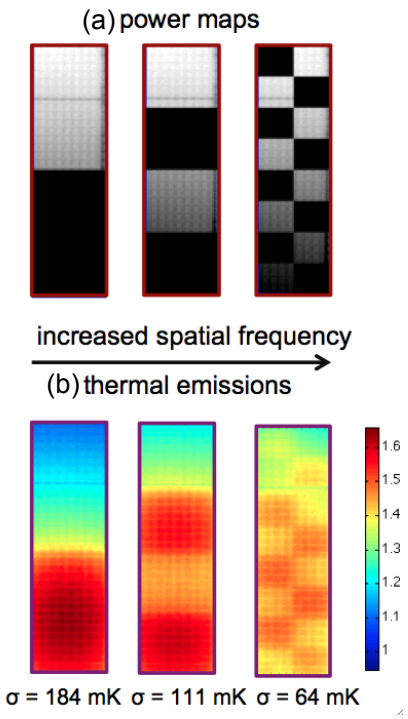

Figure 3.1: Illustrating the impact of spatial low pass filtering on the intensity and variations of thermal emissions.

trated in Figure 3.1, where we create checkerboard power maps with increasing spatial frequency in a test chip and measure their emissions. The maps in Figure 3.1.a have the same amount of total power but they differ in their spatial frequencies. Figure 3.1.b gives the resultant emissions demonstrating that the variations are attenuated as the spatial frequency increases. For example, the standard deviation drops from 184 mK to 111 mK and 64 mK as the spatial frequency is increased. In a simulation-based environment with double precision floating point numbers, attenuation is not an issue, but in real systems with physical and technological limits on their detectors and analog to digital converters, attenuation is a major problem as it degrades the signal to noise ratio. This degradation can attenuate the signal to a level below the detection sensitivity of the infrared imaging equipment, leading to irreversible loss of information. In the discrete form of the heat diffusion equation (Equation (2.6)), the low-pass filtering effect implies that there exists at least one power vector $\mathbf{p_g}$ that belongs to the null space of $\mathbf{R}$. If vector $\mathbf{p}$ satisfies Equation

(2.6), then $\mathbf{p} + \mathbf{p_g}$ also satisfies the equation. In practical terms, $\mathbf{p_g}$ is in the null space of $\mathbf{R}$, if $\mathbf{Rp_g}$ is below the detection sensitivity of the imaging system. Thus, the low-pass filtering effect could render the temperature to power inversion problem ill-posed as the uniqueness condition is violated.

3.2.2 Second Challenge: Noise in Measurements

The second limit arises from discretization and measurement noise introduced during infrared imaging. A number of noise phenomena could lead to errors in the measurements [86]. Sources of noise include (i) dark noise caused by random generation of electron-hole pairs in quantum detectors; (ii) thermal noise caused by the agitation of charge carrier in the electronic readout circuitry; (iii) flicker noise which is inversely proportional to the emission frequency; and (iv) discretization errors introduced by the analog to digital converters of the imaging system. Mathematically, the impact of these errors can be expressed as

$$\mathbf{Rp} + \mathbf{e} = \mathbf{T} + \mathbf{e} = \mathbf{T_m}, \tag{3.5}$$

where the vector $\mathbf{e}$ denotes the errors in measurements introduced during imaging, and the vector $\mathbf{T_m}$ denotes the measured temperatures. To understand the impact of measurement errors on the inversion problem, we will use the Singular Value Decomposition (SVD). SVD decomposes a matrix into a weighted sum of ordered matrices. That is

$$\mathbf{R} = \sum_i s_i \mathbf{u}_i \mathbf{v_i}^T, \text{ and } \mathbf{R}^{-1} = \sum_i \frac{1}{s_i} \mathbf{v}_i \mathbf{u_i}^T \tag{3.6}$$

where the s_i's are the singular values, the $\mathbf{u_i}$'s form a set of orthonormal vectors, and the $\mathbf{v_i}$'s form a set of orthonormal vectors such that $\mathbf{Ru_i} = s_i\mathbf{v_i}$. The operator ^{T} denotes the transpose operation. Using the non-zero singular values of the SVD of $\mathbf{R}$ and Equation (3.5), we find that

$$\begin{aligned} \text{estimated } \mathbf{p} &= \mathbf{R}^{-1}\mathbf{T_m} = \mathbf{R}^{-1}(\mathbf{T} + \mathbf{e}) \\ &= \mathbf{p} + \sum_i \frac{1}{s_i}(\mathbf{u_i}^T\mathbf{e})\mathbf{v}_i \end{aligned} \tag{3.7}$$

Thus, small singular values amplify the impact of noise during inversion [8]. Using Equation (3.7), we can bound the error in power estimation, $\delta\mathbf{p}$, as follows. Since $\delta\mathbf{p} = \mathbf{R}^{-1}\mathbf{e}$, $||\delta\mathbf{p}||_2 \leq ||\mathbf{R}^{-1}||_2||\mathbf{e}||_2$, where $||\cdot||_2$ is the ℓ_2 norm. Thus,

$$\begin{aligned} \frac{||\delta\mathbf{p}||_2}{||\mathbf{p}||_2} &\leq \frac{||\mathbf{R}^{-1}||_2||\mathbf{e}||_2}{||\mathbf{p}||_2} = \frac{||\mathbf{R}^{-1}||_2||\mathbf{R}||_2||\mathbf{e}||_2}{||\mathbf{p}||_2||\mathbf{R}||_2} \\ &\leq \frac{||\mathbf{R}^{-1}||_2||\mathbf{R}||_2||\mathbf{e}||_2}{||\mathbf{T}||_2} \\ &\leq \frac{s_{\max}}{s_{\min}}\frac{||\mathbf{e}||_2}{||\mathbf{T}||_2}, \end{aligned} \tag{3.8}$$

where the ratio $s_{\max}/s_{\min}$ is the *condition number* of $\mathbf{R}$, and is equal to $||\mathbf{R}^{-1}||_2||\mathbf{R}||_2$ [95]. The condition number controls the propagation of errors from the measurements to the estimated power. For example, if the condition number of our chip's $\mathbf{R}$ matrix is in the order of 10^3 then the inversion algorithm could transform a tiny milli-Kelvin error – a typical number in cryogenic thermal quantum detectors – in temperature measurement into a significant error in power estimation. This amplification of noise leads to ill-posed

inversion as the stability condition is violated, where slight changes in the measured temperatures lead to the large changes in the estimated power.

3.3 Proposed Methodology for Thermal to Power Inversion

The objective of thermal to power inversion is to find the best power map $\mathbf{p}$ that minimizes the total squared error between the true temperatures $\mathbf{T}$ as computed using Equation (2.6) and the measured temperatures; that is, $\min ||\mathbf{Rp} - \mathbf{T}||_2^2$. As mentioned in Section 3.2, this objective is challenging to achieve because (1) spatial filtering could lead to ill-posed inversion with a singular or ill-conditioned matrix $\mathbf{R}$, and (2) measurement noise can deviate $\mathbf{T_m}$ from the true temperature $\mathbf{T}$. To address these two challenges, we propose two techniques. We introduce *regularization theory* techniques to improve the numerical stability of $\mathbf{R}$ and to reduce the impact of measurement noise, and we introduce *constraints* on the solution space to reduce the possibility of getting a wrong solution due to the existence of a null space for $\mathbf{R}$. We explain the details of these two techniques in the rest of this section.

To reduce the impact of measurement noise and the potential ill-posed nature of the problem, we propose techniques from *regularization theory* [8]. Regularization techniques consider a family of approximate solutions using 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. *Tikhonov regularization* finds the power solution that gives the least total squared error while simultaneously minimizing the ℓ_2 norm; that is,

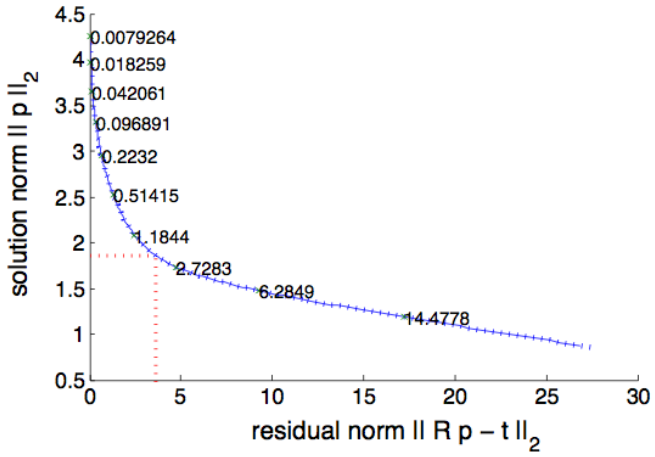

Figure 3.2: Tradeoff between $||\mathbf{Rp} - \mathbf{T}||$ and $||\mathbf{p}||$ as a function of the regularization parameter α.

$$\mathbf{p}_\alpha = \arg_{\mathbf{p}} \min ||\mathbf{Rp} - \mathbf{T_m}||_2^2 + \alpha ||\mathbf{p}||_2^2, \tag{3.9}$$

where $\alpha > 0$ is the regularization parameter that controls the minimization emphasis between the two terms of the objective function [34]. The first term $||\mathbf{Rp} - \mathbf{T_m}||_2^2$, which gives the total squared error, controls how well the power estimates $\mathbf{p}$, lead to temperatures that match the measurements (which could be noisy). If the value of this term is large then the solution is far from the true power, but a small value for this term could lead to a fitting that is driven by noise. The second term $||\mathbf{p}||_2^2$ controls the regularity of the solution. Large values for this term could lead to solutions that are dominated by high-frequency measurement noise. Using one of our captured thermal traces, the trade-off between the two terms as a function of α is illustrated in Figure 3.2. Increasing the value of α traces the curve from the top-left corner to the lower-right corner, and thus the trade-off curve is typically referred to as the *L-curve* [34]. Small values for α can lead to solutions dominated by noise, while large values for α leads to more regularized solutions.

To develop further insight into the role of the regularization parameter α, we utilize the

SVD, but it is first necessary to re-cast the objective of Equation (3.9) as a least squares problem as follows

$$\mathbf{p}_\alpha = \arg_{\mathbf{p}} \min \left\| \begin{pmatrix} \mathbf{R} \\ \sqrt{\alpha}\mathbf{I} \end{pmatrix} \mathbf{p} - \begin{pmatrix} \mathbf{T_m} \\ 0 \end{pmatrix} \right\|_2^2 \tag{3.10}$$

Thus, the solution to the least square estimation problem is

$$\begin{aligned} \begin{pmatrix} \mathbf{R} \\ \sqrt{\alpha}\mathbf{I} \end{pmatrix}^T \begin{pmatrix} \mathbf{R} \\ \sqrt{\alpha}\mathbf{I} \end{pmatrix} \mathbf{p}_\alpha &= \begin{pmatrix} \mathbf{R} \\ \sqrt{\alpha}\mathbf{I} \end{pmatrix}^T \begin{pmatrix} \mathbf{T_m} \\ 0 \end{pmatrix} \\ \mathbf{p}_\alpha &= (\mathbf{R}^T\mathbf{R} + \alpha\mathbf{I})^{-1}\mathbf{R}^T\mathbf{T_m} \end{aligned} \tag{3.11}$$

The SVD expansion of Equation (3.6) can be written as $\mathbf{R} = \mathbf{USV}^T$, where $\mathbf{U}$ and $\mathbf{V}$ are unitary matrices formed from the $\mathbf{u_i}$ and $\mathbf{v_i}$ vectors respectively, and $\mathbf{S}$ is a diagonal matrix with the diagonal elements are the non-zero singular values s_i. Using the SVD, Equation (3.11) can be further analyzed as follows

$$\begin{aligned} \mathbf{p}_\alpha &= (\mathbf{VS}^T\mathbf{U}^T\mathbf{USV}^T + \alpha\mathbf{VIV}^T)^{-1}\mathbf{VS}^T\mathbf{U}^T\mathbf{T_m} \\ &= \mathbf{V}(\mathbf{S}^T\mathbf{S} + \alpha\mathbf{I})^{-1}\mathbf{S}^T\mathbf{U}^T\mathbf{T_m} \\ &= \mathbf{V}\text{diag}(\frac{s_i^2}{s_i^2 + \alpha} \times \frac{1}{s_i})\mathbf{U}^T\mathbf{T_m} \\ &= \sum_i \frac{1}{s_i}\frac{s_i^2}{s_i^2 + \alpha}\mathbf{v}_i{\mathbf{u_i}}^T\mathbf{T_m}. \end{aligned} \tag{3.12}$$

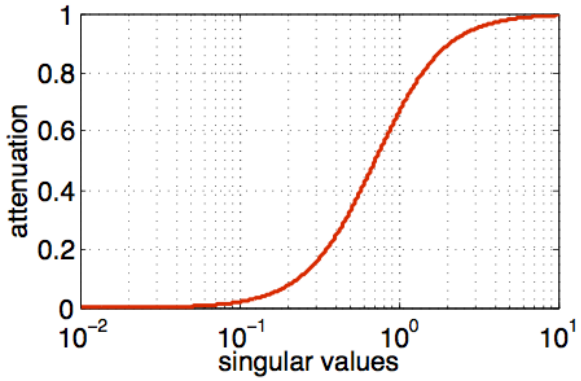

Figure 3.3: Attenuation of the singular values as a function of σ. The regularization parameter $\alpha = 0.5$.

Comparing Equation (3.12) against Equation (3.7), we find that each singular value in the SVD decomposition of Equation (3.7) is multiplied by a factor $\frac{s_i^2}{s_i^2+\alpha}$ known as *Tikhonov attenuation factor* [98, 34]. Figure 3.3 gives the value of the Tikhonov attenuation factor $\frac{s_i^2}{s_i^2+\alpha}$ as a function of the singular value s_i for $\alpha = 0.5$. The figure shows that the attenuation factor essentially functions as a *filter function* that filters out singular components that are small relative to α and passes singular components that are large relative to α [98]. Singular values of zero value are totally eliminated. The attenuation of the small singular components makes the inverse problem more well conditioned and controls the error propagation as governed by Equation (3.8). Good values for α can be found by inspecting the L-curve. One possibility is to the use the *corner* of the curve [34]. The corner is the point on the L-curve with the maximum curvature. In Figure 3.2 the corner occurs at α is equal to 1.6.

To reduce the possibility of getting a wrong solution due to the existence of a null space for $\mathbf{R}$, we introduce additional constraints on the solution space. One possible constraint is that the sum of the elements of the power map must be equal to the total power consumption of the chip p_{total} and that these elements are nonnegative; i.e.,

$$||\mathbf{p}||_1 = \sum_i p_i = p_{total} \text{ and } \mathbf{p} \geq \mathbf{0}, \tag{3.13}$$

where $||\cdot||_1$ is the ℓ_1 norm and p_{total} is the total power consumption of the chip which could be measured externally using a digital multimeter. Thus, if multiple solutions exist, then the the solution with least total power error is chosen. Note that because the total (and average) power is constrained by Equation (3.13), the regularization term in Equation (3.9) controls the variance in spatial power estimates. In practice, any digital multimeter has a tolerance, tol, in its measurements (the tolerance is typically listed in the multimeter's data sheet), and thus it is better to replace the constraint of Equation (3.13) by the following constraints:

$$\begin{aligned} ||\mathbf{p}||_1 &\leq p_{total} + tol, & (3.14)\\ ||\mathbf{p}||_1 &\geq p_{total} - tol, \text{and} & (3.15)\\ \mathbf{p} &\geq \mathbf{0} & (3.16) \end{aligned}$$

Our overall inversion methodology is summarized in the algorithm given in Figure 3.4.

3.4 Test Chip Design and Modeling

In order to validate our power characterization methodology, we designed an experimental test chip with special emphasis on the ability to estimate the spatial power consumption through two different means. We describe the design and implementation of the test chip

Procedure: Thermal to power inversion method
Input: Thermal map $\mathbf{T_m}$, p_{total}, $\mathbf{R}$
Output: $\mathbf{p}$

1. Given $\mathbf{T_m}$ and $\mathbf{R}$, construct the L-curve
2. Identify α from the corner of the L-curve
3. Compute $\mathbf{R}_\alpha = \sum_i s_i \frac{s_i^2+\alpha}{s_i^2} \mathbf{u}_i \mathbf{v}^T$
4. Solve the quadratic program: $\min ||\mathbf{R}_\alpha \mathbf{p} - \mathbf{T_m}||_2^2$ such that $||\mathbf{p}||_1 \leq p_{total} + tol$, $||\mathbf{p}||_1 \geq p_{total} - tol$ and $\mathbf{p} \geq 0$
5. Return the solution of the quadratic program $\mathbf{p}$

Figure 3.4: Proposed thermal to power inversion methodology.

as follows.

Test Chip Design and Implementation. The basic unit of our circuit is a *programmable micro heater*, which consists of a number of ring oscillators (ROs) that are controlled by flip-flops that determine the operational status of the micro-heater. Figure 3.5.a shows the programmable 15-stage RO with the D Flip-flop (DFF). Each micro-heater block consists of 3 $\times$ 3 ring oscillators, creating the array of micro-heater blocks as shown in Figure

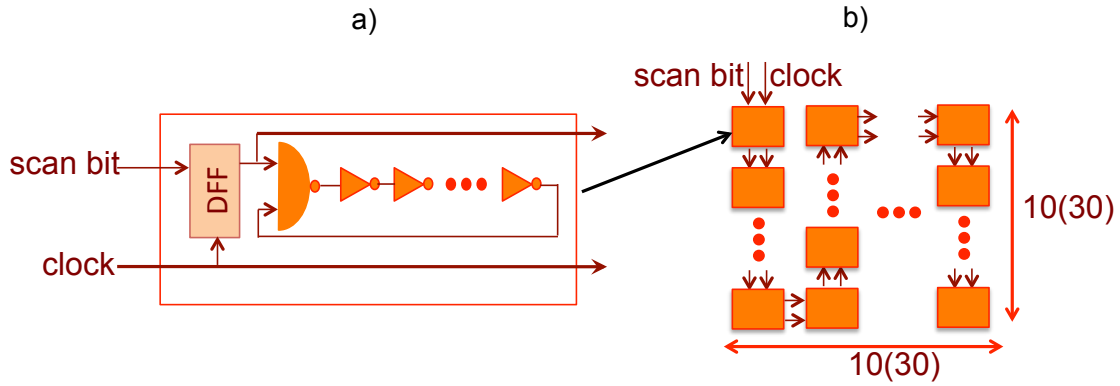

Figure 3.5: a) Programmable ring oscillators, and b) Programmable array of micro-heaters.

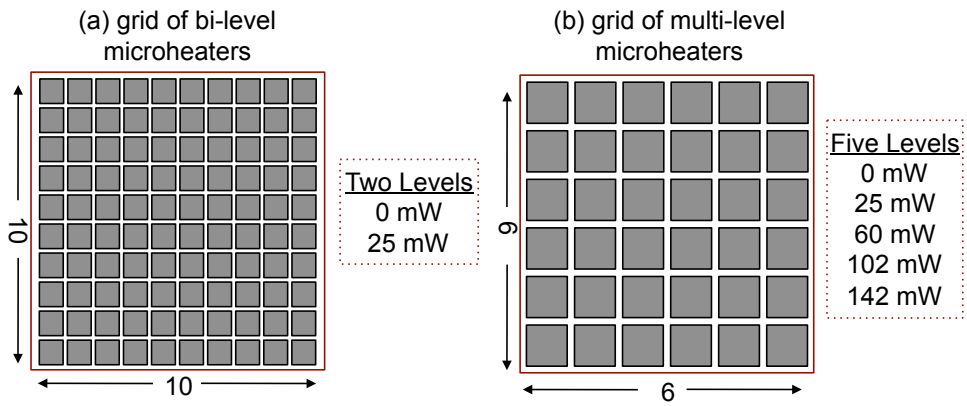

Figure 3.6: Grid of programmable micro-heaters.

3.5.b. 10(30) represents 10 micro-heater blocks, and 30 ring oscillators on one side. So, in total 100 micro-heater blocks consists of 30 $\times$ 30, nine hundred ring oscillators. We create two kinds of micro heater designs:

1. *Bi-Level Micro-Heaters:* A bi-level micro heater consists of 3 $\times$ 3 ROs resulting into nine 15-stage ROs together with one flip-flop that controls their operational status. If the DFF holds a binary value of 1 then the heater is turned on; otherwise, it is turned off. When enabled, each micro heater consumes 25 mW. Using the programmable heater, a grid that consists of 10 $\times$ 10 micro heaters is created as shown in Figure 3.6.a. In the grid structure, the output of each DFF is connected to the input of the DFF of the consecutive heater forming a scan chain.

2. *Multi-Level Micro-Heaters:* A multi-level micro heater consists of ROs that can be programmed to one of the following configurations: nine 51-stage ROs with a power consumption of 25 mW; eighteen 25-stage ROs with a power consumption of 60 mW; twenty-seven 19-stage ROs with a power consumption of 102 mW; and thirty-six 13-stage ROs with a power consumption of 142 mW. Thus, each micro-heater block offers five different power levels: 0, 25, 60, 102, and 142 mW. The DFFs associated with each micro heater determines its status. Using the programmable

micro heater blocks, a grid that consists of 6 $\times$ 6 blocks is created as shown in Figure 3.6.b. In the grid structure, the output of each DFF is connected to the input of the DFF of the consecutive heater forming a scan chain.[1]

In both designs, the advancements of the programming bits in the chain is controlled by the clock signal. 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 experimental 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 map on a real chip. Previous works used processors in which independent control of various processor blocks is infeasible. The programmable nature of the grid enables the generation of a large number of different maps that can be used to test the accuracy of the thermal to power inversion methodology.

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 total power consumption of the grid and dividing it by the number of enabled micro heaters. The locations of the enabled heaters are known by construction. Hence, we are able to construct the spatial power map through an alternative path to infrared emissions. Previous works lacked this ability to validate their spatial power estimations through different means.

[1]The multi-level micro-heater implementation in the FPGA was done by the co-author Stefan Angelevski.

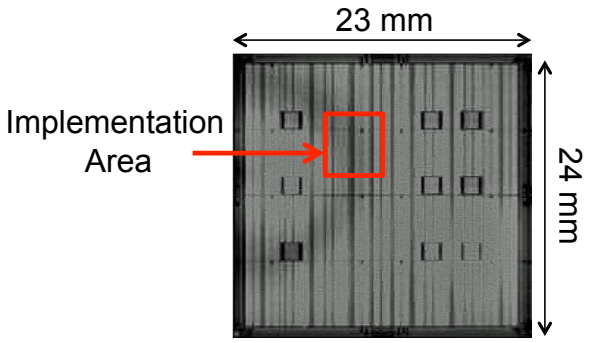

Figure 3.7: 90 nm Altera Stratix II FPGA .

For implementation, we choose a 90 nm Altera Stratix II (EP2S180) field programmable gate array (FPGA) with 180,000 logic elements with total die dimensions of 23 mm $\times$ 24 mm as shown in Figure 3.7. The regular fabric of the FPGA ideally fits our design. For our experiments, we use a relatively homogenous section of the die that spans 7.2 mm $\times$ 7.9 mm for the bi-level 10$\times$10 grid as shown in Figure 3.8.a, and a section that spans 8.8 mm $\times$ 8.9 mm for the multi-level 6$\times$6 grid as shown in Figure 3.8.b. The micro heater blocks are mapped to the logic array blocks at the precise grid locations using Altera's Quartus II placement assignment editor. In order to capture the chip's thermal emissions it was necessary to remove the heat spreader. While removing the heat spreader is going

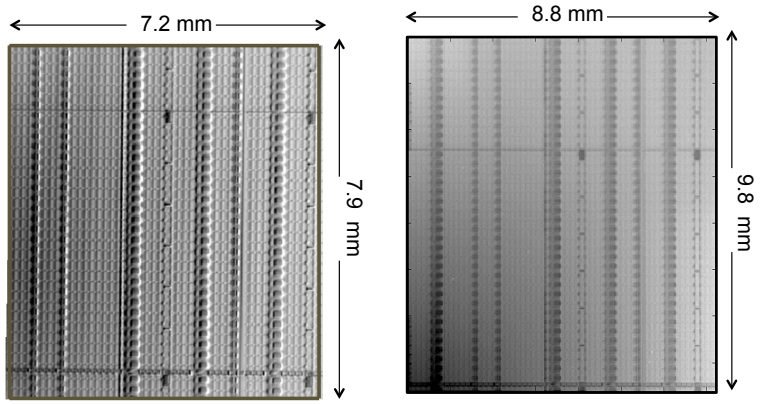

Figure 3.8: Implementation areas in Altera Stratix II EP2S180 device.

to change the spatial thermal behavior, the change in spatial thermal emissions does not change the underlying dynamic power consumption (fCV^2) [85], which is weakly dependent on temperature. 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 in the next paragraph.

Measuring Modeling Matrix ($\mathbf{R}$) of the Test Chip. Given the test chip, it is necessary to measure the modeling matrix $\mathbf{R}$. The matrix $\mathbf{R}$ can be measured in a column-by-column basis as follows. Enabling only one micro heater at a time is mathematically equivalent to setting the vector $\mathbf{p}$ to be equal to $[0 \cdots p_k \cdots 0]^T$, where p_k is the power consumption of the kth enabled block and all the other elements are zeros. For each micro heater location, we enable the selected micro heater as shown in Figure 3.9.a and record the emitted temperatures across the field as shown in Figure 3.9.b. If $\mathbf{T_k}$ denotes the thermal emissions captured from enabling the k^{th} micro-heater, then column k of matrix $\mathbf{R}$ is equal to $\mathbf{T_k}/p_k$. We automate the whole process in order to measure the columns of $\mathbf{R}$ with fast turn-around time.

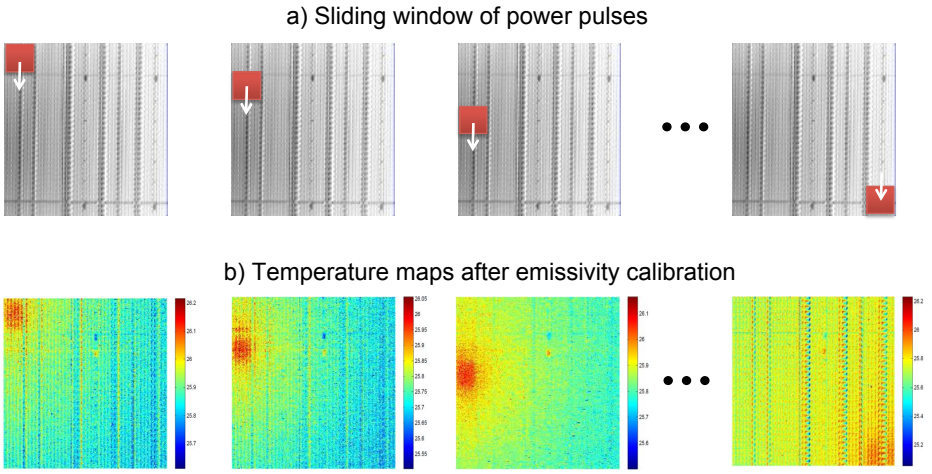

Figure 3.9: Measuring the power to thermal inversion matrix $\mathbf{R}$ using a sliding window of power pulses.

We utilized the programmability of FPGAs to estimate the model matrix $\mathbf{R}$. In custom chips with generic designs, the matrix $\mathbf{R}$ can be estimated in the same conceptual way but through a different implementation approach [33, 61]. One approach is to turn off the chip, and scan a laser beam with known power density to deliver the power from the outside to the regions of interest. The scanning of the laser system can be automated by using a pair of galvo-directing mirrors [33]. Our method uses the programmable nature of our design to get the same results of the expensive laser scanning system but in a much cheaper way. Another approach is to use the actual design and layout of the chip to conduct a fluid dynamic simulation coupled with a heat diffusion simulation to estimate the matrix $\mathbf{R}$ [33, 60].

There is always a possibility that errors might occur during the estimation of $\mathbf{R}$. If $\mathbf{R}$ is estimated from direct measurement, then measurement noise can introduce errors, and if $\mathbf{R}$ is estimated from simulations, then unrealistic simulation assumptions can lead to errors. Our experimental results in the next section show that the overall power estimation error arising from our inversion procedure is relatively small.

3.5 Validation and Experimental Results

To test our post-silicon power characterization methodology, we put together the following experimental setup which is shown in Figure 3.10. To capture the thermal emissions from the back-side of our die, we use a FLIR SC5600 infrared camera with a mid-wave spectral range of 2.5 μm – 5.1 μm. The camera is run at a rate of 100 Hz with a spatial resolution of 30 μm with a 0.5$\times$ microscopy kit. Silicon is transparent to infrared emissions in the 1.5 μm – 5.1 μm range with a 55% transmittance. Thus, the infrared camera can measure the temperature through the chip under test [60]. The camera's cryogenic InSb detectors

cooled to 77 K (-196 °C) have a sensitivity of about 15 mK noise equivalent temperature difference (NETD). We use an Agilent E3634A power supply to supply and measure the total power consumption of FPGA chip.

Thermal Calibration. Measuring the temperature is slightly complicated because 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 measurements 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, which is defined as the ratio of that material's thermal emission to that of a perfect black-body at the same temperature [86]. 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 held at isothermal temperature by external means.

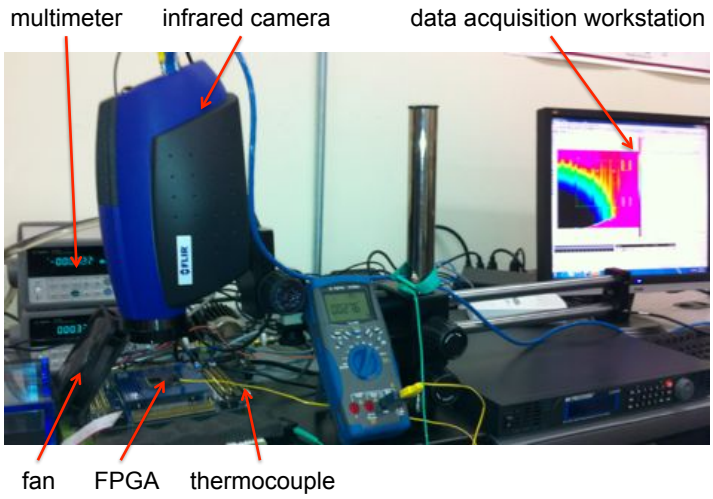

Figure 3.10: Experimental setup showing the different equipment that make up our measurement system.

To handle the emissivity problem, previous approaches coated the backside of the die with a material of constant emissivity [33]. The downside of this approach is that it obstructs the silicon transparency and only measures the heat emissions resulting from the projections of the internal emissions on the backside of the die. Previous approaches were thus forced to thin the backside silicon to reduce the amount of internal projections and spatial heat diffusion. To get a more accurate thermal imagery, we avoid coating the backside of the silicon and instead devise a pixel-by-pixel calibration process that translates the captured emissions (as measured by the digital levels of the camera's A/D converters) to temperatures.

The relationship between the digital level D_i and temperature T_i of a pixel i can be modeled by an exponential function

$$D_i = \alpha_i e^{\beta_i T_i}, \tag{3.17}$$

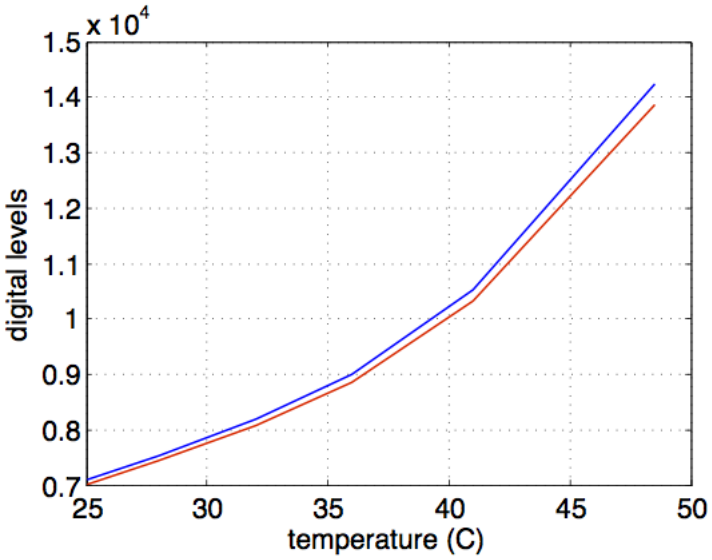

Figure 3.11: Relationship between camera's digital levels and temperatures for two pixels.

where α_i and β_i are per-pixel coefficients that are functions of many factors including emissivity, path to length, and integration time [36]. The exponential relationship arises from the physics of photon detectors in which the current of an infrared-sensitive diode depends exponentially on the incident radiation. If there were no emissivity differences between the different pixels then α_i and β_i would be chip-wide constants, but instead they must be computed for each pixel. For example, Figure 3.11 shows two curves relating the temperatures to the measured digital levels for two different pixels on the test chip. Our pixel-by-pixel calibration procedure is simple. We first turn the test chip off and then force it to an isothermal status through external means.

Two thermocouples placed at opposite ends of the die could be used to verify chip-wide steady-state attainment. Once steady-state is reached, the digital levels of all pixels are captured using the camera. This process is repeated for a few temperatures and then the calibration curves (as the ones given in Figure 3.11) are constructed, and the α_i and β_i for every pixel i are found through curve fitting. Figure 3.12 contrasts the raw digital thermal levels before calibration and the thermal images after calibration. The images show that the calibration process successfully removes imaging artifacts introduced by emissivity variation. For high-power chips, thermally-controlled infrared transparent oil can be used to force the isothermal status, and then measure the temperature of the chip

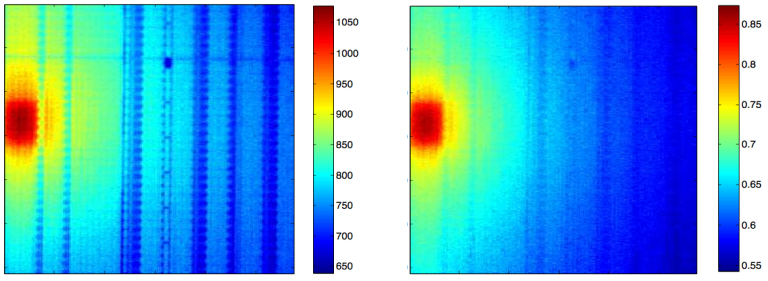

(a) Digital levels without emissivity calibration

(b) Temperature after emissivity calibration

Figure 3.12: Thermal images before and after emissivity calibration.

through available thermal sensors.

We conduct and report the results of four experiments:

1. The first experiment compares the estimation accuracy of our proposed method against previous techniques using the bi-level micro heater block grid. We analyze the sources of errors using statistical and frequency domain techniques.

2. The second experiment assesses the effectiveness of our method using random spatial maps constructed using the bi-level micro heater grid.

3. The third experiment evaluates the accuracy of our power estimation method as a function of the spatial frequency of the power maps.

4. The fourth experiment assesses our power mapping method using the multi-level micro heater block grid. We demonstrate that our method is capable of handing underlying power maps with various intensities and spatial constructions.

Experiment 1. In the first experiment we assess the accuracy of our thermal to power inversion methodology by evaluating its power estimates for a number of reference spatial power maps generated using the bi-level micro heater grid. The locations of the activated micro heaters of these maps are illustrated in Figure 3.13.a, and the measured temperature maps after emissivity calibration are illustrated in Figure 3.13.b. Given our low-power micro heater designs, the temperature range, i.e., the difference between the maximum temperature and the minimum temperature in each map, is about 1.5–2 °C. The power maps estimated by just minimizing total squared error as proposed by previous approaches ([33]) are given in Figure 3.13.c; and the estimated power maps from our methodology are given in Figure 3.13.d. The power mapping results in Figure 3.13.c and Figure 3.13.d are rounded to the nearest power level. We report the *estimation error* in percentage which

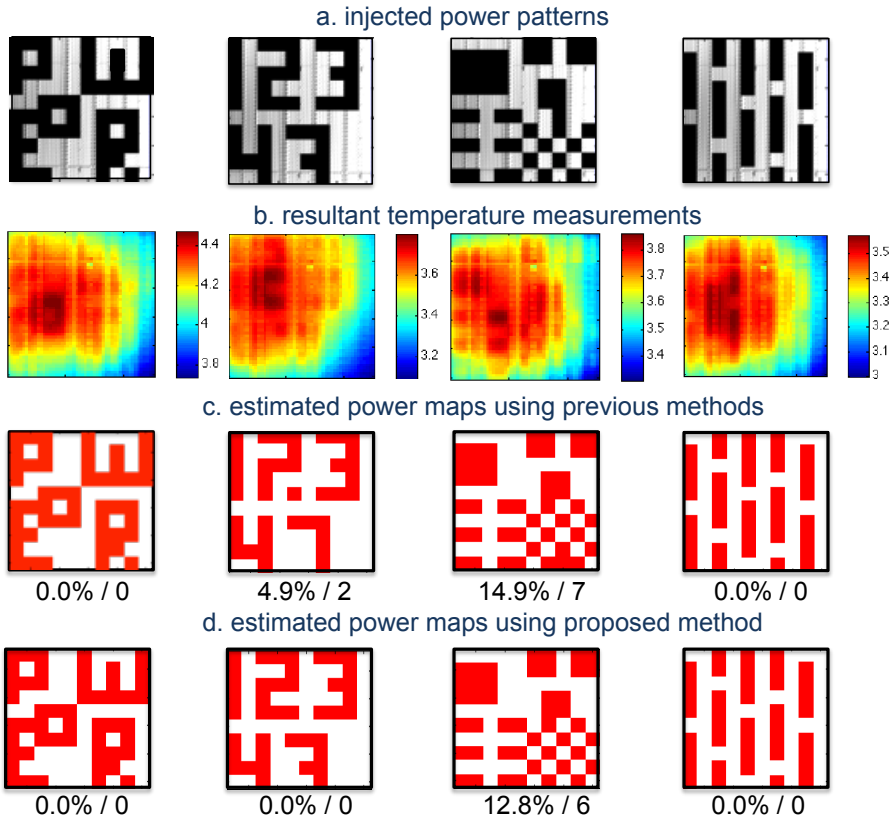

Figure 3.13: Accuracy of estimating arbitrary power maps using thermal emissions. Power maps are rounded to the nearest level.

is equal to the sum of the absolute differences between the power estimates and their true values divided by the total power. The average error from using previous approaches is 4.95% while the average error from using our proposed approach is 1.5%. We also report the number of block heaters that were not estimated correctly (either turned on and estimated to be off or vice versa). Previous techniques give a total of 9 incorrectly estimated blocks, whereas our technique give 6 incorrectly estimated blocks for reduction of 30%. By visually comparing the injected power maps and the estimated maps, we notice that our technique recovers the injected maps to a very good extent.

To understand the source of errors, we conduct the following two analyses. First, we simulate the resultant temperatures if the true power maps are used as inputs; the simula-

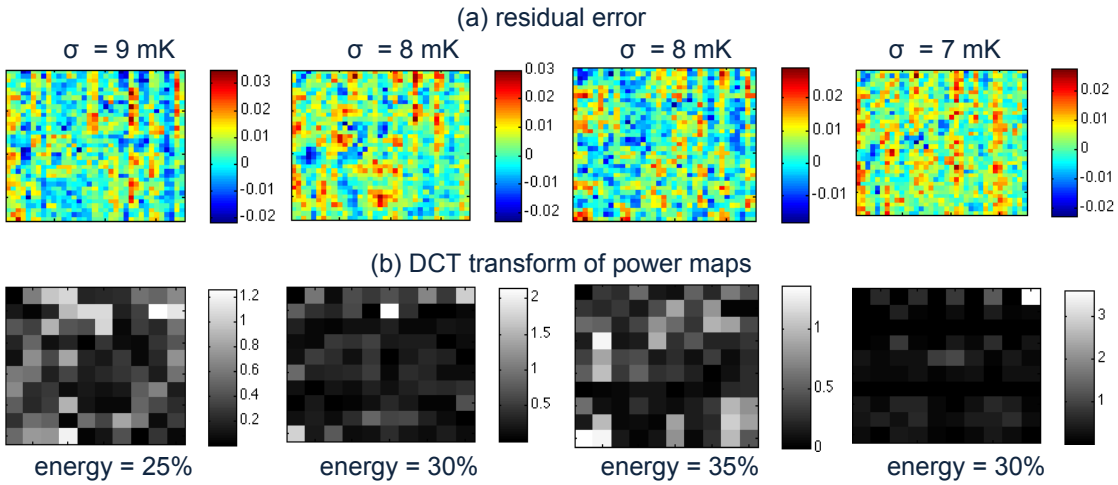

Figure 3.14: Analysis of errors of maps in Figure 3.13.

tion is basically the result of multiplying the injected power maps by the matrix $\mathbf{R}$. We plot in Figure 3.14.a the residual error between the measured temperatures and the simulated measurements. We verify that the errors form a normal distribution using the Kolmogorov-Smirnov test, and we compute the standard deviation for each residual distribution. The standard deviations are given as labels (σ) in Figure 3.14.a. For instance, the first pattern, which has an error of 2.0% in its power map estimates, has a residual standard deviation of 9 mK which means that the large majority of errors (97%) are between ± 18 mK. The residual errors fall within the sensitivity limitation (15 mK) of our camera detectors. The first pattern has the highest standard deviation in residuals which implies that it had the highest amount of noise. Our second analysis technique computes Discrete Cosine Transform (DCT) of the true power maps and report them in Figure 3.14.b (the top-left corner is the lowest frequency). The DCT quantifies the spatial frequencies of an input power map. For each power map, we compute the percentage of the power signal energy that is present in higher frequencies of the DCT.[2] The third pattern, which has 12.8% error in its power map estimates, has the largest amount of energy, 35%, in the high-frequency range. Thus, we attribute the power estimation error to such high-frequency energy content.

[2]The energy of a signal is the sum of squares of its frequency domain coefficients. We compute the energy in coefficients with frequencies ≥ 7 and divide the result by the total energy.

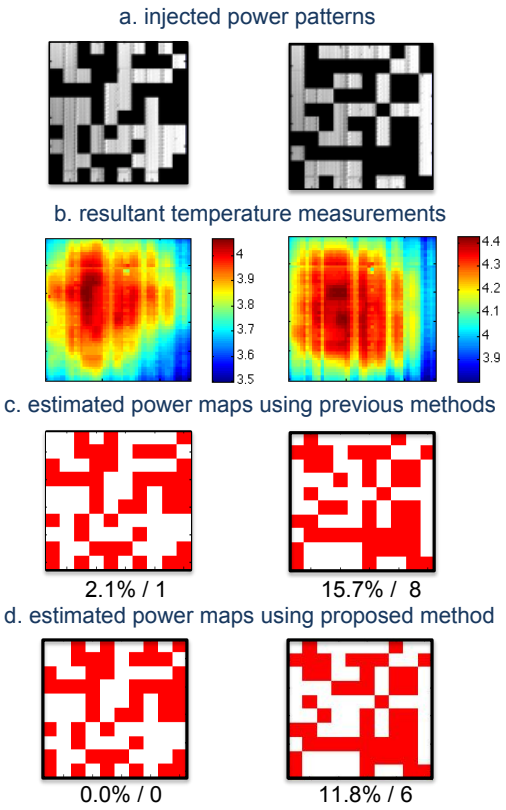

Figure 3.15: Accuracy of estimating spatial power estimates from thermal emissions of random power maps. Power maps are rounded to the nearest level.

Experiment 2. In the second experiment we use random spatial power maps for testing our methodology in addition to the maps of Experiment 1. The use of random maps provides a more complete assessment of our proposed method. In Figure 3.15.a we provide the random maps tested, and the resultant thermal images are provided in Figure 3.15.b. Figure 3.15.c gives the resultant power maps from previous techniques, while Figure 3.15.d gives the results from using our inversion methodology. The results show estimation errors of 0% and 11.8% from our techniques versus 2.1% and 15.7% from previous techniques. We also report the number of block heaters that were not estimated correctly. Previous techniques give a total of 9 incorrectly estimated blocks, whereas our technique give 6 incorrectly estimated blocks for reduction of 30%. Overall, the results of experiments 1 and 2 show an average error of about 4.4% from our techniques versus

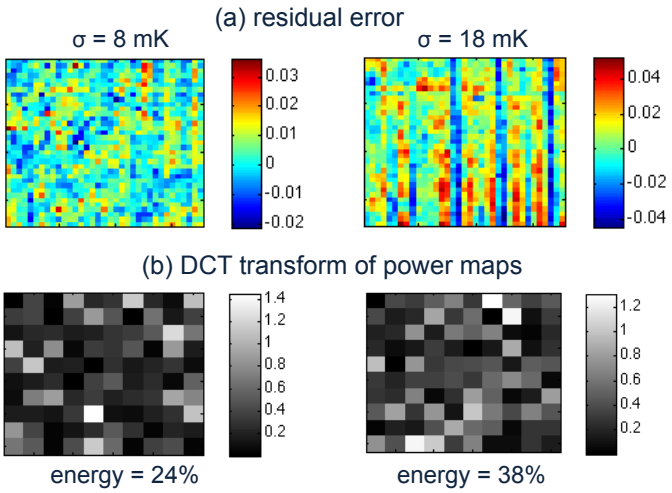

Figure 3.16: Analysis of errors of maps in Figure 3.15.

6.6% from previous techniques, where our technique gives a total of 13 incorrectly estimated blocks and previous techniques give 19 incorrectly estimated blocks. Thus, our technique gives 31% improvement in power mapping.

To further understand the reason for the errors in some of the maps, we conduct analyses similar to the ones of Experiment 1. First, we compute the residual error between the simulated temperatures and the measured temperatures and plot the results in Figure 3.16.a. We verify that the errors form a normal distribution, and we compute the standard deviations for the distributions which are given as labels (σ) in Figure 3.16.a. For instance, the second pattern, which has a large error of 11.8% in its power map estimates, has a residual standard deviation of 18 mK. Second, we compute the Discrete Cosine Transform (DCT) of the true power maps and give them in Figure 3.16.b (the top-left corner is the lowest frequency). Again, the second pattern shows the largest amount of energy, 38%, in its higher-frequency components. Thus, we attribute both the noise and high-spatial frequencies to the large error in power estimation.

Experiment 3. 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 Section 3.2, the nature of heat conduction on chips leads to a low-pass filtering effect. Hence, we create checker-board maps of increasing spatial frequencies as illustrated in Figure 3.17.a, which lead to the thermal emissions illustrated in Figure 3.17.b. The estimated spatial power maps are given in Figure 3.17.c. The average errors are: 0.0%, 0.0%, and 11.5%. To provide further analysis into the source of error, we plot the 2-D DCT of the power maps in Figure 3.18. The figures clearly show the trend of increased frequencies in the power pattern, where the third pattern has the highest frequency and error. The results agree with our earlier discussions in Section 3.2 that concluded that increasing the spatial frequencies of power maps can lead to a deterioration in the accuracy of thermal to power inversion due to the impact of low pass filtering.

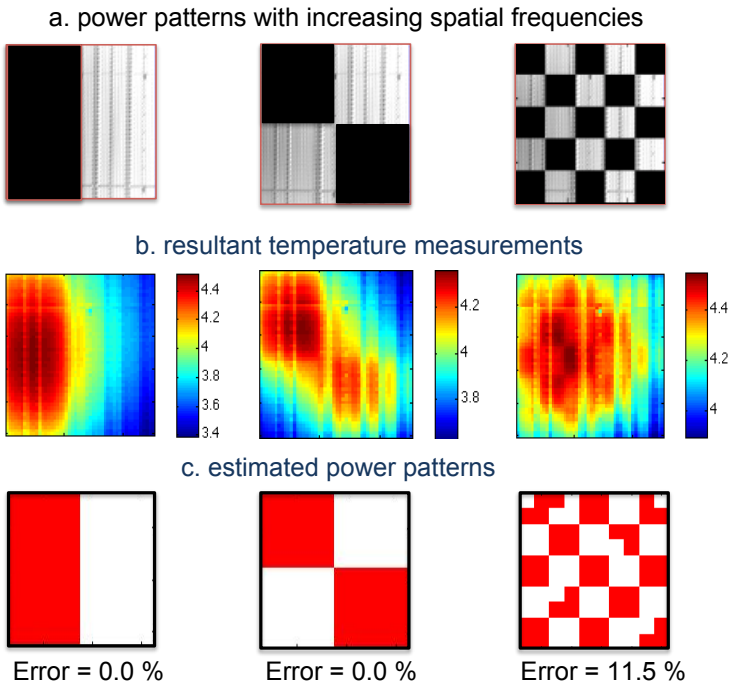

Figure 3.17: Impact of increasing spatial frequencies of power maps on the accuracy of thermal to power inversion. Power maps are rounded to the nearest level.

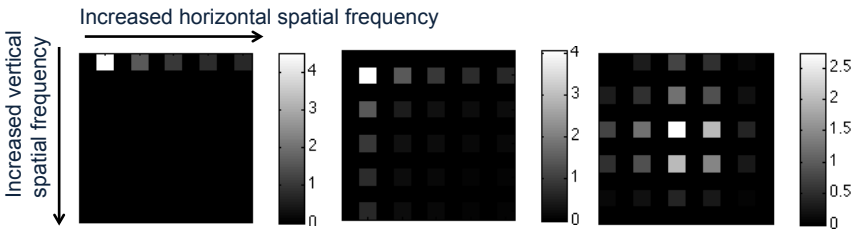

Figure 3.18: DCT of checkerboard maps.

Experiment 4. In this experiment, we provide results using the multi-level micro heater grid design. In contrast to the first three experiments where the power of each micro heater was only restricted to two levels (0 and 25 mW), the micro-heaters in the constructed spatial power maps of this experiment have power levels of multiple intensities (0, 25, 60, 102, and 142 mW). Our inversion procedure can naturally handle any number of levels, and the results of this experiment confirm this capability. The number and step size of different levels are limited by the noise level of the camera. The state-of-the-art infrared camera can detect temperature difference as small as 20 mK, as long as the power sources

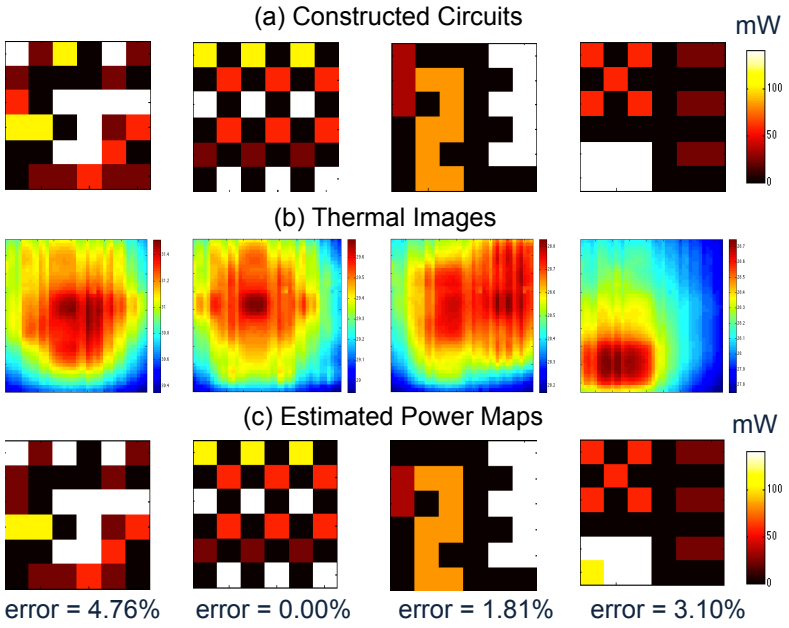

Figure 3.19: Results from multi-level power estimation. Power maps are rounded to the nearest level.

produce temperature difference that can be detected by the infrared camera, it can be used as a different power level. Figure 3.19.a gives the constructed multi-level power maps. The resultant thermal maps are given in Figure 3.19.b, and the estimated power maps from our inversion procedure are given in Figure 3.19.c. The power mapping results are rounded to the nearest power level. The power estimation errors of the four maps are 4.76%, 0.00%, 1.81% and 3.10%, with an average of 2.41%. The magnitudes of the power estimation errors from the multi-level power mapping experiment are comparable to the magnitude of the errors in the first three experiments using bi-level micro heaters. The results of this experiment confirm that our technique is capable of handling a large range of power intensities and spatial power maps.

3.6 Summary

In this chapter we presented a new methodology for spatial post-silicon power characterization using the infrared emissions from the back of the silicon die. We elucidated the various challenges that underlie thermal to power inversion. We demonstrated mathematically and experimentally how low-pass filtering, discretization, and measurement errors could all compromise the accuracy of power estimation. We proposed new techniques from regularization theory to reduce the impact of noise and improve the numerical instability in the model matrix $\mathbf{R}$ by eliminating zero singular values and filter small singular values. We also introduced constraints on the solution space to eliminate the possibility of getting multiple solutions. Furthermore, we provided experimental techniques to compensate for the varying emissivity of different chip materials and to measure the thermal resistance model matrix. For experimental validation of our proposed technique, we designed a highly modular programmable test chip to create sets of known power maps. Our test chip and realistic infrastructure enabled us to validate our methodology by comparing

its power estimates against the known injected spatial power maps. Compared to previous approaches, our experiments demonstrate consistent improvement in power estimation accuracy where we improve the power mapping accuracy by 30%. We also analyzed the residual errors between the simulated temperatures and the measured temperatures to understand the error sources. Our statistical and frequency domain analyses quantify the roles of noise and spatial filtering on the accuracy of power estimates.

Chapter 4

Fundamentals of Post-silicon AC-based Power Mapping

4.1 Introduction

In Chapter 3, we have proposed a detailed technique for DC-based post-silicon power mapping, and demonstrated the accuracy our technique using a programmable test chip. As elucidated DC-based power maping faces two major challenges: (1) spatial low-pass filtering of the underlying power map arising from heat diffusion can cause information loss, and (2) the measurement noise in the setup of the infrared imaging system limits the obtainable accuracy. In this chapter, we present a post-silicon power mapping methodology, where AC excitation signals are used instead of DC excitation signals to improve mapping accuracy. We present analytical derivations and experimental validations, demonstrating that the proposed AC-based approach leads to significant improvements in the power mapping compared to DC-based methods as discussed in Chapter 3. The major

contributions of this work are as follows.

1. We analyze and quantify the role of noise on the quality of thermal imaging. We breakdown the noise contribution into flicker and white noise, and analyze their dependencies on integration time and excitation frequencies.

2. We analyze the impact of AC excitation frequency on the extent of spatial heat diffusion and the strength of the thermal signal. We combine this analysis with the noise analysis to elucidate the relationship between the signal-to-noise (SNR) ratio, the excitation frequency, and the integration time.

3. We quantify the thermal to power model parameters of real chips at desired AC frequencies. Our method captures all modes of heat transfer.

4. Given the AC thermal emissions and model parameters, we devise a constrained convex optimization inversion procedure to estimate the post-silicon power maps. In our procedure additional constraints obtained from lumped physical measurements, e.g., total power, are imposed on the solution space to reduce numerical error.

5. To scientifically validate our post-silicon power mapping approach, we utilize the programmable test chip described in Chapter 3. Due to the programmable nature of the design, it is possible to precisely control the AC excitation frequency of the power sources. The test chip has micro-heater blocks which can be switched ON and OFF at desired frequencies to create different excitation frequencies. Thermal emissions from the test chip are captured using infrared imagery and then processed to reveal the estimated post-silicon power maps. We quantify the impact of excitation frequency on the accuracy of post-silicon power maps and relate this accuracy to the signal-to-noise ratio.

The rest of this chapter is organized as follows. Section 4.2 provides the motivation and necessary background information on AC-based thermography. In Section 4.3, we analyze the impact of excitation frequency and integration time on flicker noise and white noise. In Section 4.4, we analyze the impact of excitation frequency on spatial heat diffusion and thermal signal strength. We combine the analyses of Section 4.3 and Section 4.4 to analyze the SNR trends in Section 4.5. We present our numerical inversion optimization method in Section 4.6. We provide an extensive set of experimental results on a test chip in Section 4.7. Finally, Section 4.8 summarizes our main results.

4.2 Motivation and Background for AC Thermography

The are two major challenges in post-silicon power mapping: i) spatial heat diffusion [33, 27, 38, 17], and ii) measurement noise in the thermal imaging system [17, 10]. Heat diffusion blurs the underlying power map and reduces the accuracy of post-silicon power maps as it filters out the spatial high-frequency power patterns. AC excitation reduces the amount of spatial heat diffusion as the AC excitation frequency increases [55, 10]. As a result, AC-based thermography improves the resolution of thermal images and the detection of weak emission sources, which makes it a valuable tool in device characterization [41] and failure analysis of integrated circuits [10].

In the power mapping procedure described in Chapter 3, a DC excitation source, e.g., a workload of a stable nature or a test pattern, is applied to the chip under characterization, and the infrared emissions are captured from the back of the die using infrared imaging. The emissions are then inverted to get the power maps. Figure 4.1.a illustrates the main framework of DC post-silicon power mapping [33, 60, 78, 17, 71]. The general framework for AC thermography is given in Figure 4.1.b, where an AC excitation source rather than

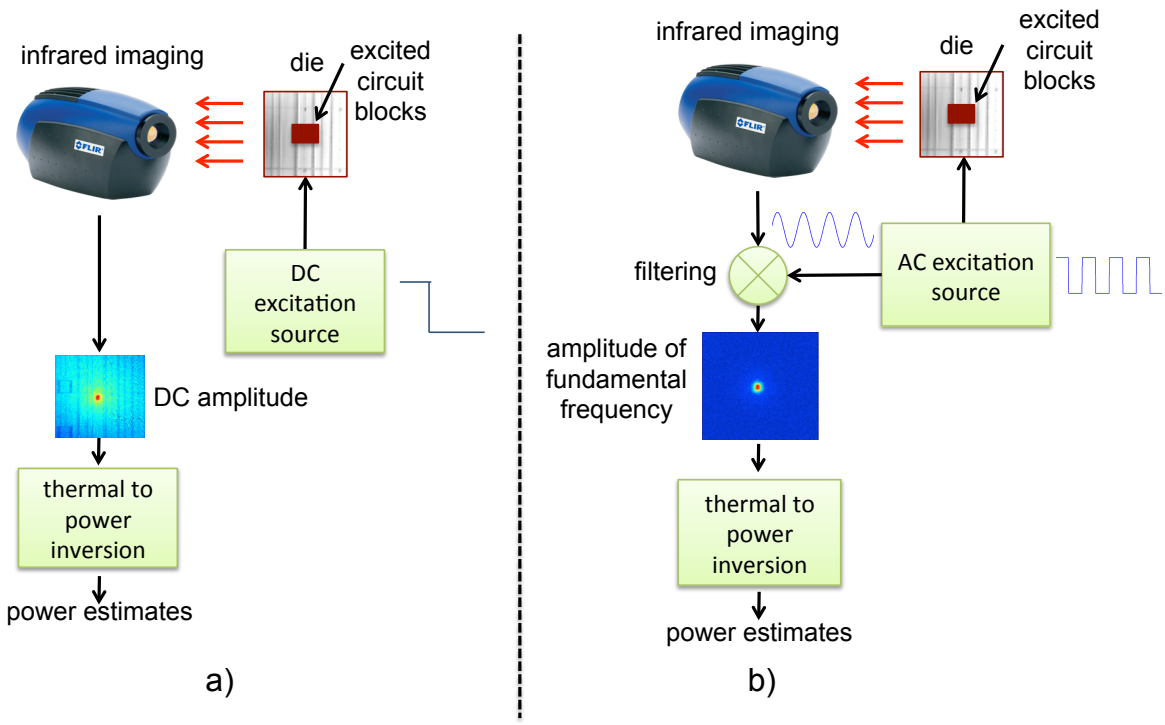

Figure 4.1: a) DC post-silicon power mapping framework, and b) AC post-silicon power mapping based on thermography techniques using infrared emissions.

a DC excitation is applied to the integrated circuit. Applying a true sinusoidal AC source to excite a digital circuit is impossible. Instead a square wave is applied, and because a square wave can be represented by a Fourier series, whose dominant component is the fundamental frequency, such technique does not alter the results as long as the acquired infrared emissions are filtered to only extract the fundamental frequency [10]. Creating AC square-wave excitations in digital circuits can be implemented by a number of techniques such as: (1) toggling enable signals of circuit blocks while keeping the operating voltage constant; (2) alternating the voltage supply signal between two operational values (e.g., 0.9 V and 1 V); or by (3) executing workloads (for the case of processors) that alternate between an activity phase and inactivity phase. For example, Figure 4.2.a shows the thermal emissions of just one pixel over time arising from a test circuit where a circuit block is toggled at 2 Hz, and Figure 4.2.b shows the Fourier spectrum of the waveform, clearly showing the amplitude at the fundamental frequency. The amplitudes of all pixels

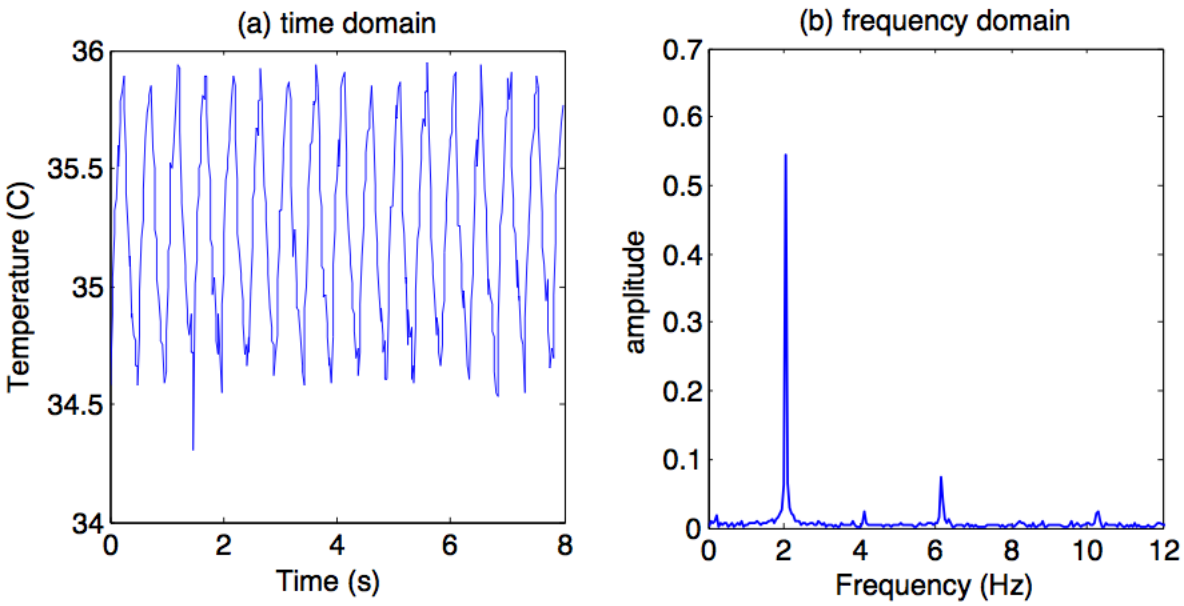

Figure 4.2: Steady-state AC emissions at 2 Hz in time domain and frequency domain for one pixel.

at the fundamental frequency form the AC-based thermal image as shown in Figure 4.1.b.

In the context of fault isolation, Breitenstein [10] discusses methods to use AC lock-in techniques to deblur thermal images by deconvolution to reveal sharper power images that identify a failure site with excessive leakage power consumption. However, no practical way of measuring the thermal to power modeling parameters are presented and instead the parameters are derived from analytical derivations that do not capture all modes of heat transfer (e.g., convection at surface and radiation) in a real chip. Furthermore, the reported experimental setup and results are not devised to quantify intricate power maps, but rather to isolate faults where only a handful of locations at most are actively consuming power. In contrast to previous usages of AC-based thermography, we use AC-based thermography for the purpose of post-silicon power mapping where intricate spatial power maps produced from tens of circuit blocks can be simultaneously estimated from their combined thermal emissions.

4.3 Impact of Using AC Excitation on Noise

The objective of this section is to analyze and quantify the impact of noise on AC-based thermography. Previous works assumed white noise on the measurements [10]. We will show in this section that $\frac{1}{f}$ flicker noise has a significant contribution to the measurement noise.

Let $T_i(t)$[1] denote the temperature of pixel i at time t as recorded by the thermal imaging system, and let P denote the integration period of the measurements. Then the *temperature magnitude*, T_i, of pixel i is given by

$$T_i = \begin{cases} \frac{1}{P}\int_0^P T_i(t)dt & \text{DC case} \\ \frac{2}{P} \parallel \int_0^P T_i(t)e^{-2\pi j f_0 t}dt \parallel_2 & \text{AC case,} \end{cases} \tag{4.1}$$

where f_0 is the fundamental frequency of the AC excitation source, and $\parallel \cdot \parallel_2$ is the ℓ_2 norm which gives the magnitude of a complex number in this case. If there is no noise in the measurements, then we expect T_i to exhibit no stochastic behavior. However, noise in the measurements leads to a stochastic process where T_i is a random variable. The main sources of noise in the infrared imaging system are: (1) thermal noise, (2) digitization noise, (3) dark noise, and (4) flicker noise [36]. Thermal noise is caused by agitation of charge carriers, and is present in all electronic devices. Digitization noise arises from the use of analog-to-digital converters in the infrared camera. Dark noise is caused by random generation of electron-hole pairs in the quantum detectors which is usually present in photosensitive devices. Flicker noise is also present in almost all

[1]Note that, we denoted $\mathbf{T}$ as thermal maps without noise and $\mathbf{T_m}$ for infrared thermal maps with noise in Chapter 3 previously. We drop the subscript from thermal map $\mathbf{T_m}$, here onwards $\mathbf{T}$ represents the infrared thermal map with noise for simplicity.

electronic devices, which is related to the trapping and detrapping fluctuations of charge carries at the transistor interfaces [15]. We can divide the noise sources into two categories, frequency-dependent and frequency-independent. The first three sources of noise do not show dependency on frequency, and thus, they can be modeled as white noise. White noise has a flat power spectral density that does not vary with the frequency. The last source of noise has frequency dependency, and usually termed as a $\frac{1}{f}$ noise.

The *amplitude of noise* is equal to the standard deviation of the values of T_i. The noise amplitude is commonly referred to as *noise equivalent temperature difference* (NETD) [86, 10], and is given by

$$NETD_{f_0,P} = \sqrt{\frac{1}{k}\sum_{i=1}^{k}(T_i - \bar{T})^2}, \tag{4.2}$$

where $\bar{T}$ is the mean value of the pixel over k successive measurements. If all the pixels are experiencing the same constant temperature, then it does not matter whether the NETD is calculated from k successive measurements of one pixel or from evaluating k pixels on one image.

4.3.1 Noise Reduction using Higher Frequencies

One of the main advantages of using AC excitation is that the frequency-dependent flicker noise component reduces significantly as frequency increases. To analyze the noise in our infrared system, we implement a simple test chip where a circuit block is placed at the center of an otherwise idle chip (the full details of the test chip are given in Section 4.7). The block is toggled at different frequency rates 0.25 Hz, 0.5 Hz, . . ., 8 Hz. At each

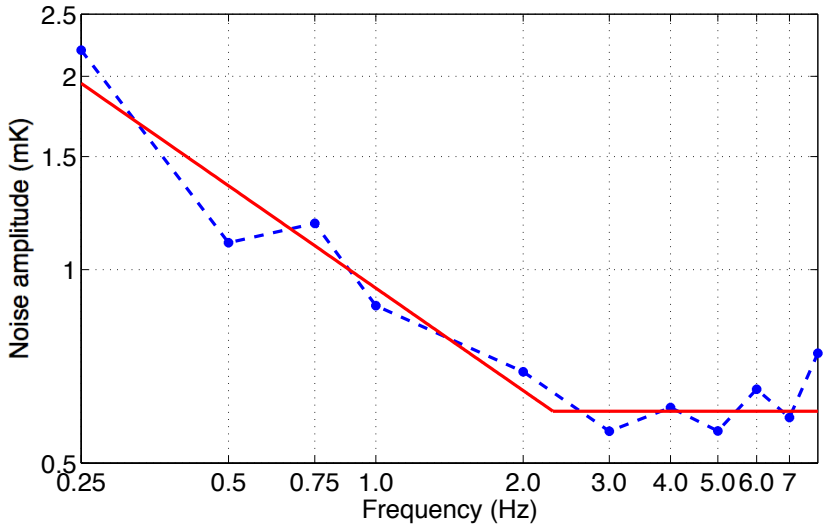

Figure 4.3: log-log plot of noise amplitude as a function of frequency. Dashed blue line gives the noise amplitude from measurements using an integration time of 16 seconds. Red line gives the fitting to measurements. Corner frequency is observed between 2 and 3 Hz. Noise amplitude fitting before corner frequency yield an amplitude of $9.5 \times 10^{-4} f^{-0.51}$.

frequency rate, we capture the thermal emissions for 400 seconds at steady state using a frame rate of 100 Hz. We then use Equation (4.1) (with an integration time of $P = 16$ seconds) to compute the magnitude of the thermal signal at the fundamental frequency. Thus, each pixel yields $\frac{400}{16}$ magnitudes at the fundamental frequency. The noise amplitude is then computed as the standard deviation of these temperature magnitudes as given by Equation (4.2). There is a tradeoff between data acquisition time and the noise of the image. As the integration time increases noise reduces which is discussed in the following section. The integration time of 400 seconds was chosen for this experiment to optimize between data acquisition time and noise.

We plot the noise amplitude (dashed blue line) as a function of frequency in Figure 4.3 on a log-log scale. The plot shows that the noise decreases as a function of frequency until it reaches a *corner frequency*, beyond which the noise amplitude does not show improvement with frequency increase. In our system, this corner frequency is somewhere

between 2 – 3 Hz. Fitting the measurements gives a noise amplitude equal to

$$NETD_{f,P} = \begin{cases} 9.5 \times 10^{-4} f^{-0.51} & f < \text{corner frequency} \\ 6.2 \times 10^{-4} & f \geq \text{corner frequency}. \end{cases} \tag{4.3}$$

The fitted trends are given by the solid red line in Figure 4.3. Thus, the noise spectral density, which is equal to the variance, has a near perfect $1/f$ characteristic as it is equal to the square of the NETD. This indicates strong dominance of flicker noise at low frequencies. Thus, increasing the excitation frequency reduces the impact of flicker noise. However, increasing the excitation frequency beyond the corner frequency has no benefit as noise becomes dominated by white noise which has flat frequency characteristics.

4.3.2 Noise Reduction using Larger Integration Times

As discussed earlier, the second component of noise is white noise. White noise has a Gaussian distribution, and thus by increasing the integration time, we reduce the standard deviation of the thermal signal which is responsible for the amplitude of noise. From the central limit theorem, we know that the standard deviation of the average of a number of a samples of a random variable has $\frac{1}{\sqrt{n}}$ dependency on the number of samples n. Thus, by increasing the integration time, we can reduce white noise proportionally to the square root of integration time [10].

To analyze the white noise in our infrared system, we implement a simple test chip where a circuit block is placed at the center of an otherwise idle chip. The block is toggled at 2 Hz, and the thermal emissions are captured for 400 seconds at steady state using a camera's frame rate of 100 Hz. We then use Equation (4.1) to compute the magnitude of

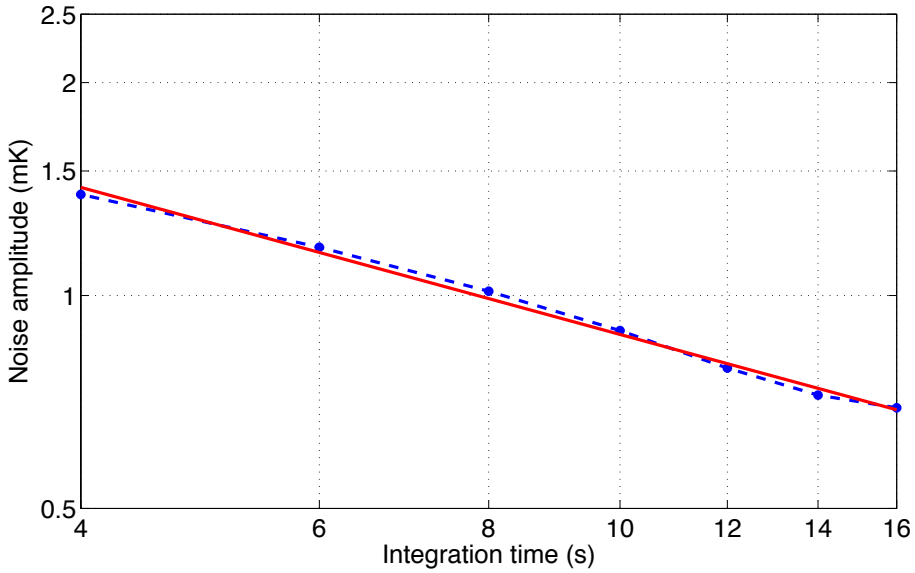

Figure 4.4: log-log plot of noise amplitude as a function of integration time at excitation frequency 2 Hz. Dashed blue line gives actual measurement. Red solid line gives results from fitting.

the thermal signal at 2 Hz for various integration times of 4, 6, . . ., 16 seconds. For each integration time, the noise amplitude is computed as the standard deviation of the thermal magnitudes as given by Equation (4.2).

We plot the noise amplitude (dashed blue line) as a function of integration time in Figure 4.4 on a log-log scale. The plot shows the dependency on the reciprocal of the square root of integration time (solid red line with a fit $\propto 1/n^{0.52}$). As a result, as we increase the integration time, the noise amplitude of our image reduces.

4.4 Impact of using AC Excitation on Spatial Temperature Signal

In this section we analyze the impact of AC excitation on spatial heat diffusion and the strength of the thermal signal. We first mathematically analyze this phenomenon and then demonstrate it experimentally with a test chip.

We first assume a semi-infinite isotropic and homogeneous silicon substrate with standard wafer thickness of $750\,\mu\mathrm{m}$. Heat transfer inside the substrate is governed by the heat diffusion equation:

$$D\nabla^2 T(\vec{r},t) + p(\vec{r},t) = \frac{\partial}{\partial t}T(\vec{r},t), \tag{4.4}$$

where $T(\vec{r},t)$ is the temperature as a function of the radial distance, $\vec{r}$, from the center of the substrate at time t, $p(\vec{r},t)$ is the power function, and D is the thermal diffusivity of silicon ($D = 0.88\,\mathrm{cm}^2/\mathrm{s}$). The standard thickness of 200 and 300 mm wafers is about $750\,\mu\mathrm{m}$, and memory and mobile chips are commonly thinned to less than $100\,\mu\mathrm{m}$. By comparison, at 20 Hz the thermal diffusion length $L_D \equiv \sqrt{D/2\pi f}$ is more than $800\ \mu$m. This justifies treating heat transfer as a two-dimensional problem.

We first consider a sinusoidal point heat source at $\vec{r} = 0$ of unit amplitude that is toggled in time with an angular frequency $\omega = 2\pi f$; i.e., $p(\vec{0},t) = e^{j\omega t}$. From a practical perspective, a point heat source has physical extent smaller than the spatial resolution of the thermal imaging equipment, and thus occupies at most one pixel. We can also express the temperature function in polar coordinates: $T(\vec{r},t) = T(r,t)$, where r is the radial distance from the source and without any dependency on the angle, due to the rotational

symmetry of the single-point excitation. The thermal response $T(r,t)$ to the unit point heat source is the Green's function $g(r,t)$. At steady-state AC, the boundary conditions of the setup are

$$g(r \to \infty, t) = 0 \quad \text{and} \quad \lim_{r \to 0} r \frac{\partial}{\partial r} g = e^{j\omega t}. \tag{4.5}$$

The standard Green's function solution for Equation (4.4) subject to the boundary constraints of Equation (4.5) is given by

$$g(r,t) = K_0 \left(r \sqrt{\frac{j\omega}{D}} \right) e^{j\omega t}, \tag{4.6}$$

where $K_0(\cdot)$ is the modified Bessel function of the second kind. $K_0(\cdot)$ is a complex function and its amplitude and phase cannot be described analytically. Thus, we plot in Figure 4.5 the amplitude and phase of $K_0(r\sqrt{j\omega/D})$ as a function of $r\sqrt{\omega/D}$. For a given r, the figure shows that the amplitude of the temperature decays at a faster rate as ω increases. That is, the temporal excitation frequency controls the extent of spatial heat diffusion. In AC thermography, the thermal image at the fundamental frequency is the one that is used as it has the strongest signal value; thus, in the frequency-domain, only the amplitude of $g(r,t)$ at the fundamental frequency, f_0, is used. That is,

$$g_{\omega_0}(r) = K_0 \left(r \sqrt{\frac{j\omega_0}{D}} \right), \tag{4.7}$$

where w_0 is equal to $2\pi f_0$. Figure 4.6.a shows in a visual illustration the amplitudes of the

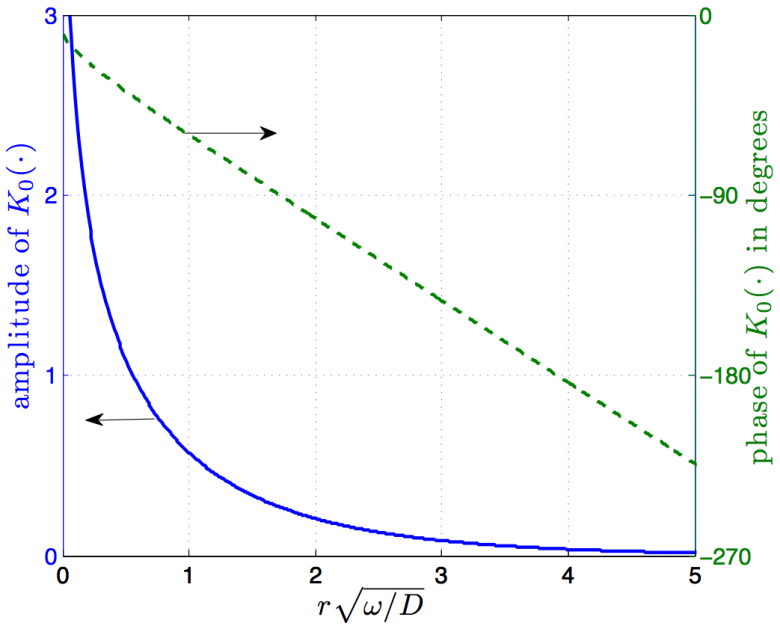

Figure 4.5: Amplitude and phase of the Green's function $g_\omega(r)$ as a function of $r\sqrt{\omega/D}$

fundamental frequency component of $g_{\omega_0}(r)$ for a number of excitation frequencies DC (0 Hz), 1 Hz, 2 Hz, 4 Hz, and 8 Hz. The images in Figure 4.6.a show a clear trend where the spatial extent of heat diffusion reduces as the excitation frequency increases. To quantify this reduction, we plot in Figure 4.6.b and Figure 4.6.c the temperature as a function of the distance (in mm) from the center of the point source and the excitation frequency. Figure 4.6.c gives normalized results of Figure 4.6.b. We observe the following.

- The thermal "inertia" of silicon reduces the amplitude of the temperature change at higher AC frequencies. For example, the plot of Figure 4.6.c shows that the signal drops to 40% of its peak value at a distance of 3.5 mm for DC, 0.6 mm for 1 Hz and 0.35 mm for 8 Hz. Thus, using higher excitation frequencies has the advantage of reducing the extent of thermal diffusion and as a result the contrast of the captured thermal images is improved.

- Increasing the frequency has the disadvantage of reducing the absolute signal value as shown in Figure 4.6.b. This reduction in average signal value ultimately reduces

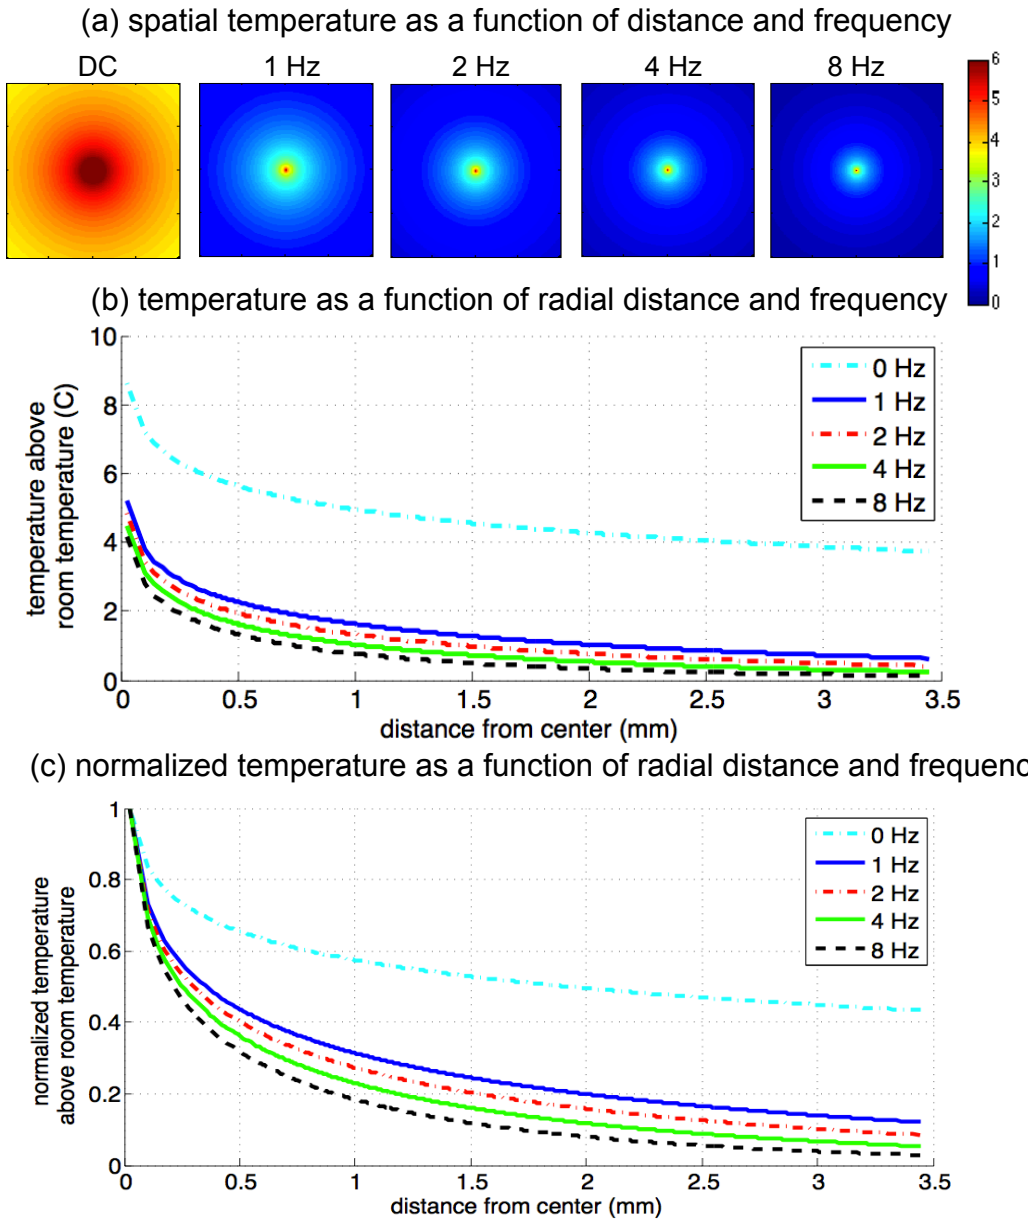

Figure 4.6: Impact of increasing excitation frequency on spatial heat diffusion as computed from the analysis.

the benefit of AC-based thermography at higher frequencies.

To demonstrate experimentally the impact of excitation frequency on spatial heat diffusion, we implement a simple test circuit where a circuit block (approximating a point

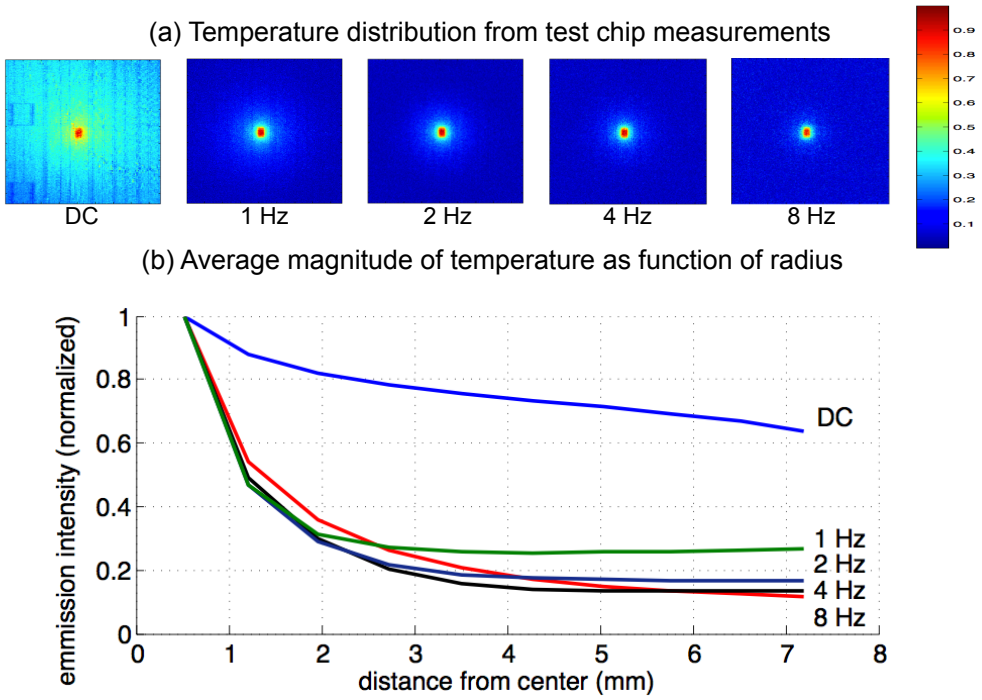

Figure 4.7: Impact of increasing excitation frequency on spatial heat diffusion as measured from the test chip.

source) is placed at the center of an otherwise idle chip. The block is toggled through an enable signal at various excitation rates, and the AC thermography procedure is used to compute the amplitude of every pixel at the fundamental frequency according to Equation (4.1). The thermal amplitudes at the fundamental frequency of all pixels are plotted in Figure 4.7.a, and the normalized thermal signals as a function of distance from the block are given in Figure 4.7.b. The results confirm that using AC excitation reduces the extent of heat diffusion.

To understand the impact of using AC excitation on the final quality of the thermal image, we combine next the analysis of this section, which focused on estimating the signal value, with the noise analysis in Section 4.3 to analyze the signal-to-noise ratio (SNR) as a function of the excitation frequency.

4.5 Signal to Noise Ratio (SNR) Analysis

In Section 4.3, we observed that increasing the excitation frequency decreases the flicker noise up to the corner frequency. Beyond the corner frequency, noise is dominated by white noise rather than flicker noise. Thus, increasing the excitation frequency beyond the corner frequency does not reduce the noise; instead, increasing the integration time is more beneficial. In Section 4.4 we observed that increasing the AC excitation frequency reduces the extent of heat diffusion which reduces the blurring of the underlying power map; however, increasing the excitation frequency has the disadvantage of reducing the absolute temperature signal value at the fundamental frequency. In this section, we investigate the impact of the excitation frequency on both the signal and the noise. The signal-to-noise (SNR) is a combined metric that gives the quality of the thermal image.

Theoretical analysis of SNR. To analyze the SNR, we first compute the average thermal signal over a radius of 3.5 mm using Equation (4.7) at various frequencies from DC to 8 Hz and then divide the results at each frequency by the corresponding noise amplitude as given by Equation (4.3). The SNR values as a function of frequency are plotted in Figure

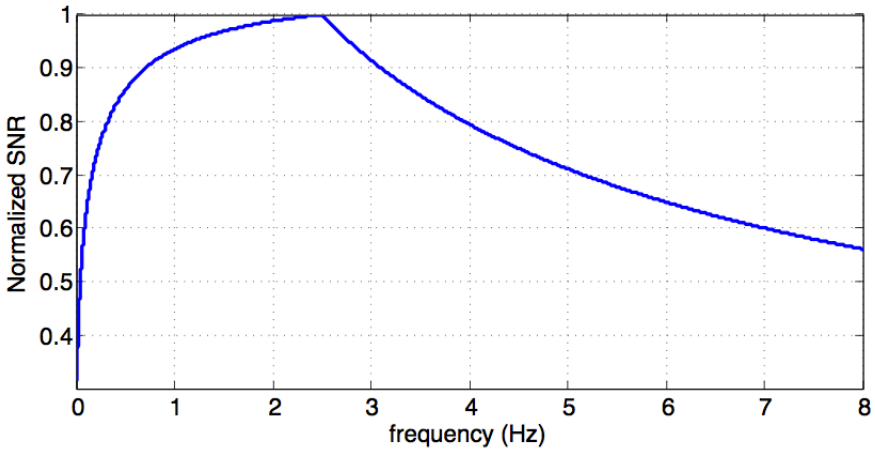

Figure 4.8: Results from theoretical signal-to-noise analysis.

4.8. Clearly, the SNR improves as the frequency increases at the beginning because the reduction in flicker noise and in spatial heat diffusion outweigh the reduction in signal value. But as the frequency is increased beyond the corner frequency, the SNR worsens due to the reduction of the signal strength and lack of reduction in noise.

Empirical analysis of SNR. To measure the SNR, we use the thermal maps measured in the context of the experiment corresponding to Figure 4.7, where we calculate the average temperature over a radius of 3.5 mm from the center of our chip area. The 3.5 mm radius is chosen to cover our experimental test chip area. We divide the results at each frequency by the measured noise amplitude as given in Figure 4.3. The SNR values as a function of frequency are plotted in Figure 4.9. The empirical results show the same trends as the theoretical SNR plot of Figure 4.8, where both plots relatively agree with each other. For instance, the empirical results of Figure 4.9 shows a degradation of the SNR by 57% from its peak value of 140 to 80 at 8 Hz. The same reduction percentage is predicted from the theoretical results given in Figure 4.8.

Our analysis and results show that for our test chip and imaging equipment, the peak

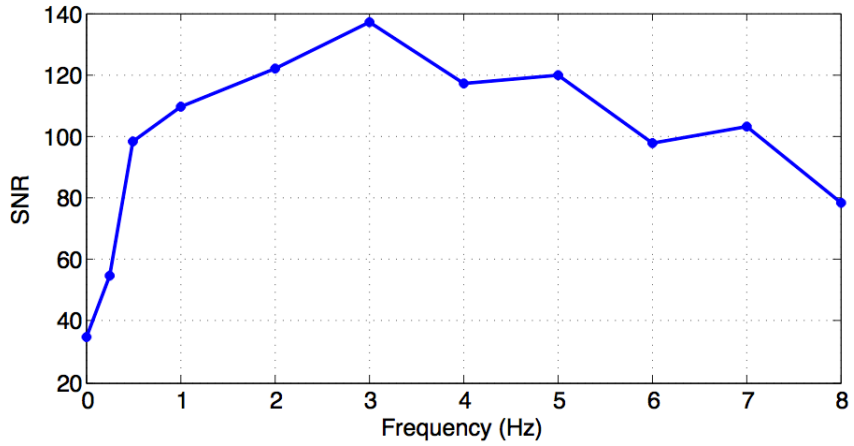

Figure 4.9: Empirical Signal-to-noise ratio at different frequencies with fixed integration time (16s).

SNR occurs around 2-3 Hz. Our analysis procedure is applicable for other chips and other imaging systems, though the exact peak SNR frequency might differ. This difference could arise because of different flicker noise characteristics in the imaging equipment readout electronics and from the dimensions (in particular the thickness) of the test chip. Reducing the thickness of the chip can shift the peak SNR to higher frequencies.

4.6 Thermal to Power Inversion Method

In a discretized form, the steady-state DC relationship between power and temperature can be succinctly described using

$$\mathbf{Rp} + \mathbf{e} = \mathbf{T}, \tag{4.8}$$

where $\mathbf{R}$ is the power-to-thermal modeling matrix, vector $\mathbf{p}$ is the power vector that gives the power consumption of every circuit block, $\mathbf{e}$ is the measurement noise vector, and vector $\mathbf{T}$ is the temperature vector that gives the recorded temperatures at every pixel of the thermal image [33]. The length of vector $\mathbf{p}$ is determined by the number of circuit blocks, and the length of $\mathbf{T}$ is determined by the camera's spatial resolution and the dimensions of the die. In steady-state AC, Equation (4.8) is still valid except that the elements of vectors $\mathbf{p}$ and $\mathbf{T}$ are complex numbers that give the the amplitudes and phases of power and temperature at the fundamental frequency. If the phase of the power signal is considered as reference, then $\mathbf{p}$ is real and $\mathbf{T}$ is complex. The modeling matrix $\mathbf{R}$ is now frequency dependent and its elements are complex numbers as described in Section 4.4. We denote the frequency-dependent model matrix by $\mathbf{R}_f$, and the DC model matrix by $\mathbf{R}_0$.

For power mapping purposes, three steps need to be conducted: i) $\mathbf{T}$ needs to be measured for the chip under characterization after applying the appropriate workload and waiting for AC or DC steady state; ii) $\mathbf{R}$ needs to be estimated for the actual chip in its deployed environment; and iii) given $\mathbf{T}$, $\mathbf{R}$ and the total power consumption, a numerical inversion procedure must be carried to find $\mathbf{p}$. We explain each of these steps in the remainder of this section.

Measuring T. The thermal imaging system captures a discretized thermal emission field with an image sampling rate that is at least twice larger than the AC excitation frequency according to the Nyquist sampling criterion. To measure $\mathbf{T}$, we capture a number of consecutive frames for an integration time, P, as illustrated earlier in Figure 4.2. The temperature at each pixel is computed as

$$T_i = \begin{cases} \frac{1}{P}\int_0^P T_i(t)dt & \text{DC case} \\ \frac{2}{P}\int_0^P T_i(t)e^{-2\pi j f_0 t}dt & \text{AC case,} \end{cases} \tag{4.9}$$

In the case of the AC case, the temperature T_i of pixel i is a complex number that gives the amplitude and phase of the thermal wave at pixel i. The temperatures from all pixels form the elements of the vector $\mathbf{T}$.

Estimating the Modeling Matrix R. For practical power mapping, it is necessary to estimate the modeling matrix $\mathbf{R}$ accurately. The model parameters must capture all modes of heat transfer (e.g., conduction throughout the solid, convection and radiation at the surface). Thus, analytical derivations as described in Section 4.4 are not sufficient. We estimate the model parameters directly from the test chip using the following procedure. For each excitation frequency f, the matrix $\mathbf{R}_f$ can be estimated in a column-by-column basis as follows. Exciting only the k^{th} circuit block is mathematically equivalent to setting

the vector $\mathbf{p}$ to be equal to $[0\ 0 \cdots p_k \cdots 0\ 0]'$, where p_k denote the total additional power (at the fundamental frequency) incurred from exciting block k. The total power of a test chip can be readily measured using an external digital ammeter. Dividing the captured temperature vector $\mathbf{T}$ by p_k gives the values of the k^{th} column of the matrix $\mathbf{R}_f$. Thus, we can measure $\mathbf{R}_f$ column-by-column by enabling each block one by one and repeating the described procedure. It is also possible to carry the same procedure in simulation. In this case, it is first necessary to construct a finite-element model (FEM) of the test chip, its substrate board, and its environment, and then the use the FEM within a numerical simulation environment to estimate the model parameters in the same conceptual way as the described procedure. The matrix $\mathbf{R}_f$ is a function of f and thus it needs to be measured at each desired excitation frequency.

Inverting Temperature to Power. Given the amplitude measurements $\mathbf{T}$, $\mathbf{R}$ and total power, the objective is to find the best power map vector $\mathbf{p}$ that minimizes the total squared error between the temperatures as computed from the estimated power $\mathbf{p}$ and the measured temperatures $\mathbf{T}$; that is,

$$\arg_{\mathbf{p}} \min \parallel \mathbf{R}_f\mathbf{p} - \mathbf{T} \parallel_2^2, \tag{4.10}$$

where $|| \cdot ||_2$ denotes the ℓ_2 norm, and under the constraints that the sum of the elements of the power vector is equal to the total power consumption of the chip p_{total} and that the individual power estimates of circuit blocks must be greater than zero; i.e.,

$$\parallel \mathbf{p} \parallel_1 = \sum_k p_k = p_{total} \tag{4.11}$$

$$\forall k : p_k \geq \mathbf{0}, \tag{4.12}$$

where $||\cdot||_1$ denotes the ℓ_1 norm. The total power is readily measured through an external digital multimeter. Any multimeter has its own tolerance (tol), so for practical purposes we change the constraints in Equation (4.11) to be two inequalities:

$$\sum_k p_k \geq p_{total} - tol \tag{4.13}$$

$$\sum_k p_k \leq p_{total} + tol \tag{4.14}$$

In our implementation, we use MATLAB's quadratic optimization solver (`lsqlin`) to minimize (4.10) under the constraints of Equation (4.12), Equation (4.13) and Equation (4.14). The solver uses the active-set strategy (also known as a projection method) which relies on a two-step solution. The first step calculates a feasible solution point, and the second phase generates an iterative sequence of feasible solution points that converge to the final solution.

4.7 Experimental Results

The objective of our experiments is to assess the improvement in power mapping accuracy using the proposed AC-based framework. For validation purposes, we design a test chip where we can control exactly the switching activity, which yields *reference power maps*. By knowing the reference power maps, we can scientifically validate our thermal-to-power

inversion technique as it is applied to estimate the spatial power maps. We utilize the same test chip which is 90 nm Altera Stratix II (EP2S180) field programmable gate array (FPGA) with 180,000 logic elements as in Chapter 3, but instead of DC excitation sources, we use AC excitation source for the following experiments. The programmability of the test chip allows us to excite the micro-heater block at desired AC frequencies. The basic unit of our test circuit is a *programmable micro heater*, which consists of a number of ring oscillators (ROs) that are controlled by flip-flops that determine the operational status of the micro-heater. We create two kinds of micro-heater designs: *Bi-Level Micro-Heaters* and *Multi-Level Micro-Heaters*. We describe the design of the test chip in detail in Chapter 3.4.

Experiment 1. For this experiment we utilize our bi-level microheater grid (10×10 micro-heaters; each consumes 25 mW when enabled) to create a number of reference maps and then we capture the resultant thermal emissions using the infrared camera. The reference power maps are given in the first column of Figure 4.10. We use our thermal to power inversion method to estimate the spatial power maps and compare them with the reference maps for validation. We conduct power estimation for traditional DC excitation and AC excitation with frequencies of 0.5, 1, 2, 4, 8 Hz using an integration time of 16 seconds. We report the estimated power maps in Figure 4.10. The percentage error for each individual power map is computed as the absolute error between the reference power map and the estimated power map normalized by the total power of the map. That is,

$$Error = \frac{\sum_k |p_k - Pcorrect_k|}{\sum_k Pcorrect_k} \tag{4.15}$$

where $Pcorrect$ is the reference power map and p_k is the value of the k^{th} element in the vector $\mathbf{p}$. The results show that AC-based inversion gives significant reduction in power

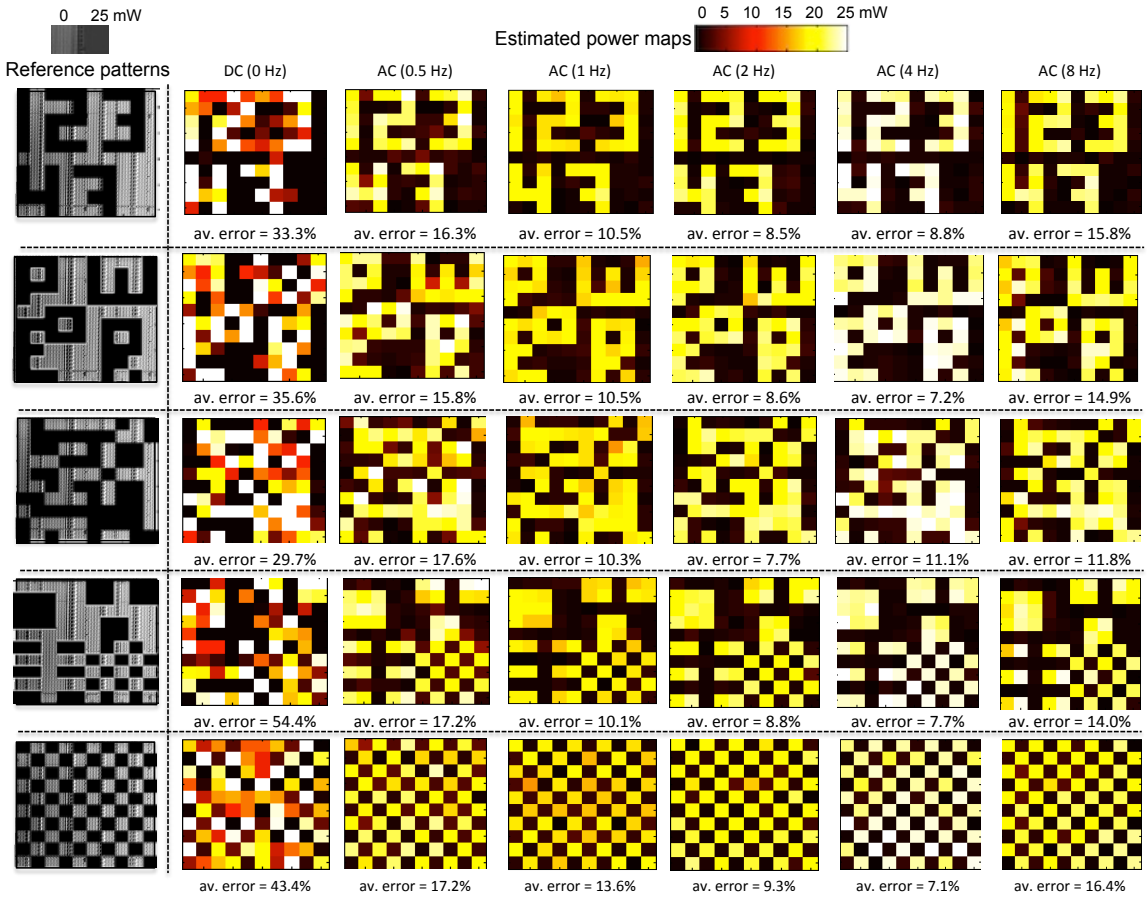

Figure 4.10: Error in post-silicon power mapping for DC and AC excitation without rounding.

mapping error compared to DC-based inversion. The average error decreases from 40% at DC to about 8.5% in the AC (2 Hz) method. Rounding the results of the AC method to the nearest level (0 or 25 mW) yields a perfect estimation of the reference maps with no errors.

Comparison of power mapping error against excitation frequencies. We summarize the power mapping percentage errors of Figure 4.10 in Figure 4.11, which gives the average power mapping percentage error for the five shown maps against all the excitation frequencies including DC. The trends in the power mapping accuracy results are in agreement with the SNR results provided earlier in Section 4.5. It is clearly observed that the

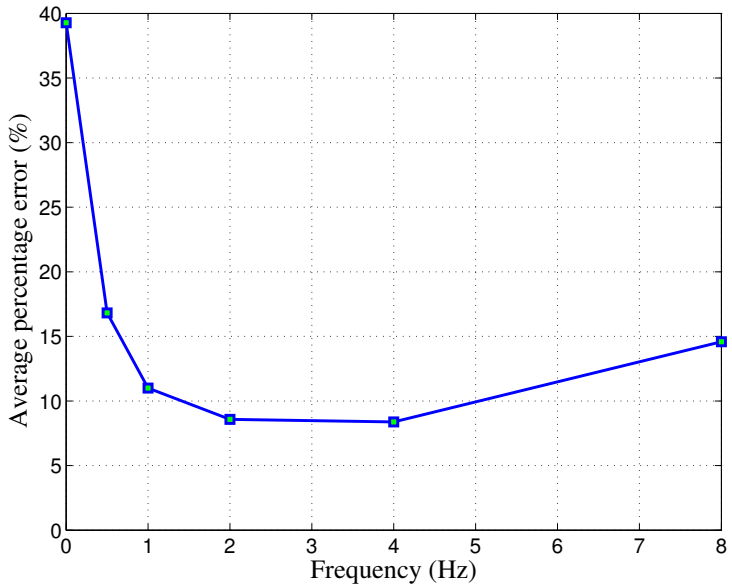

Figure 4.11: Percentage of error versus different frequencies with fixed integration time (16s).

power mapping error reduces drastically when we move from DC excitation to AC excitation until the 2-4 Hz range. Increasing the excitation frequency beyond that range leads to a loss in power mapping accuracy, which is consistent with the SNR results.

Comparison of power mapping error against integration times. In Section 4.3, we

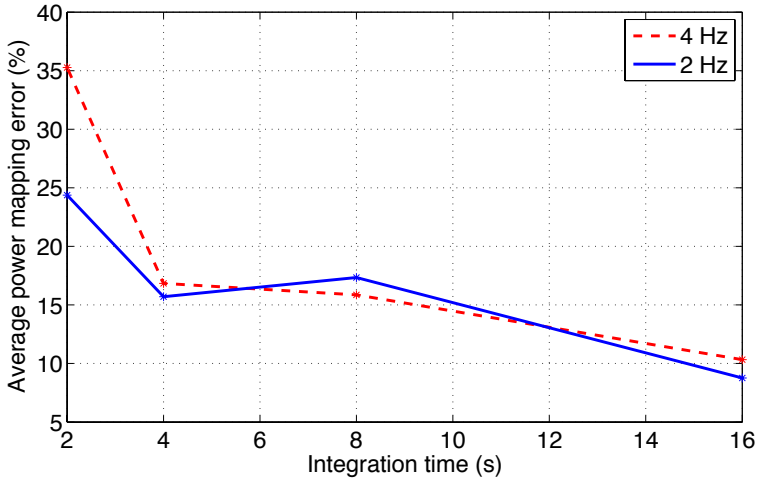

Figure 4.12: Percentage of error versus different integration time.

analyzed the dependency of noise on the integration time. Increasing the integration time (especially at and beyond the corner frequency) reduces the noise, which translates into improved power mapping results. In Figure 4.12, we plot the average percentage error for the five spatial maps as a function of the integration time for 2 Hz and 4 Hz power mapping results. The plot shows consistent improvements in accuracy as the integration time is increased. Increasing the integration time requires additional time for data collection and processing and additional space for storing the measurement data.

Experiment 2. In this experiment, we provide results using the multi-level micro-heater grid design (6×6 micro-heaters; each can consume 0, 25, 60, 102, 142 mW when enabled). In contrast to the previous experiment where the power of each micro heater was only restricted to two levels (0 and 25 mW), the constructed reference spatial power maps in

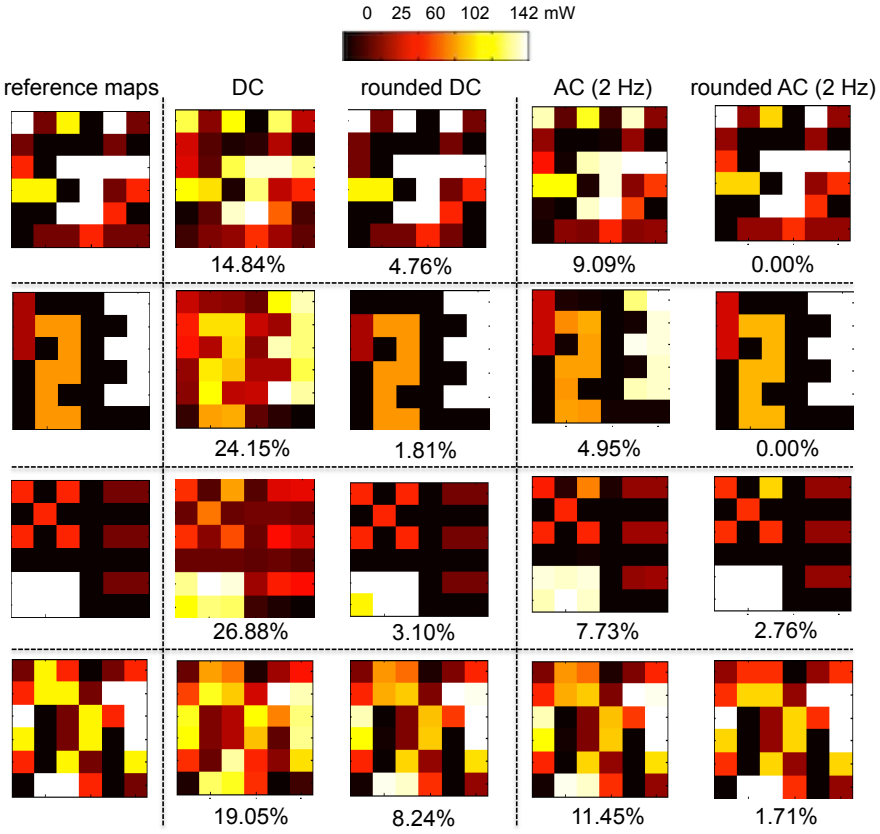

Figure 4.13: Results from multi-level power estimation.

this experiment comprises of power levels of varying intensities (0, 25, 60, 102, and 142 mW). Our inversion procedure can naturally handle any number of levels, and the results of this experiment confirms this capability. Figure 4.13 gives the reference multi-level power maps in the first column. The second column gives the estimated power maps using DC-based excitation. The average error is 21.23%. We give in the third column the results of the DC-based power maps when they are rounded to the nearest power level. In the rounded case the average error drops to an average of 4.48%. In the fourth column we give the results from using the proposed AC-based technique at 2 Hz; the average error is 8.30% which again confirms the drastic improvement compared to the DC case. The last column gives the rounded power maps from AC-based estimation with an average error of 1.11%. The results of this experiment confirm that our technique is capable of handling a large range of power intensities and spatial power maps, and that AC-based power mapping consistently outperforms DC-based power mapping with a large margin.

4.8 Summary

In this chapter we have investigated the challenges in power mapping and proposed AC-based thermography techniques to overcome the challenges. We have presented theoretical analysis and experimental validation for the impact of AC excitation on measurement noise and spatial heat diffusion. To quantify the noise in our system, we have analyzed and quantified the signal-to-noise ratio. We have devised techniques for realistic estimation of the parameters of the thermal to power modeling matrix, and we have devised numerical techniques to invert the thermal emissions into power estimates. We have crafted a programmable test chip to scientifically evaluate the accuracy of our thermal to power inversion methods. Our test chip enables us to create any desired spatial power map. Using a number of constructed intricate power maps, we have demonstrated that AC-based power

mapping dramatically improves post-silicon power mapping by reducing the average error from 40% at DC to about 8.5% in the AC method. We can also observe that rounding the results of the AC method to the nearest level yields a perfect estimation of the reference maps with no errors. We analyzed the power mapping results for different AC excitation frequencies and integration times, and linked these results to the SNR analysis.

Chapter 5

Applications of Post-silicon Power Mapping

5.1 Introduction

In recent years, post-silicon power mapping has emerged as a technique to mitigate the uncertainties in design-time power models and enable effective post-silicon power characterization [33, 61, 78, 17, 84, 72]. We proposed DC and AC-based post-silicon power mapping framework in Chapter 3 and Chapter 4 respectively, which solve many of the open challenges in this area. Using the proposed power mapping techniques, it is possible to estimate detailed runtime spatial power maps of various computing devices accurately. These post-silicon power estimation techniques can be applied to various computing substrates, such as, multi-core processors, FPGAs, mobiles processors and SoCs. In this chapter we present three different applications of the proposed post-silicon power mapping.

1. **Power Mapping of Real Processors:** To test the DC and AC-based post-silicon power mapping framework proposed in Chapter 3 and 4, we utilize three different computing devices. In contrast to previous chapters, where we have created power patterns in a programmable test chip to verify the proposed methods, in this section, we apply our methods in real processors while running practical applications. First, we apply our DC-based power mapping methods on a soft processor embedded in a FPGA, while running practical applications. Second, the DC-based power mapping techniques are applied to a real quad-core processor, while running different real benchmarks. The power maps from the quad-core processor are decomposed into dynamic and leakage power per block. Third, we demonstrate the basic applicability of our AC-based power mapping on a real dual-core processor. For all three computing devices, we present extensive set of experimental results.

2. **Power Mapping using Sparse Thermal Sensors:** In this application, we extend our previous methods by devising a framework for post-silicon power mapping using thermal maps reconstructed from thermal sensor measurements. We give a detailed description of frequency domain based full characterization methods. These reconstructed thermal maps can be used to locate chip thermal hot spots. Using the reconstructed thermal maps, we provide an alternative way for the post-silicon power mapping procedure. The infrared imaging which is most commonly used for post-silicon power mapping can be very expensive. This approach of using thermal sensor measurements reduces the cost of power mapping drastically, since on-chip thermal sensors are already present in the chip for various thermal management purposes, and do not require any extra cost. We present thermal sensor based power mapping results utilizing our programmable test chip and a real quad-core processor.

3. **Power Mapping for Hardware Trojan Detection:** Vulnerability of modern integrated circuits (ICs) to hardware Trojans has been increasing considerably due to the globalization of semiconductor design and fabrication processes. The large number of parts and decreased controllability and observability to complex ICs internals make it difficult to efficiently perform Trojan detection using current detection methods. We propose a completely new post-silicon multimodal approach using runtime thermal and inverted power maps for Trojan detection and localization. Utilizing our novel power mapping framework, we propose various Trojan detection methods involving two-dimensional principal component analysis to reduce dimensionality of thermal and power maps. The *supervised thresholding* method uses a training data set, and the *unsupervised clustering* method require no prior characterization data of the chip. To characterize real-world ICs accurately, we perform our experiments in presence of 20 - 40% CMOS process variation. Our experimental evaluations reveal that our proposed methodology can detect very small Trojans with 3-4 orders of magnitude smaller power consumptions than the total power usage of the chip.

The organization of this chapter is as follows. In Section 5.2 we present power mapping framework and results for real processors, which comprises of embedded soft processor and multi-core processors. Section 5.3 describes the full thermal characterization methods and post-silicon power mapping using reconstructed thermal maps along with experimental results from a FPGA test chip and a real quad-core processor. In Section 5.4, we describe various hardware Trojan detection and localization techniques and simulation results. Section 5.5 summarizes the novelties of all three applications.

5.2 Post-silicon Power Mapping of Real Processors

The proposed post-silicon power mapping in Chapter 3 and 4 can be applied to various computing devices to get access to detailed runtime spatial power maps. In this section, we describe methodologies for applying the DC-based power mapping techniques to two different computing substrates; a soft processor in Section 5.2.1 and real quad-core processors in Section 5.2.2. We apply our AC-based techniques to a real dual-core processor in Section 5.2.3.

5.2.1 DC-based Power Mapping of a Soft Processor

Given that FPGA-based implementations are less area and power efficient than their ASIC counterparts, power characterization for FPGAs is an active topic of research in recent literature [90, 50]. We apply our proposed power characterization methodology that inverts the spatial thermal emissions into power estimates to a soft processor embedded in a FPGA chip. A soft processor (also called soft-core microprocessor) is a microprocessor core that can be wholly implemented using logic synthesis. It can be implemented via different semiconductor devices containing programmable logic (e.g., ASIC, FPGA, CPLD), including both high-end and commodity variations. In many applications, soft-core processors provide several advantages over custom designed processors such as, reduced cost, flexibility, platform independence and greater immunity to obsolescence. These soft processors are optimized for various specific requirements, such as, layout area, high performance, low power consumption, and versatility in a wide range of applications. The power consumption of these processors is not entirely determined during design time as runtime workloads can impact the exact power consumption. To validate each of the claimed optimized features, there is a great need to perform post-silicon characterization of these soft

processors. We propose to apply our thermal to power inversion techniques on soft processors. To evaluate our power mapping methods proposed in Chapter 3, we embed a soft processor in a FPGA test chip. We estimate the spatial power maps of the soft processor during runtime. Following we describe our experimental setup and results.

a. Experimental Setup and Results

In this experiment, we evaluate our spatial power mapping technique by inverting the thermal emissions of the Nios II soft processor, while running the standard Dhrystone 2.1 application, into spatial power estimates. The Nios embedded processor is a general-purpose RISC CPU implemented as a soft core in Altera FPGAs. We use 90 nm Stratix II FPGA for our experiment. We consider four different configurations of the Nios II processor: the economy model Nios II/e, the standard model Nios II/s with multipliers implemented in the FPGA's logic blocks, the standard model Nios II/s with multipliers implemented in the FPGA's DSP blocks, and the full-performance model Nios II/f. For all four cases, the processor frequency is set to 150 MHz. We first constrain the logic blocks of the Nios II processor to fit into a 30×30 array of logic blocks (about 6.4 mm $\times$ 7.0 mm) of the layout as shown in Figure 5.1.a and then capture the steady-state thermal emissions $\mathbf{T}$ from this area as shown in Figure 5.1.b. The total power consumption of the four configurations are 171 mW, 324 mW, 315 mW and 477 mW respectively. For the purpose of spatial power mapping, we discretize the layout area of the soft processor into 6×6 regions, and thus the $\mathbf{p}$ vector is comprised of 36 elements that need to be estimated.

To compute the power map $\mathbf{p}$ from the thermal emissions $\mathbf{T}$, we need to estimate the modeling matrix $\mathbf{R}$ as described in earlier chapters. We estimate the matrix $\mathbf{R}$ in a similar way as discussed in Chapter 3 on a column-by-column basis. We note that the k^{th} column of matrix $\mathbf{R}$ can be obtained by setting the vector $\mathbf{p}$ to be equal to $[0\,0 \cdots 1 \cdots 0\,0]^T$, where

the "1" is at the k^{th} location of the $\mathbf{p}$ vector, and then use the resultant emissions $\mathbf{T_k}$ directly as the k^{th} column of matrix $\mathbf{R}$. To realize this setting, we utilize the fact that FPGAs are programmable. For each power region k, we embed, and turn on, ring oscillators precisely into the logic array blocks that are available in the region, while the rest of the blocks in the design are inactive. Such precise embedding is possible with Altera's Quartus II tool. The resultant thermal emissions $\mathbf{T_k}$ from such embedding are then normalized by the total power p_k that is measured externally through the digital multimeter. Thus, column k of matrix $\mathbf{R}$ is equal to $\mathbf{T_k}/p_k$. We automate the whole process in order to measure the 36 columns of $\mathbf{R}$ with fast turn-around time.

With the estimated matrix $\mathbf{R}$ and thermal emissions from the Nios II processor (Figure 5.1.b), we compute the spatial power maps using the optimization formulation as follows,

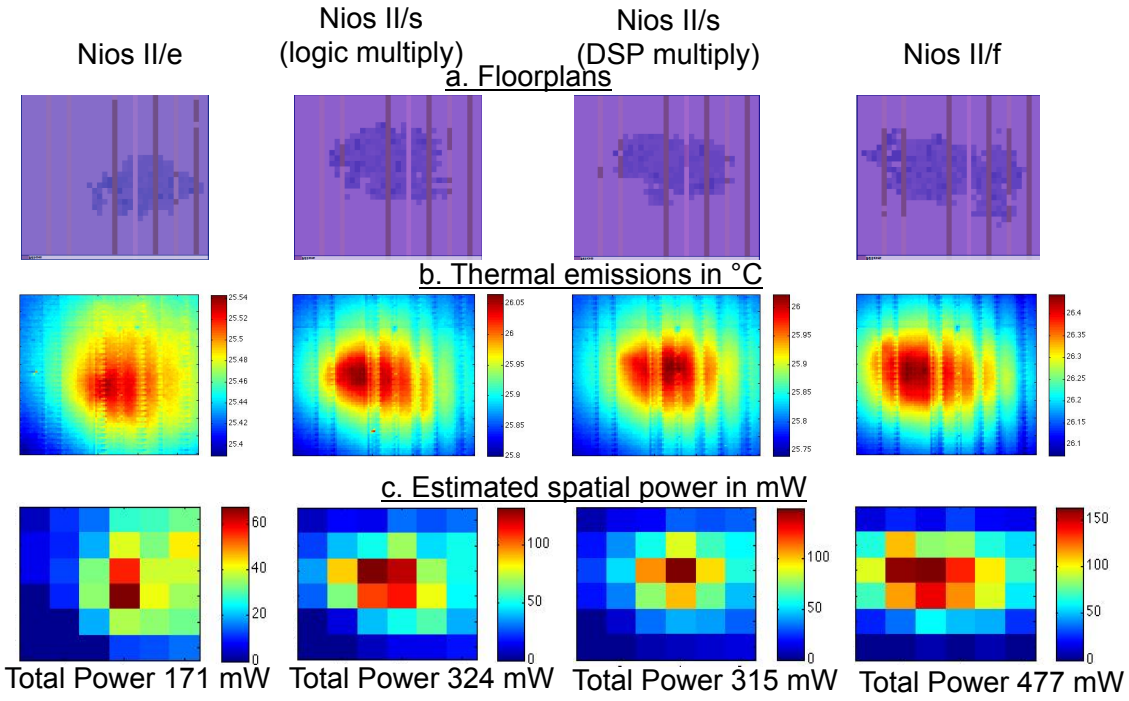

Figure 5.1: Spatial power estimates of Nios II processors running Dhrystone 2.1 application.

$$\mathbf{p} = \arg_{\mathbf{p}} \min ||\mathbf{Rp} - \mathbf{T}||_2^2, \tag{5.1}$$

where the vector $\mathbf{p}$ denotes the power, $\mathbf{R}$ is the modeling matrix, and the vector $\mathbf{T}$ denotes the measured thermal map using infrared imaging. The total power constraints are also imposed to make the solutions unique as follows.

$$\begin{aligned} ||\mathbf{p}||_1 &\leq p_{total} + tol, & (5.2)\\ ||\mathbf{p}||_1 &\geq p_{total} - tol, \text{and} & (5.3)\\ \mathbf{p} &\geq \mathbf{0} & (5.4) \end{aligned}$$

where p_{total} is the total measured multimeter power and tol is the tolerance of digital multimeter. The detailed description of the procedure is given in Section 3.3. The estimated spatial power maps in mW are plotted in Figure 5.1.c. Our estimated spatial maps augment the floorplan with valuable spatial power density estimates. The revealed detailed power estimation maps can be used to calibrate the estimates from high-level dynamic power modeling tools. Given that the design of the Nios II processor is proprietary, it is not possible for us to match the spatial power consumption estimates to the various functional blocks of the processor.

5.2.2 DC-based Power Mapping of a Multi-core Processor

Modern multi-core processors designs are highly complex, incorporating a number of independent cores with billions of transistors. This complexity makes accurate pre-silicon

power modeling a very difficult task for multi-core processors [13, 77]. Furthermore, workloads and process variability alter the power consumption during runtime, making it harder to accurately estimate power consumption during design time. In this section, we extend our framework for DC-based post-silicon power mapping and modeling for real multi-core processors. We add to our framework the capability to identify the dynamic and leakage power consumption of blocks of multi-core processors under different workloads, while simultaneously analyzing the impact of process variability.

a. Power Mapping Framework

Figure 5.2 illustrates the framework of the proposed power mapping method. In our setup, the processor's regular fan and metal heat spreader are removed and replaced by an infrared-transparent heat sink with silicon windows. Laminar mineral oil flow is pumped through the heat sink on top of the processor's die with high flow rate to remove its heat [33, 60, 59, 82]. During runtime, realistic workloads are applied to the processor and the steady-state or averaged thermal map is captured with the infrared camera. We will use $\mathbf{T}_{oil}$ to denote the vector that corresponds to the captured temperature map of the processor.

Replacing the fan and copper (Cu) heat spreader with an oil-based infrared transparent heat sink changes the thermal map of the die [40]. The changes in the thermal map have no impact on dynamic power, but they could change the leakage power characteristics slightly. This impact of the change in leakage power on the power mapping results is very small for our purpose and ignored for experimental results. For power mapping, it is necessary to have an accurate modeling matrix $\mathbf{R}$ that relates temperature to power at the steady state. This modeling matrix can be estimated experimentally [33, 17] or numerically using finite-element modeling (FEM) methods [45]. We propose to use FEM

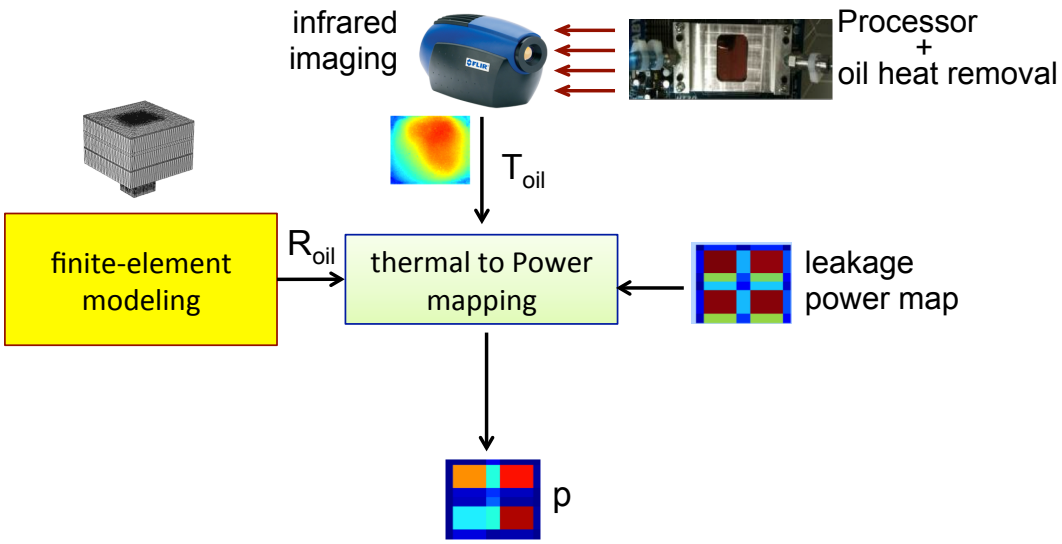

Figure 5.2: Power mapping framework.

methods to accurately estimate the modeling matrix $\mathbf{R}_{oil}$ for the case of the oil-based heat removal sink[1]. In particular, we solve following optimization problem to reconstruct the power map of the die.

$$\begin{aligned} \mathbf{p} \quad &= \quad \arg_{\mathbf{p}} \min \|\mathbf{R}_{oil}\mathbf{p} - \mathbf{T}_{oil}\|_2^2 \\ &\text{such that} \quad p_{lkg,i} \leq p_i \quad \forall i \end{aligned} \tag{5.5}$$

where, $\mathbf{p}$ is the reconstructed power vector, $\mathbf{T}_{oil}$ is the thermal map, $p_{lkg,i}$ denotes the leakage power in the i^{th} die-block, and p_i denotes i^{th} element of $\mathbf{p}$, i.e., the power in the i^{th} block of the die. By solving the above optimization problem, we obtain the total power of each block for the die. The dynamic power of each block is readily obtained by subtracting the leakage power from the reconstructed total power. Using the $p_i \geq p_{lkg,i}$ constraints helps in ensuring that dynamic power for all blocks is always positive. Our

[1]The FEM modeling of the $\mathbf{R}_{oil}$ matrix is done by Kapil Dev. Details can be found in the work [21]

power mapping framework provides the dynamic, leakage, and total powers for each block of a processor given its thermal map.

b. Mapping Leakage Variability

Leakage power, especially its dominant sub-threshold component, depends exponentially on temperature. But within the typical chip operation range 25 - 85 °C, it has a quadratic dependency on temperature, which can be modeled by second-order Taylor series expansion at a reference temperature. In order to compute the chip's spatial leakage power map, we divide the die area into a grid with large number of locations n. For each location i, we develop a second-order Taylor expansion model for leakage power, $p_{lkg,i}$, as a function of the average temperature, T_i, of location i. The expansion around a reference power, $p_{ref,i}$, and temperature, $T_{ref,i}$, is given by

$$p_{lkg,i} = p_{ref,i} + \alpha_{1,i}(T_i - T_{ref,i}) + \alpha_{2,i}(T_i - T_{ref,i})^2 \tag{5.6}$$

where $\alpha_{1,i}$ and $\alpha_{2,i}$ are the model coefficients for location i that depend on the voltage, process variability, and structure of devices. The total leakage power, P_{lkg} is the sum of leakage of the chip's n locations, which can be written as:

$$P_{lkg} = \sum_{i}^{n} p_{ref,i} + \sum_{i=1}^{n} [\alpha_{1,i}(T_i - T_{ref,i}) + \alpha_{2,i}(T_i - T_{ref,i})^2]$$

which can be re-arranged as

$$\Delta P = \sum_{i=1}^{n} \alpha_{1,i} \Delta T_i + \alpha_{2,i} \Delta T_i^2, \tag{5.7}$$

where $\Delta P = P_{lkg} - \sum_i p_{ref,i}$ and $\Delta T_i = T_i - T_{ref,i}$. Note that ΔP, which is the change in total power, is readily obtained using an external multimeter that measures the total power of the processor, and ΔT_i is measured using the thermal maps provided from our infrared imaging or from the translated thermal maps.

To learn the model coefficients, we repeat the thermal conditioning experiment m times with different ambient temperatures, and for each experiment, we measure the change in total power and change in the thermal map. The j^{th} thermal conditioning experiment provides a thermal image which consists of $\Delta T_{j,i}$ at each chip location i and an incremental total leakage power ΔP_j, which creates an instance of Equation (5.8).

$$\Delta P_j \quad = \quad \sum_{i=1}^{n} \alpha_{1,i} \Delta T_{j,i} + \alpha_{2,i} \Delta T_{j,i}^2 \tag{5.8}$$

The results from the m thermal conditioning experiments can be assembled into a system of equations as follows,

$$\begin{bmatrix} \Delta T_{1,1} & \Delta T_{1,1}^2 & \cdots & \Delta T_{1,n} & \Delta T_{1,n}^2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \Delta T_{m,1} & \Delta T_{m,1}^2 & \cdots & \Delta T_{m,n} & \Delta T_{m,n}^2 \end{bmatrix} \begin{bmatrix} \alpha_{1,1} \\ \alpha_{2,1} \\ \vdots \\ \alpha_{1,n} \\ \alpha_{2,n} \end{bmatrix} = \begin{bmatrix} \Delta P_1 \\ \vdots \\ \Delta P_m \end{bmatrix} \tag{5.9}$$

We solve the above system of equations using least-square regression to find the $2n$ α first-order and second-order model coefficients. To compute the total reference leakage power, $\sum_i^n p_{ref,i}$ in Equation (5.7), one can change the ambient temperature of the chip while keeping the dynamic power constant (by running a stable workload), and measuring the total power consumption and the average chip temperature simultaneously. To estimate the total reference leakage power, an exponential model of the measured power to the

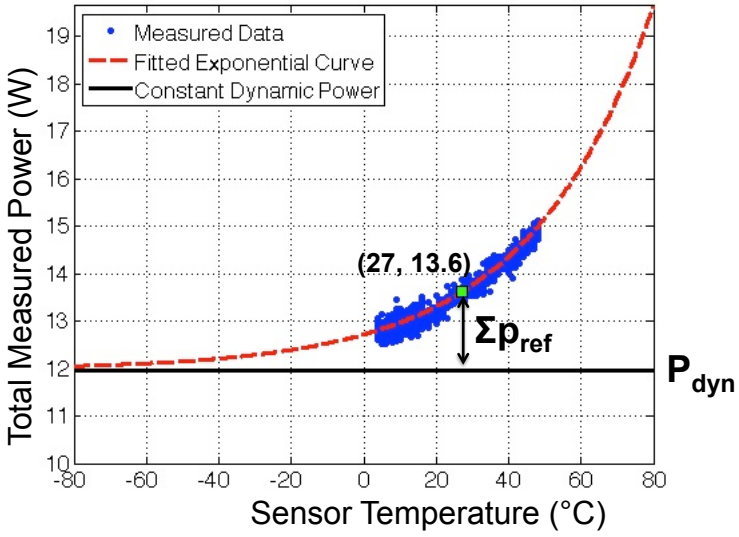

Figure 5.3: Measured power vs. average chip temperature, while keeping dynamic power unchanged. P_{dyn} denotes the dynamic power and $\sum p_{ref}$ denotes the total leakage power at reference temperature 27 °C.

chip's average temperature can be used to extrapolate it to the point where leakage power tapers off. As shown in Figure 5.3, for our experimental quad-core processor, we estimate the total reference leakage power at 27 °C as 1.6 W, and stable dynamic power, P_{dyn} as 12 W. For a particular chip, these coefficients need to be computed only once, and then the same coefficients are used for estimating fine-resolution leakage maps for any thermal map of the chip as given in the framework of Figure 5.2.

Process Variability Mapping. In a typical power mapping experiment, the temperature of location i is plugged into Equation (5.6) to estimate the leakage power of location i. If it is desired to estimate the inherent spatial leakage variability arising from process variability, then a fixed temperature could instead be plugged into the equations of all chip locations. By using the same temperature everywhere, the leakage variations that arise will be due to the coefficients $\alpha_{1,i}$ and $\alpha_{2,i}$ which are dependent on the inherent process variability, assuming a fixed operating voltage.

c. Experimental Setup and Results

Our experimental system consists of a motherboard fitted with a 45 nm AMD Athlon II X4 610e quad-core processor and 4 GB of memory. The motherboard runs Linux OS with 2.6.10.8 kernel. The floorplan of the processor with 11 different blocks is shown in Figure 5.4. We treat each core as one block, as we could not find public-domain information on the make-up of blocks within each core. It is worth mentioning that our proposed technique of power mapping is generic and will work for any arbitrary layout details we use for reconstructing power maps. The processor has 4×512 KB L2 caches, but it lacks a shared L3 cache. The area in the center is occupied by the northbridge and other miscellaneous components such as the main clock trunks, the thermal sensor, and the built-in thermal throttling and power management circuits. The periphery is composed of the devices for

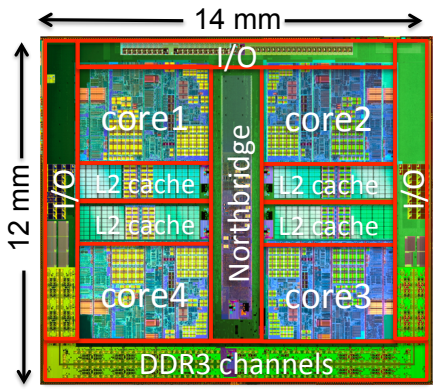

Figure 5.4: Layout of the quad-core AMD Athlon II X4 processor.

I/O and DDR3 communication. The processor supports four distinct DVFS settings. We set the DVFS to 1.7 GHz.

We image the processor using a mid-wave FLIR 5600 camera with 640×512 pixel resolution. We also intercept the 12 V supply lines to the processor and measure the current through a shunt resistor connected to an external Agilent 34410A digital multimeter, which enables us to log the total power measurements of the processor. To implement thermal conditioning in our experimental setup, we use a thermoelectric device and a fluid monitoring device in line with the oil flow as shown in Figure 5.5. By changing the voltage and current of the thermoelectric device, we can either cool or heat the fluid to any desired temperature. Thus, we setup a feedback control system to control the fluid temperature

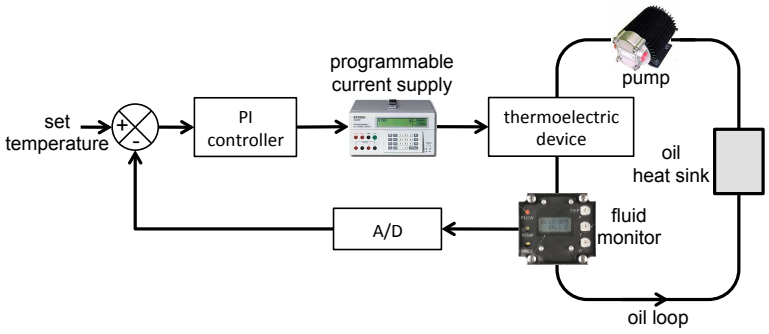

Figure 5.5: Experimental setup for thermal conditioning.

	memory bound	**processor bound**
Integer point	*omnetpp*	*hmmer*
Floating point	*soplex*	*gamess*

Table 5.1: Selected SPEC CPU2006 benchmarks.

to any desired set temperature point. In the feedback loop, the fluid temperature is compared to the set point and the error is fed to a PI controller, the output of which derives the programmable power supply of the thermoelectric device.

Experiment 1: The goal of the first experiment is to demonstrate the results of power mapping for the processor using different number of workloads and different workload characteristics. Our workloads come from widely used SPEC CPU2006 benchmark suite. We selected four benchmark applications, which cover both integer point and floating point computations and processor-bound and memory-bound characteristics. These benchmarks are listed in Table 5.1. We ran different cases of workload sets. For each experiment, we captured the steady-state thermal image using an infrared-camera and reconstructed the underlying power maps from the thermal maps. We decomposed the total power maps into dynamic and leakage power dissipation of each block of the processor and analyzed the spatial leakage variability. The per-block power results for 10 sample workload cases are presented in Table 5.2. We also report the total dynamic power, total leakage power, and the sum of leakage and dynamic power. The results show that the leakage power is on the average about 11% of the total power. We also report in the last column the total measured power through the external multimeter. We notice that our total estimated power through infrared-based mapping achieve very close results with an average absolute error of 0.84 W of the measured power. The differences could be either due to modeling inaccuracies or due to the fact that the measured total power also include the power consumed by the off-chip voltage regulators, and thus, it does not represent the net power consumed by the processor.

				Reconstructed total power (W) for each block											Total power (W)			
c1	c2	c3	c4	c1	L2-1	c2	L2-2	c3	L2-3	c4	L2-4	I/O	N. B.	DDR3	dyn	lkg	dyn+lkg	meas
o	-	-	-	3.21	0.78	1.13	0.45	1.58	0.24	2.11	0.49	0.82	4.62	0.92	14.16	2.18	16.35	16.77
h	-	-	-	4.43	0.90	0.90	0.47	1.33	0.28	2.04	0.47	0.81	5.20	0.79	15.36	2.24	17.62	18.45
s	-	-	-	3.24	0.84	1.03	0.47	1.56	0.23	2.10	0.46	0.80	4.78	0.88	14.19	2.18	16.39	17.06
s	s	-	-	3.06	0.95	2.52	0.92	1.63	0.26	2.32	0.47	0.88	6.20	0.89	17.70	2.37	20.09	19.75
o	-	h	-	3.07	0.93	0.74	0.60	4.81	0.49	2.24	0.47	0.83	6.74	0.86	19.31	2.46	21.78	21.57
s	-	g	-	3.31	1.01	0.75	0.61	4.72	0.58	2.30	0.47	0.83	6.83	0.85	19.76	2.49	22.25	21.86
s	s	s	-	2.92	1.00	2.29	1.01	3.34	0.49	2.50	0.49	0.89	7.24	0.94	20.57	2.53	23.11	22.26
o	s	g	-	3.02	1.08	2.39	1.16	4.89	0.67	2.54	0.49	0.91	8.36	0.84	23.66	2.71	26.36	24.77
o	s	h	-	2.98	1.05	2.31	1.13	5.01	0.58	2.55	0.49	0.92	8.28	0.88	23.49	2.70	26.19	24.58
s	s	s	s	2.88	1.12	2.17	0.99	3.08	0.54	4.46	0.56	0.92	8.14	0.96	23.13	2.68	25.83	24.31

Table 5.2: Power-mapping results for 10 test cases. c1, c2, c3 and c4 stand for core 1, core 2, core 3 and core 4 respectively; `o`, `h`, `s` and `g` stand for `omnetpp`, `hmmer`, `soplex` and `gamess` respectively; N.B. stands for north bridge block; dyn stands for dynamic; lkg stands for leakage; dyn+lkg is the total power reconstructed from post-silicon in infrared imaging; meas is the total power measured through the external digital multimeter.

Experiment 2: To see the impact of increasing number of applications on the power consumption of different blocks, such as, core, cache, northbridge, I/O, DDR3 channels, we run `soplex` in four different ways. First, we run one instance of `soplex` on core 1, second, we run two instances of `soplex` on core 1 and core 2, third we run three instances of `soplex` on core 1, core 2 and core 3 and last we run four instances of `soplex` on all four cores. Figure 5.6 shows the trend of power consumption of different blocks in the processor as we increase the number of applications. When a core is idle it usually clock

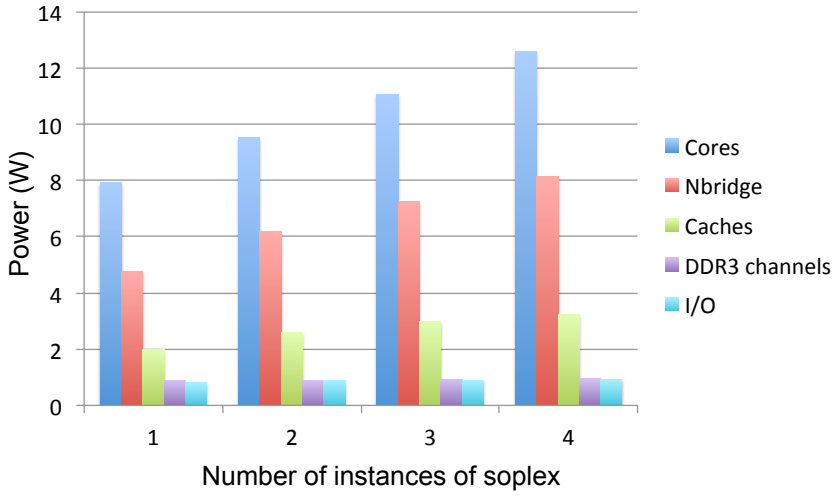

Figure 5.6: Increasing number of instances of `soplex` in the quad-core processor

gates, and consumes minimum power, but as we increase the number of applications, the total power of the four cores increases proportionally. In contrast, the power consumption from other blocks such as the northbridge, I/O, DDR3 do not change as much depending on the number of workloads, because those blocks do not clock gate and they are always operational.

Experiment 3: To estimate the leakage profile for the AMD quad-core processor, we perform the thermal conditioning techniques described earlier in this section and shown in Figure 5.5, where we increase the chip temperature from 27 °C to 55 °C by increasing the infrared transparent cooling fluid temperature from 18 °C to 45 °C, and measuring the associated changes in power consumption and thermal profiles of the chip using infrared imaging. We divide our chip into small blocks of size about 0.4 mm^2 resulting into approximately 418 first-order and 418 second-order coefficients. In order to maintain the stability of the least square estimation, the maximum number of coefficients i.e. the leakage power resolution is limited by the available number of instances of Equation (5.8). We collected approximately 2000 data points to solve our least square estimation. The total reference leakage power, $\sum p_{ref}$ in Equation (5.7) is estimated by changing the die ambient temperature as shown earlier in Figure 5.3.

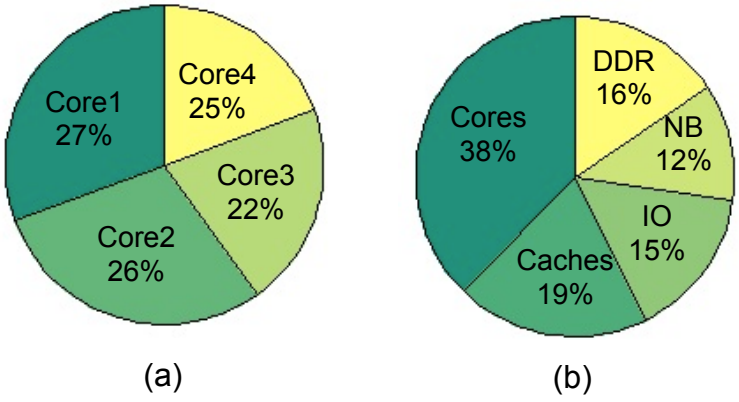

Figure 5.7: a) Percentage Leakage power per core with its L2 cache, and b) Percentage Leakage power per block type.

To uncover the underlying leakage spatial-variability introduced by process variability, we assume constant temperature across the die, and measure the leakage power for each grid location. Figure 5.7.a shows the percentage of leakage power for each core with its L2-cache. Core 1 has approximately 5% more leakage than the lowest power core. This result for instance can be used to bias the operating system scheduler to allocate applications on the lower-leakage cores before the higher-leakage cores. Figure 5.7.b gives the total leakage power distribution among different blocks. There is approximately 10.3% within-die variations among all the blocks.

5.2.3 AC-based Power Mapping of General-Purpose Processors

In this section we demonstrate the basic applicability of our AC-based power mapping technique on general-purpose processors. We focus on demonstrating the ability to (i) alternate between two discrete power levels at different excitation frequency, and to (ii) analyze the resultant thermal images to show the reduction of the spatial heat diffusion as a function of the excitation frequency.

For our demonstration we utilize an AMD Athlon II dual-core processor that is em-

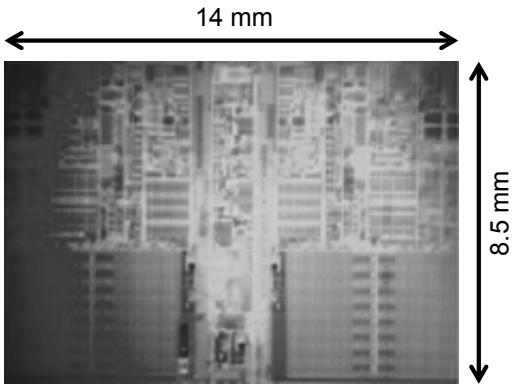

Figure 5.8: AMD Athlon II dual-core processor chosen for demonstration.

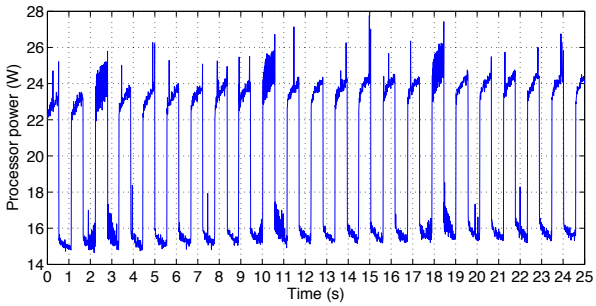

Figure 5.9: Impact of alternating DVFS between two levels of power.

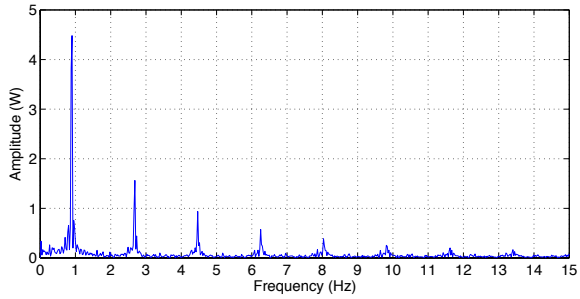

Figure 5.10: Frequency domain representation of the power signal.

bedded in a motherboard with 4 GB of memory and running the Linux operating system. The layout of the processor is given in Figure 5.8. We execute a floating-point application of a stable nature on one of the cores. To alternate the power level between two values, we execute a script that alternates the dynamic voltage-frequency setting (DVFS) between two levels. For example, Figure 5.9 shows the power signal of the processor alternating between two values completing a full cycle every 1.1 seconds. Figure 5.10 gives the amplitude of the power signal in the frequency domain, clearly showing the fundamental component of the square wave at 0.9 Hz and its odd harmonics. Switching the DVFS settings takes a few microseconds, which is a negligible amount of time compared to the frequency of alternating the DVFS setting, which results in excellent frequency domain characteristics.

Figure 5.11.a gives the DC thermal map when the DVFS settings does not alternate.

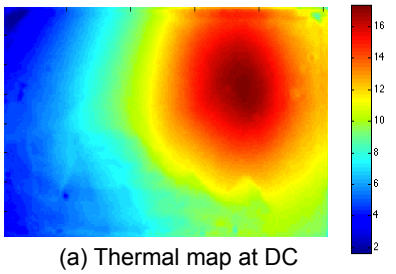

(a) Thermal map at DC

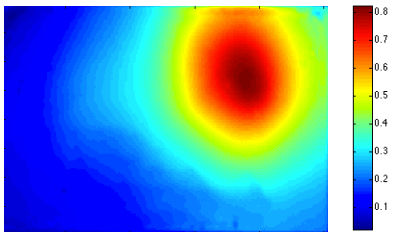

(b) Thermal map at 0.9 Hz

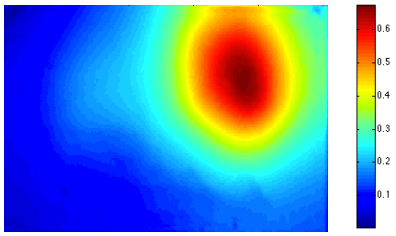

(c) Thermal map at 1.8 Hz

Figure 5.11: Thermal maps at various excitation frequencies.

Using our DVFS alternating script, we collect the thermal emissions and process them as described in Section 4.2 for the cases of 0.9 Hz and 1.8 Hz. Figure 5.11.b gives the AC thermal map when DVFS settings alternate at 0.9 Hz, and Figure 5.11.c gives the the AC thermal map when DVFS settings alternate at 1.8 Hz. The thermal maps show that the extent of spatial heat diffusion reduces as the alternation frequency of the DVFS setting increases. The results confirm the applicability our AC-based technique on general-purpose processors.

To perform full AC-based post-silicon power mapping to any general-purpose processor, there are two main challenges, (1) to estimate an accurate modeling matrix $\mathbf{R}$ corresponding to the excitation frequencies, and (2) to be able to control the excitation frequencies of the power sources in the chip. The modeling matrix $\mathbf{R}$ can be obtained

by using the actual design and layout of the chip to conduct a FEM simulation coupled with a heat diffusion simulation which is discussed in Section 5.2.2 [33, 61, 21]. The only difference for the AC-case is that the modeling matrices have to be estimated by FEM at various excitation frequencies. The second requirement is more challenging because we do not have control over the excitation frequencies of each floorplan block in a real processor. By switching DVFS settings, or running applications between high and low activity phase, it is possible to control the voltage/frequency for some parts of the chip, e.g. cores, but other blocks of the chip e.g. cache memories, I/Os, Nbridge remain at DC state. So, we cannot filter out the temperature corresponding to the fundamental frequency as discussed in Chapter 4, as it will filter out thermal responses corresponding to the blocks at DC. These hurdles need to be overcome in order to perform the full AC framework to modern processors to improve post-silicon power estimates. One emerging future direction to tackle power management concern in high performance microprocessor designs is to accommodate fast on-chip voltage regulator modules (VRMs) that provide fine-grain on-chip voltage switching capabilities [44, 47]. If available, these VRMs could be utilized to control the AC excitation of the power sources of modern general-purpose processors.

5.2.4 Conclusions

In this work, we utilized our proposed power mapping framework and devised techniques to apply it to power mapping applications for real processors. We have introduced multiple novel techniques that advance the state-of-the-art post-silicon power mapping and modeling for both the case of soft processors embedded in FPGAs or real multi-core processors. For power characterization of FPGAs, we proposed new techniques for thermal to power inversion using quadratic program optimization. We used our DC power mapping methodology to estimate the spatial power distribution of an embedded soft processor dur-

ing operation. For multi-core processors, we have devised accurate finite-element models that relate power consumption to temperatures. We have proposed techniques to model leakage power through the use of thermal conditioning. These leakage power models were used to yield fine-resolution leakage power maps and within-die variability trends for multi-core processors. We analyzed the power consumption of different blocks of a real quad-core processor under different workload scenarios from the SPEC CPU2006 benchmarks. Our results reveal a number of insights into the make-up and scalability of power consumption in modern processors. We have also demonstrated the applicability of our AC-based techniques on a real dual-core processor.

5.3 Post-silicon Power Mapping from Sparse Thermal Sensor Measurements

Hot spots impact directly all key circuit metrics, including: lifetime and reliability, speed, power, and costs. Hot spots reduce the mean time to failure as most failure mechanisms (e.g., electromigration, time dependent dielectric breakdown, and negative bias temperature instability) have strong temperature dependencies [74]. Furthermore, different thermal expansion coefficients of chip materials cause mechanical stresses that can eventually crack the chip/package interface [11]. Elevated temperatures also slow down devices and increase interconnect delays which might lead to timing failures [11]. High temperature also increase leakage power, which could lead to thermal runaway [52]. Many-core processors with 100s - 1000s of cores will localize power consumption which further increases thermal problems [38].

To manage runtime thermal variations, circuit designers embed within-die thermal sensors that acquire temperatures at few selected locations. The acquired temperatures are then used to guide runtime thermal management techniques [87, 56]. Inaccurate thermal monitoring arises from: limited number of sensors, remoteness of sensors from hot spots locations, manufacturing variability, and analog-to-digital conversion accuracy [87]. While increasing the number of thermal sensors can reduce these errors, thermal sensors and their support circuitry (e.g., A/D converters) and wiring occupy valuable silicon area, and thus designers tend to limit their numbers.

We propose to utilize the temperature measurements from the sparse thermal sensors for the purpose of post-silicon power mapping. Our technique for post-silicon power mapping involves capturing the thermal emissions from the back of the die and inverting the captured images to get power estimates. But this requires the full infrared imaging

setup. The high resolution infrared camera can be very expensive, making the cost of the imaging setup very high. In this subsection, we provide an alternative way for the post-silicon power mapping procedure. We describe the framework for post-silicon power mapping using thermal maps reconstructed from temperature readings of sparse thermal sensors. We use the thermal sensor measurements to perform full thermal characterization of the chip, and use the reconstructed full thermal map for post-silicon power mapping purposes. The major contributions of this section are as follows.

- We elaborate the frequency domain characterizations of thermal signals directly acquired from a real processor using a thermal infrared camera. The sparsity of these signals is analyzed in the frequency domain. We also explore different choices of frequency domain bases.
- Using the few measurements of thermal sensors, we propose a number of full signal reconstruction techniques that are capable of locating the sparse components of the temperature signal in the frequency domain and determining their magnitudes.
- We propose to use the full reconstructed thermal maps using sensor measurements for the purpose of post-silicon power mapping, giving an alternate way to the infrared imaging for post-silicon thermal and power characterization.
- In contrast to previous works which relied on simulators, we directly verify our full thermal characterization method on an operational 45-nm dual-core processors using an infrared camera. We elucidate the trade-off between the number of thermal sensors and the accuracy of full thermal characterization.
- We further utilize the various artificial power patterns created using a FPGA test chip in Chapter 3 and 4 for thermal reconstruction based post-silicon power mapping. In order to apply our thermal sensor based power mapping to a real processor, we utilize the real quad-core processor from Section 5.2.2. We extend the post-silicon

power mapping results from previous chapters, where instead of using the thermal maps obtained by infrared imaging, we use samples of the thermal traces as thermal sensor measurements and apply the proposed DC -based power mapping framework.

- We analyze the accuracy of power mapping as a function of the number of sensors and the spatial frequency of the power maps.

The organization of this section is as follows. Section 5.3.1 overviews some of the recent relevant methods in the literature. Section 5.3.2 introduces the frequency-domain as a method to analyze temperature signals. In Section 5.3.3 we propose a number of new, accurate methods for full thermal characterization. In Section 5.3.4, we propose techniques to utilize the full reconstructed thermal maps for post-silicon power mapping purposes. We verify our full thermal characterization methods through experimental results and evaluate the accuracy of the power mapping techniques with two different computing chips, FPGA and multi-core processor in Section 5.3.5. Finally, Section 5.3.6 summarizes the main results of this work.

5.3.1 Previous work

A number of recent papers discuss *design methods* that allocate sensors near potential hotspot locations, and *runtime methods* that estimate the *hotspot* temperature and *full chip* temperatures during runtime using the measurements of the thermal sensors [53, 58, 64]. For runtime hotspot and full thermal characterization, two techniques have been proposed. In the first technique, Long *et al.* [53] advocate using a grid-based interpolation scheme that identifies the hotspot around each sensor by interpolating the measurements at its immediate neighbors. In the second technique, Cochran *et al.* [18] advocated spectral techniques that are capable of fully reconstructing the temperature at all locations of the

processor die. The main idea of [18] is to regard the spatial temperature as a space-varying signal and to utilize the Nyquist-Shannon sampling theory to devise methods that can reconstruct the full thermal status from the measurements of the thermal sensors. None of the previous works in thermal characterization utilized the reconstructed thermal maps for post-silicon power mapping purposes.

5.3.2 Frequency Domain Techniques

Exact thermal estimation requires solving the partial differential heat equation, or a lumped first-order approximation for it, using as input the detailed power consumption of various processor units together with a model of the chip-package structure. This detailed information is not available during runtime, and furthermore, there is no sufficient computational resources to estimate the temperature in real time. Thus, we focus on devising computationally efficient *frequency domain* techniques that only use the measurements of the thermal sensors to fully characterize the processor's temperature. If T is the spatial temperature signal which is expressed as a $N \times 1$ vector, then the frequency domain representation of T can be expressed as

$$T = \Phi C, \text{or } C = \Phi^{\dagger} T \tag{5.10}$$

where Φ is an $N \times N$ matrix with columns that form a orthonormal basis, $\Phi^{\dagger}$ is the conjugate transpose of Φ, and $C = \{C_1, \ldots, C_N\}$ is an $N \times 1$ vector that has the frequency-domain coefficients of T. There are many choices for orthonormal bases. One choice for Φ is the 2-D Discrete Fourier Transform (DFT) matrix, where an element at row u and

column v of Φ is computed as,

$$\Phi_{uv} = \frac{1}{N \times N} e^{-2\pi i q_u q_v / N} e^{-2\pi i r_u r_v / N}, \tag{5.11}$$

where r_u and q_u are the remainder and quotient of dividing u by N respectively, and r_v and q_v are the remainder and quotient of dividing v by N respectively. Another choice is the 2-D Discrete Cosine Transform (DCT) matrix , where each element in the matrix is computed as,

$$\Phi_{uv} = \alpha_u \alpha_v \cos \frac{\pi(2q_v + 1)q_u}{2N} \cos \frac{\pi(2r_v + 1)r_u}{2N}, \tag{5.12}$$

where α_u and α_v are normalization factors. The biggest advantage of using frequency-domain techniques is that they transform a non-sparse signal, T in this case, to a sparse signal C with mostly zero coefficients. The sparsity of the spatial thermal signal in the frequency domain is confirmed as follows. Using our infrared camera, we capture the thermal maps of an AMD Athlon II dual-core processor during operation as shown in Figures 5.12.a and 5.12.c. We plot the the DCT transform of the two thermal maps in Figure 5.12.b and 5.12.d. Note that we only plot few frequency coefficients at low frequencies as the rest of the coefficients are very close to zero. Furthermore, Figure 5.12.d displays coefficients that are larger in magnitude and at a larger frequency ranges than the coefficients of Figure 5.12.b. These differences in the magnitudes and locations of the frequency domain coefficients are a quantitive metric for the visually apparent strong spatial gradients of the temperature signal of Figure 5.12.c in comparison to Figure 5.12.b.

Signal *energy* is an important concept in frequency-domain analysis, where it is de-

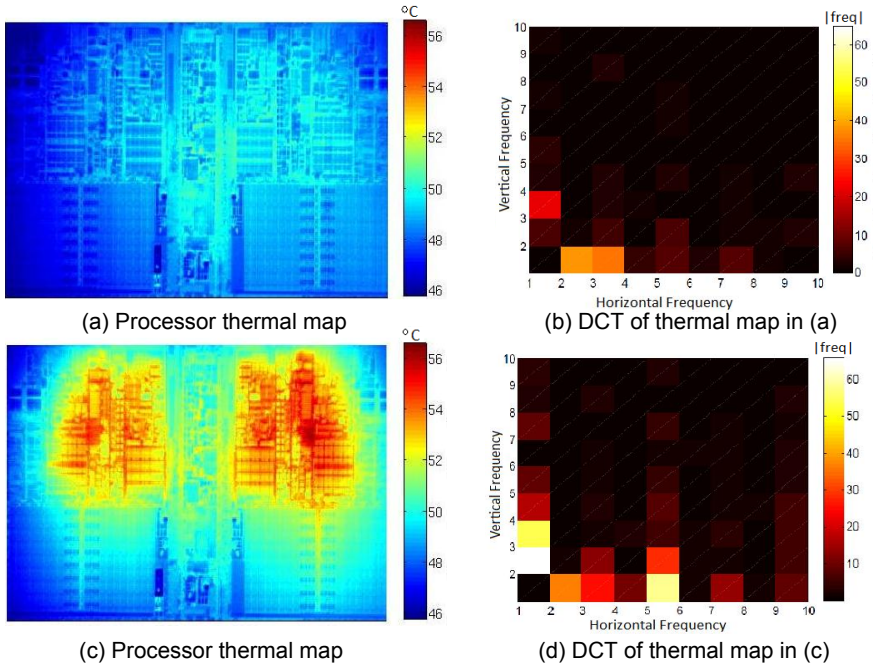

(a) Processor thermal map

(b) DCT of thermal map in (a)

(c) Processor thermal map

(d) DCT of thermal map in (c)

Figure 5.12: Thermal maps of Athlon dual II processor and their corresponding DCT frequency-domain representations.

fined as the sum of squares of the magnitudes of its coefficients in the frequency domain. The concept of energy gives a quantitive metric to compare various frequency-domain representations. We use this concept to confirm that DCT is a better basis than DFT for thermal signals. Using the thermal maps of Figure 5.12, we plot in Figure 5.13 the percentage of energy captured as a function of the number of frequency domain coefficients, for both the FFT and DCT bases. From the plot, it is evident that it is possible to capture most of the energy content of the thermal signal with only few coefficients. Furthermore if we compare DFT and DCT, we can see from the plot that the DCT captures the same energy as the DFT using a fewer number of coefficients. This result shows that DCT has more energy concentrated in fewer number of coefficients making it a sparser representation. Thus, we choose DCT as our orthogonal basis for further analyses.

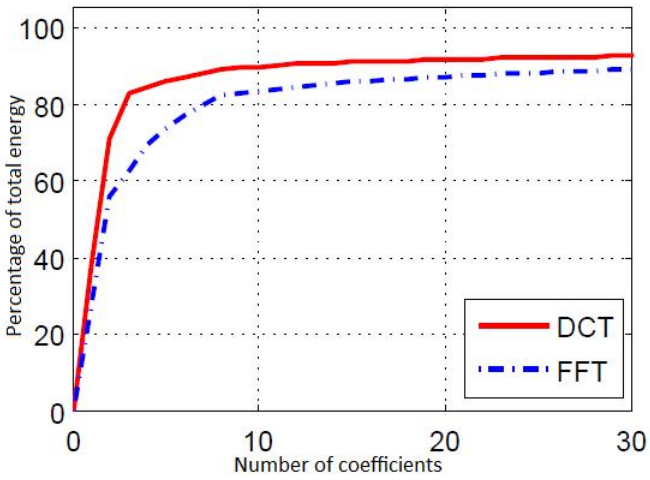

Figure 5.13: Fraction of signal energy captured as a function of the number of frequency-domain coefficients.

5.3.3 Proposed Full Runtime Thermal Characterization Techniques

During regular processor operation, the full temperature field $T = \{T_1, T_2, \ldots, T_N\}$ of the die at all N locations is not available; instead, only n samples obtained from the n thermal sensors are available. Let $y = \{i_1, i_2, \ldots, i_n\}$ denote the indices or locations of the n sensors in the temperature field, and $T(y)$ denote the temperature measurements at these locations. Then the objective of full thermal characterization is to estimate the temperature at each of the N points given the measurements of the n sensors. If the thermal signal is k-sparse in the frequency domain, and if the locations of the non-zero signal coefficients in the frequency domain are known, then let $s = \{j_1, j_2, \ldots, j_k\}$ denote the locations or indices of these non-zero coefficients. Given the sensor samples $T(y)$ and the orthonormal basis Φ, the signal samples $T(y)$ can be expressed

$$\begin{aligned} T(i_1) &= \Phi_{i_1,j_1} C_{j_1} + \Phi_{i_1,j_2} C_{j_2} + \cdots + \Phi_{i_1,j_k} C_{j_k} \\ \vdots &= \vdots \\ T(i_n) &= \Phi_{i_n,j_1} C_{j_1} + \Phi_{i_n,j_2} C_{j_2} + \cdots + \Phi_{i_n,j_k} C_{j_k} \end{aligned}$$

These set of equations can be written succinctly using matrix notation as

$$T(y) = \Phi(y, s)C(s), \tag{5.13}$$

where $\Phi(y, s)$ denote the matrix formed using Φ's rows with indices y and Φ's columns with indices s, and $C(s)$ is the vector formed of the C elements at indices s. In this case the best C that gives the total least square errors (LSE) of Equation (5.13) is given by

$$C_{LSE} = (\Phi(y, s)^{\dagger}\Phi(y, s))^{-1}\Phi(y, s)^{\dagger}T(y). \tag{5.14}$$

Full temperature characterization is achieved by multiplying the original basis set Φ with C_{LSE}; i.e., $T = \Phi C_{LSE}$. While Equation (5.14) gives a convenient closed-form to find the best frequency-domain representation, it is also possible to find the LSE solution of Equation (5.13) iteratively using gradient descend methods [5]. The advantage of gradient descend methods is that they allow a relatively smooth trade-off between computational runtime and solution accuracy. In all cases, solving Equation (5.13) hinges on the ability to determine the locations of the non-zero coefficients of C; that is, s must be determined. We propose two approaches to tackle this problem:

k-LSE using Pre-determined Thermal Characterization: In the first approach, we utilize our observation of Section 5.3.2 that the energy of the temperature signals acquired from real processors are mostly concentrated in the low frequency range. Consequently, the locations of the k non-zero coefficients can be picked from the low-frequency range. The number of coefficients picked depend on the number of available thermal sensors where $k < n$ to make sure that the solution to Equation (5.14) is stable. To prioritize low-frequency coefficients over high-frequency coefficients, we propose the order given

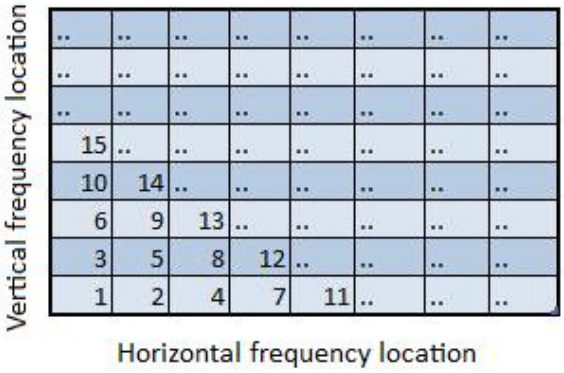

Figure 5.14: Order of coefficients for the k−LSE method.

in Figure 5.14. Essentially our proposed method picks coefficient locations to match the frequency-domain representations of thermal signals acquired from real processors.

Compressive Sensing: Compressive sensing techniques attempts to simultaneously find the locations and magnitudes of the non-zero components of the sparse signal [14]. We investigate two compressive sensing techniques. In the first technique (CS−L1_MIN), the following Second-Order Cone Programming (SOCP) formulation

$$\min \quad ||C||_1 \tag{5.15}$$

$$\text{subject to} \quad T(y) = \Phi(y,:) \times C, \tag{5.16}$$

is solved to minimize the ℓ_1 norm of C, where $\Phi(y,:)$ is the matrix formed from the full rows of the Φ matrix. Minimizing the ℓ_1 norm forces the SOCP solver to choose the solutions that are sparsest. For a real-time setting, solving a SOCP formulation can be unacceptable from an computational perspective. Thus, we explore a second technique (CS−STOMP) based on the greedy iterative procedure Stagewise Orthogonal Matching

Procedure: STOMP$(\cdot)$
Input: Initialize $r = T(y)$, C as a zero $N \times 1$ vector, and $s = \emptyset$.
Output: Full characterization of temperature signal.

While $|r|$ ¡ noise threshold

1. $C_r = \Phi(y, :)^{\dagger} r$
2. Find $p = \{j | C_r(j) > \max(C_r) -$ noise threshold$\}$
3. Let $s = s \cup p$
4. Compute $C(s) = (\Phi(y, s)^{\dagger} \Phi(y, s))^{-1} \Phi(y, s)^{\dagger} T(y)$
5. Let $r = T(y) - \Phi(y, :) C$

Return $T = \Phi C$

Figure 5.15: Procedure STOMP $(\cdot)$ for full thermal characterization using compressive sensing.

Pursuit (StOMP) [22]. The StOMP algorithm transforms the sampled signal $T(y)$ into a negligible residual by identifying the significant non-zero components from the signal in the frequency domain one at a time in a greedy fashion as given in Figure 5.15. In Step 1, the algorithm computes the frequency domain representation C_r of the residual signal (the residual signal r is initialized to $T(y)$ at the beginning of the algorithm), and then in Step 2, it determines the locations of non-zero coefficients from C_r that are above a certain noise threshold. Using the selected coefficient locations, the algorithm then solves the least squares estimation formulation in Step 4 to determine the magnitude of these selected coefficients. Finally, it reconstructs the signal in Step 5 and subtracts it from the original sampled signal to produce a new residual. The algorithm is iterated until the residual goes close to zero and all the significant non-zero components in the signal are recovered.

A fundamental assumption for both the compressive sensing techniques is that the

sampling is performed randomly. In real processors thermal sensors are placed at strategic locations near hot spots, and consequently, the theoretical attractiveness of compressive sensing is no longer guaranteed.

5.3.4 Post-silicon Power Mapping using Reconstructed Thermal maps

We propose to perform post-silicon power mapping with full thermal maps reconstructed using frequency domain techniques as described in Section 5.3.3. The procedure for post-silicon power mapping is similar to the framework described in Chapter 3, where we find the best power map $\mathbf{p}$ that minimizes the total squared error between the true temperatures $\mathbf{T}$ and the measured temperatures; that is, $\min ||\mathbf{Rp} - \mathbf{T}||_2^2$. In contrast to previous techniques, the thermal map $\mathbf{T}$ is obtained by reconstruction using measurements from thermal sensors, rather than the commonly used infrared imaging. The matrix R is estimated either based on measurements as described in Chapter 3 or based on equivalent finite-element simulations as described in Chapter 5.2.

For the thermal to power inversion, the following optimization formulation is used, where $\mathbf{T} = \Phi C_{LSE}$; Φ is the original basis formed for the full thermal characterization and C_{LSE} are the estimated frequency domain coefficients from Equation (5.14).

$$\mathbf{p} = \arg_{\mathbf{p}} \min ||\mathbf{Rp} - \Phi C_{LSE}||_2^2 , \tag{5.17}$$

with following constraints,

$$||\mathbf{p}||_1 \leq p_{total} + tol, \tag{5.18}$$

$$||\mathbf{p}||_1 \geq p_{total} - tol, \text{and} \tag{5.19}$$

$$\mathbf{p} \geq \mathbf{0} \tag{5.20}$$

where $\mathbf{p}$ is the power map using reconstructed thermal map, $\mathbf{R}$ is the thermal resistance matrix, and ΦC_{LSE} is the reconstructed thermal map. $||\cdot||_1$ is the ℓ_1 norm and p_{total} is the total power consumption of the chip which could be measured externally using a digital multimeter. Note that, we do not use a regularization parameter for the thermal to power inversion as in Chapter 3 formulation, which was used to mitigate the thermal noise present in infrared images. In this case, the reconstructed thermal images will not have measurement noise as in the case of infrared imaging. We present experimental results from two different computing substrates, FPGA and multi-core processor as follows.

5.3.5 Experimental Results for Full Thermal Characterization and Power Mapping

To test the effectiveness of our post-silicon power mapping using full thermal characterization methods, we use the following experimental setup:

- For thermal imaging, we utilize 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 100 Hz with a spatial resolution of 30 μm with a 0.5$\times$ microscopy kit. Details of the experimental setup is given in Chapter 3.

- For full thermal characterization, we use a real dual-core 45 nm AMD Athlon II X2 240 processor running at 2.1 GHz. The processor is cooled with a oil based cooling system as described in Chapter 2. To collect the thermal traces involved in the experiments we use the SPEC CPU2006 benchmark set which has 29 applications. We collect 100 temperatures traces obtained from the infrared camera after executing each benchmark individually and after executing two benchmarks at a time.

- For validating the proposed thermal reconstruction based post-silicon power mapping, we utilize the FPGA test chip described in Chapter 3, which allows us to create artificial patterns with different frequencies. We further apply the thermal reconstruction based techniques on a real quad-core using SPEC CPU2006 benchmarks described in Section 5.2.2.

a. Full thermal characterization results

The objective of the first experiment is to demonstrate the effectiveness of our frequency-domain techniques in fully reconstructing the thermal signal. We report the following metric in our experiment:

- *Full thermal characterization error.* For each temperature trace, we compute the average absolute error between the true temperatures as measured by the infrared camera and as estimated by our signal reconstruction methods. We normalize the average error by the difference between the maximum and minimum temperatures in the trace. This normalization helps put the characterization error in perspective: a 0.5 °C error is more significant if the difference between maximum and minimum temperatures is 5 °C rather than 20 °C. We report the average absolute error computed for all 100 temperature traces.

In this experiment we compute the full thermal characterization error as a function of the number of sensors which are placed randomly. We compare here four full thermal characterization methods: the `spectral` method [18], the proposed `k-LSE` method, the `CS-L1_MIN` method, and the `CS-STOMP` method. The performance of these methods as a function of the number of sensors is given in Figure 5.16. The results show that the proposed `k-LSE` method gives superior results compared to other methods. Note that `CS-L1_MIN` and `CS-STOMP` fail to produce meaningful results when the number of sensors is relatively low. The main reason for the poor performance of generic compressive sensing techniques is that they attempt to compute both the locations and magnitudes of the frequency-domain coefficients, and hence there is a chance they end up picking wrong high-frequency coefficients. In contrast our proposed `k-LSE` method is devised with the nature of thermal characterizations encountered in real processors in mind, and hence it is more powerful than a generic technique. Our method `k-LSE` also outperforms the `spectral` method ([18]) as it uses a sparser basis (DCT instead of FFT) and a more accurate direct algebraic approach that guarantees minimizing the total square error.

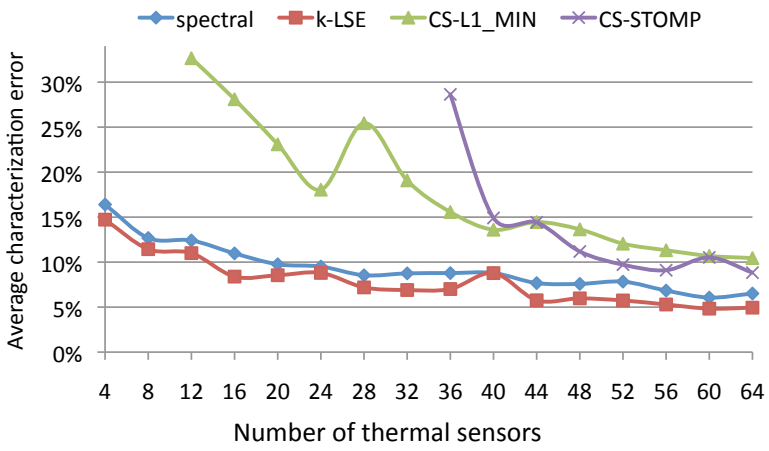

Figure 5.16: Average error in full thermal characterization using various temperature reconstruction methods.

b. FPGA power mapping results

We have used 12 different power patterns created in a FPGA test chip as described in Section 3.4 to test our proposed technique. Figure 5.17.a shows six different arbitrary injected power patterns, Figure 5.17.b shows the infrared thermal maps, Figure 5.17.c shows the sampled temperature values at 36 uniform thermal sensor locations, and Figure 5.17.d shows the reconstructed thermal maps using techniques described in Section 5.3.3. We chose `k-LSE` as our reconstruction method, because it gave us the best reconstruction error as shown in the experimental results in Section 5.3.5.a.

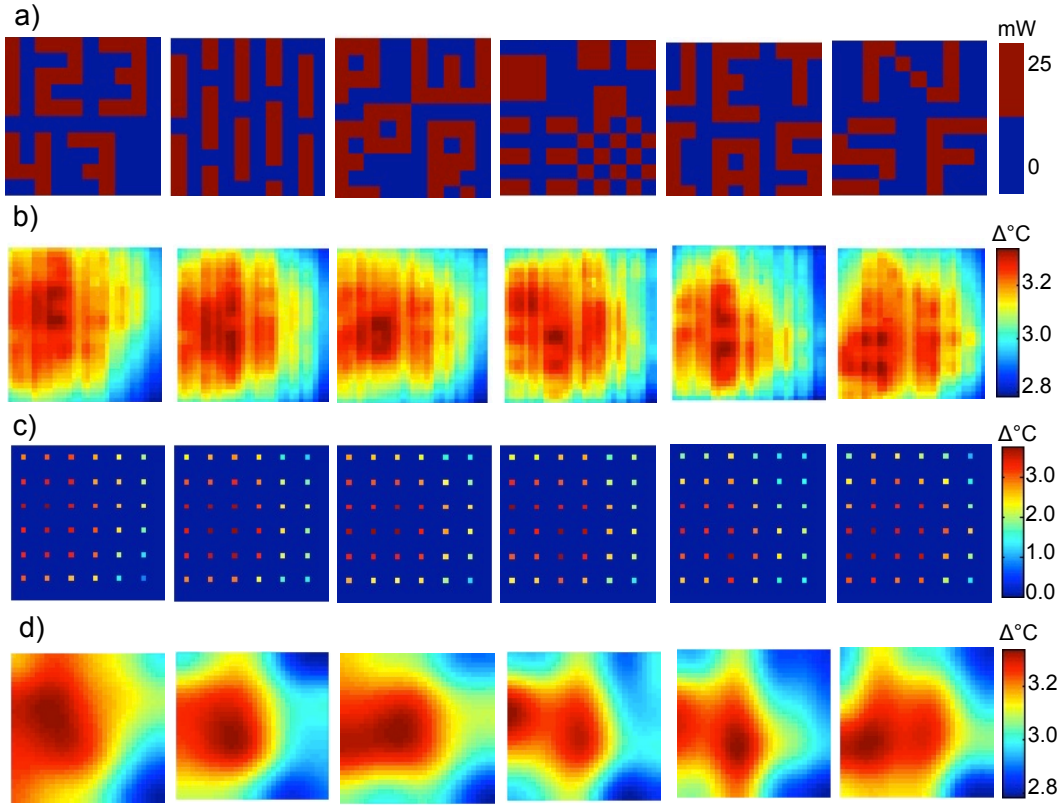

Figure 5.17: Reconstructed thermal maps of arbitrary power maps using thermal emissions. (a) Injected power patterns, (b) Resultant temperature measurements by infrared imaging, (c) Sampled temperature values at uniform sensor locations (sensor sizes have been magnified for visibility), and (d) Reconstructed thermal maps using measurements of 36 sensors.

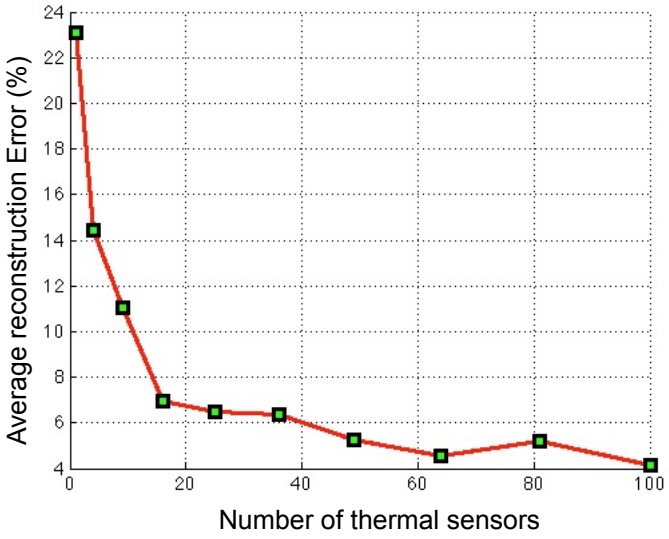

Figure 5.18: Thermal reconstruction error vs. number of thermal sensors.

We define the average reconstruction error as the mean of the absolute error between the full characterization results as computed by the `k-LSE` method and the true thermal map as given by the thermal camera. Figure 5.18 shows the average reconstruction error of all 12 cases as the number of thermal sensors increases. We observe that the reconstruction error drastically decreases as the number of sensors goes from 1 to 16. For number of thermal sensors of 16, the reconstruction error is less than 8%. In modern multicore processor with many cores, increasing number of thermal sensors are becoming available, which can be used for the full thermal characterization purposes.

We define the power mapping error as follows,

$$Error = \frac{\sum_k |\mathbf{p_k} - Pcorrect_k|}{\sum_k Pcorrect_k} \tag{5.21}$$

where $Pcorrect$ is the golden power map and $\mathbf{p_k}$ is the value of the k^{th} element in the vector $\mathbf{p}$. We estimate power mapping error for all 12 patterns. Figure 5.19 shows the

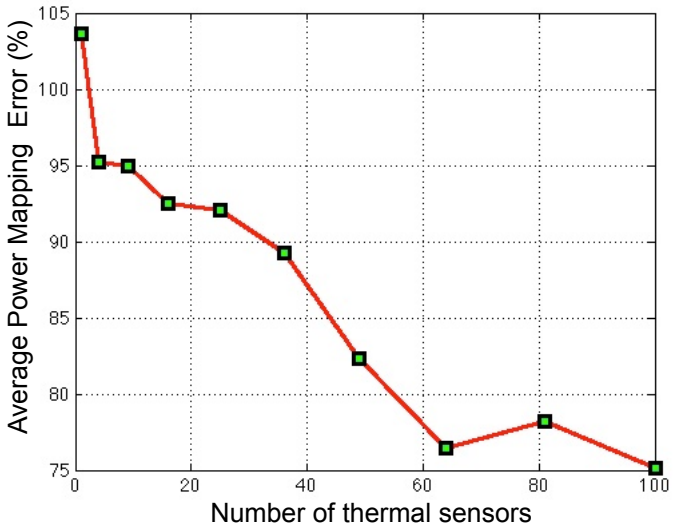

Figure 5.19: Average power mapping error (%) for 12 patterns.

average percentage error for 12 patterns. The reason for very high error is that the spatial frequency of the power patterns are intentionally created very high, which creates a twofold effect on the power mapping error. The `k-LSE` reconstruction method picks the locations of the non-zero coefficients from the low-frequency range as discussed in Section 5.3.3. The number of coefficients picked has to be less than the number of thermal sensors to make sure that the solution is stable. So, more higher frequencies are omitted from the thermal maps, which makes it difficult to reconstruct the thermal maps with high frequency. The error in the thermal reconstruction is propagated to the power mapping error. In addition to that, as discussed earlier in Chapter 3, the nature of heat conduction on chips leads to a low-pass filtering effect, which makes the thermal to power inversion harder, which in turn increases the power mapping error.

We assess the accuracy of our proposed power mapping methodology as a function of the spatial frequency of power maps. To assess the effect of the low-pass filtering, we created checker-board power maps of increasing spatial frequencies. Figure 5.20.a shows the golden spatial power maps of the checker-board patterns, Figure 5.20.b shows their infrared thermal maps, Figure 5.20.c shows the sampled temperature measurements at 36

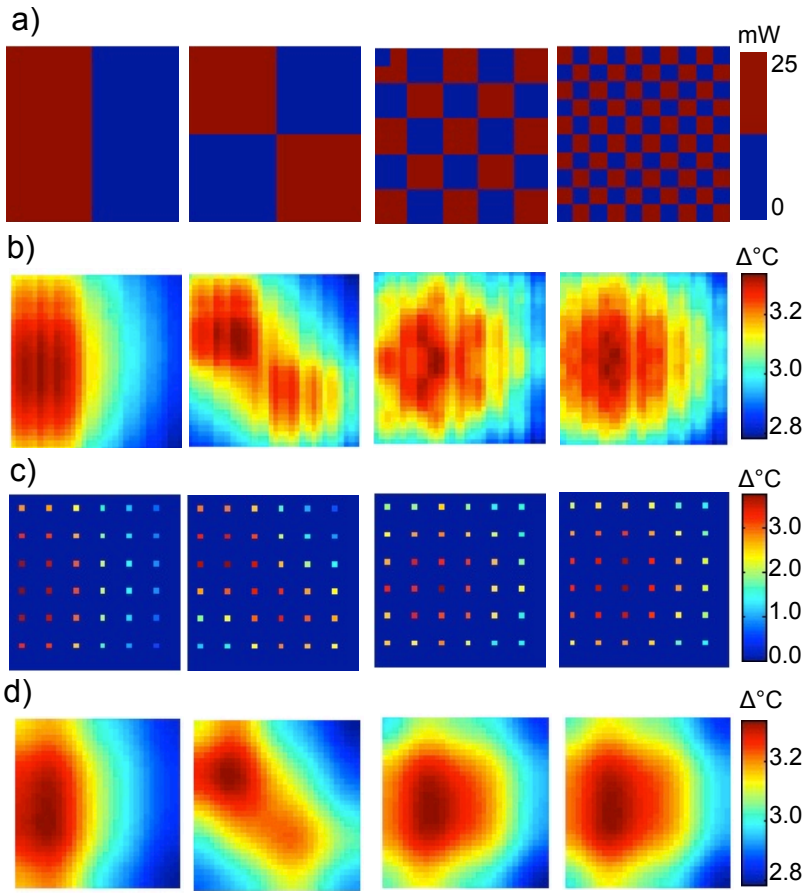

Figure 5.20: Reconstructed thermal maps of power maps with increasing spatial frequency. (a) Injected power patterns, (b) Resultant temperature measurements by infrared imaging, (c) Sampled temperature values at uniform sensor locations (sensor sizes have been magnified for visibility), and (d) Reconstructed thermal maps using values from 36 sensors.

uniform thermal sensor locations, and Figure 5.20.d shows the reconstructed thermal maps using our `k-LSE` full thermal characterization methods. We can see the trend of increasing spatial frequency in the power patterns, where the fourth pattern has the highest frequency.

To quantify the signal energy for each power pattern shown in Figure 5.20, we first perform 2-D Fast Fourier Transform of the power patterns. Then we compute the energy in coefficients with frequencies $\geq$7 for each golden power map, and divide the result by the total energy to get corresponding signal energy for all four patterns. In Figure 5.21, we plot

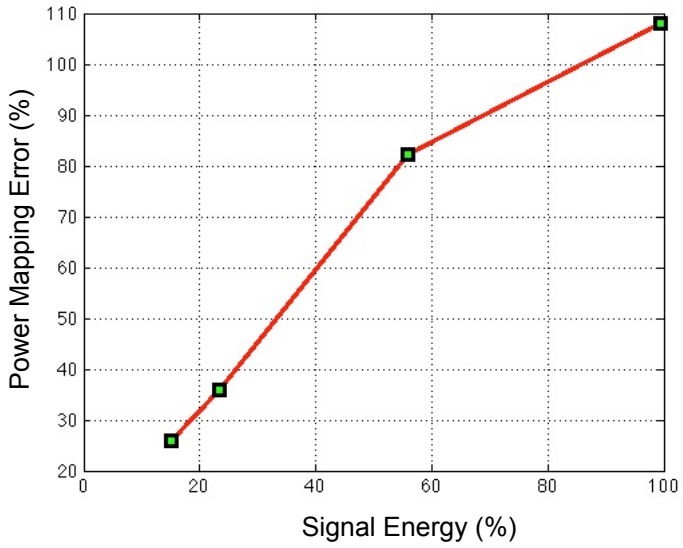

Figure 5.21: Power mapping error vs. spatial frequency of checker patterns.

the power mapping accuracy of each thermal map from Figure 5.20.d against the computed signal energy. This result confirms our earlier results in Chapter 3 that concluded that increasing the spatial frequencies of power maps can lead to a deterioration in the accuracy of thermal to power inversion. In this case, increasing the spatial frequency, leads to more low pass filtering effect and loss of information. It is more difficult to reconstruct thermal maps using sparse thermal sensor measurements as the spatial frequency of the thermal maps increases. As a result, our proposed power mapping accuracy confirms this trend.

c. Real Processor Power Mapping results

For implementing the post-silicon power mapping with reconstructed thermal map for real processor, we use AMD quad-core processor as in Section 5.2.2. We do thermal reconstruction using our full thermal characterization methods for the ten different workloads from Table 5.2. Figure 5.22 shows four of the ten cases using 36 sensor measurements for thermal reconstruction. Figure 5.22.a shows the infrared thermal map, and 5.22.b shows the reconstructed thermal maps. The thermal maps are fairly accurately reconstructed.

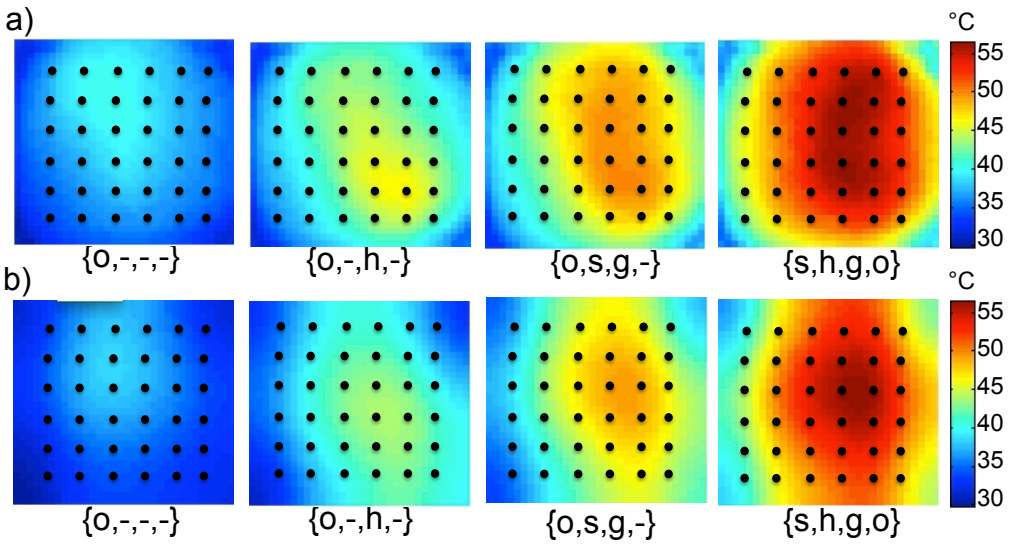

Figure 5.22: a) Infrared thermal maps and b) Reconstructed thermal maps using 36 sensors; the black dots represent the sensor locations, sizes of sensors have been magnified compare to the chip size for visibility.

The average thermal reconstruction error over ten different workload cases are shown in Figure 5.23.

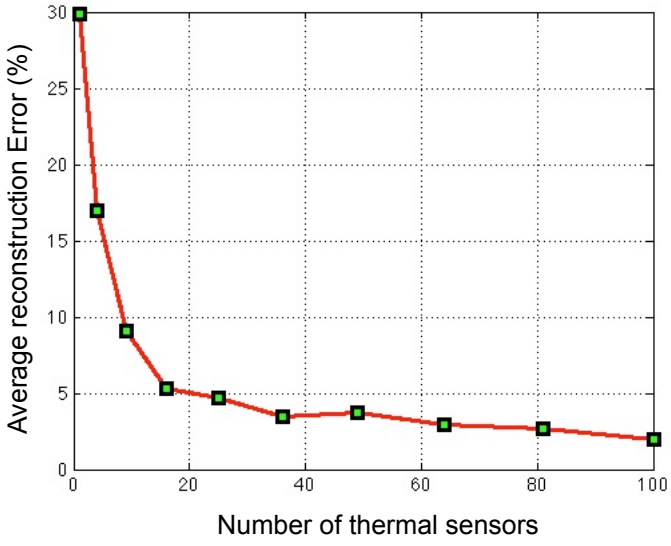

Figure 5.23: Average thermal reconstruction error over ten cases vs. number of thermal sensors.

We perform thermal to power inversion for the reconstructed thermal maps using formulation of Equation 5.17 with corresponding total power constraints. We compute the error in the power map using Equation 5.21, where $\mathbf{p}$ is the power map estimated using reconstructed thermal maps and the $Pcorrect$ is the power map estimated using infrared based thermal maps in Section 5.2.2. That is, the estimated power maps with block powers in Section 5.2.2 are considered as the golden patterns for this experiment. Figure 5.24 shows the average error over the ten cases with respect to number of thermal sensor measurements used for the reconstruction. The average error of power mapping for the real processor as shown in Figure 5.24 is much less than the manually created power patterns in the FPGA as shown in Figure 5.19. This is because the spatial frequency is less in the real processor power maps which improves the power mapping results. This shows that in a practical scenario our thermal sensor based power mapping will maintain reasonable accuracy compare to the infrared based power mapping.

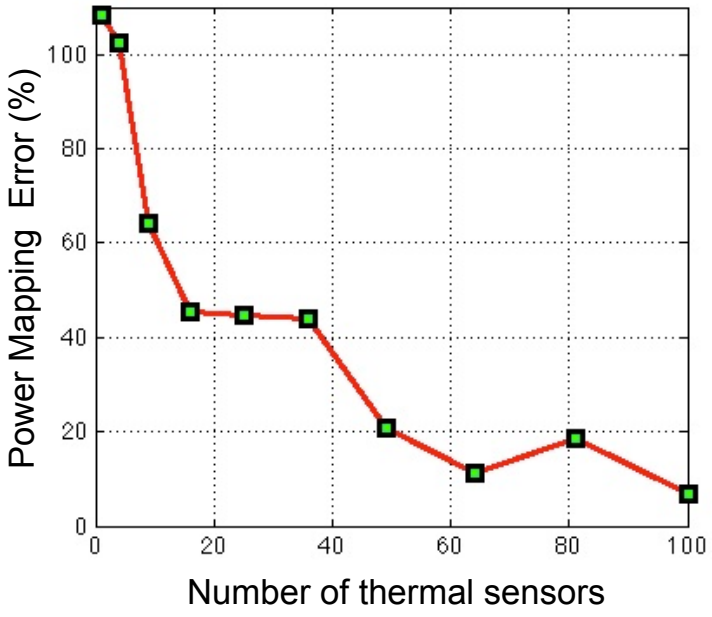

Figure 5.24: Average power mapping error for ten cases using AMD quad-core processor.

5.3.6 Conclusions

We presented a new methodology for thermal monitoring techniques for real processors and proposed runtime full thermal characterization techniques. We proposed frequency domain representations based on the DCT basis which achieves better results than the traditional FFT basis. We characterized the DCT representations of spatial thermal signals of a real processor, and utilized this characterization to devise a number of effective full thermal estimation methods such as `k-LSE`, `CS-L1_MIN`, and the `CS-STOMP`, which are designed for fine-grain thermal management techniques. Furthermore, we utilized the reconstructed full thermal maps for our proposed post-silicon power characterization, which gives an alternative to infrared imaging techniques. Using a sophisticated experimental setup with a state-of-the-art thermal infrared camera, we demonstrated the superiority of our techniques using a real 45-nm dual-core processor, a FPGA test chip with several artificially created power patterns and a multi-core processor with real applications.

5.4 Hardware Trojan Detection using Post-silicon Power Maps

5.4.1 Introduction

Globalization of the semiconductor design and fabrication process due to the ever-increasing cost of manufacturing in small-scale CMOS technology has caused imminent threat to the security of integrated circuits. Besides foundry practices, modern IC design often use intellectual properties (IP) cores and electronic design automation (EDA) software tools, which are supplied by third party vendors. While the practice saves cost by utilizing the economy of scale, involvement of third-party entities exposes the chips by authentic designers to threats including hardware malware (Trojan) insertion, unlicensed IP handling, and IP piracy [96, 23, 97]. Since ICs form the core for the computing and communication systems used in contemporary personal, commercial, and government affairs, their exposure endangers the full systems built upon them. Therefore, developing non-invasive methods for screening and interrogating ICs for maintaining integrity in presence of unreliable third-party fabrication has become essential.

We devise a novel methodology for Trojan detection using the post-silicon spatial thermal and power characterization framework proposed in Chapter 3 and 4. Chips can be thermally characterized using infrared emissions from the backside of silicon die, which then can be processed to get detailed spatial power maps [17, 33, 59]. These detailed thermal and power maps provide a much higher resolution Trojan detection method than previous methods where the total current is measured and converted to gate-level. This detection procedure is easily scalable because the chip's spatial view in thermal mapping is not limited by the size of the chip and is only dependent on the thermal mapping resolution.

The major contributions of this work can be summarized as follows.

- We propose a new direction for Hardware Trojan Detection using spatial thermal maps and inverted power maps described in Chapters 3 and 4 to detect and locate IC Trojans. Our detection framework involves acquiring post-silicon runtime thermal maps and applying residual inversion methods with l_1 regularization to obtain sparse spatial power maps which results in high sensitivity Trojan detection.

- We employ two-dimensional principal component analysis (2DPCA) in order to tackle high dimensional thermal and power maps. Utilizing the 2DPCA framework, we present two different approaches to Trojan detection, first is the *supervised thresholding* method which needs training data set, and the second is the *unsupervised clustering* method, which does not require any training data set.

- To create realistic chips, we add 20-40% process variations (PV) to gate lengths, widths and oxide thickness which can hide Trojans. To cover a wide range of variations, in our experiment we set five different PV levels with different standard variances which are obtained from realistic spatial variability models. We also add Gaussian noise to our thermal maps to mimic real noise in infrared measurements.

- We design virtual Trojans with power consumption varying from 0.05% to 0.2% of total IC power consumption. To evaluate the accuracy of our Trojan localization method, we place the virtual Trojans in ten different locations in the chip.

- We present an extensive set of simulation results with four different benchmarks with realistic chips and very small Trojan sizes. We show that our proposed methods are able to detect and locate Trojans with power consumption as small as 29.6 μW very efficiently and accurately. We also evaluate the impact of thermal noise and chip voltage on the Trojan detection accuracy.

The organization of this section is as follows. Section 5.4.2 provides the necessary background on Trojan detection. In Section 5.4.3 we outline the thermal and power framework for our proposed Trojan detection procedure and in Section 5.4.4 we describe our 2DPCA analysis. Section 5.4.5 describes various Trojan detection methods and Section 5.4.6 describes our localization procedure. In Section 5.4.7 we discuss the impact of noise in thermal maps and measures to improve detection results. In Section 5.4.8 we present our experimental setup and present our results to demonstrate the effectiveness of our approach, and finally, Section 5.4.9 summarizes our main results.

5.4.2 Background on Trojan detection

Hardware Trojans are implemented by unsought chip modifications by traitorously changing or tampering with the chips to provide opportunities for later exploits including controlling, monitoring, or spying the chip contents or secret keys [96, 97, 106]. Trojans can be very hard to detect, since they may be often inactive, only triggered as needed in target time intervals. Due to the increasing complexity of the contemporary chips and lack of controllability/observability to the chip internals post-silicon, the traditional structural and functional tests are becoming ineffective in targeting Trojans. Invasive reverse engineering methods are slow, destructive, and expensive. Thus, devising noninvasive methods for examining the ICs and detecting Trojans has been recognized as a challenging research problem.

Reports of instances of malware in military chips have triggered research and investigations into the Trojan detection problem [23]. The utilized tests for Trojan detection include current-based methods, using static or dynamic currents [80, 79, 76, 100], delay-based approaches [43, 51, 67], as well as simultaneous consideration of various current and delay testing methods [48]. In current-based approaches, both regional testing of cur-

rent sums [80, 79], and translation of the currents to the gate-level [76, 100, 48] were pursued. While current-based methods can potentially provide a good characterization on smaller circuits, the presently available methods either need additional probes to the chip for regional current measurements [80], or necessitate formation and solving very large system of equations that are highly sensitive to noise and process variation [76, 100, 48]. Delay-based detection methods have less components on each path and are easier to scale, but they suffer from the known problem of inadequacy of external test vectors for sensitizing all possible paths.

One of the early work in this area [4] utilized the dynamic current (power) measurements by destructive testing of a few ICs from the design to build signatures. The assumption was that the fingerprint did not contain any malware. The existence of Trojan(s) in other chips were verified by noninvasively comparing against the signatures formed by destructive testing. Another path taken early for Trojan detection was to use verification and functional testing methods. This approach simulates the inputs and then checks the corresponding outputs for the desired patterns [102, 6]. Functional testing suffers from the state-space explosion and lack of targeted verification output (since the Trojan behavior is not known in advance). Therefore, its scope and effectiveness is rather limited.

The typical assumption for current- and delay-based Trojan detection approaches is that a golden model of the chip can be formed by post-layout simulations. The structural properties of the manufactured chips under investigation are then compared with this model. Such detection becomes more challenging for newer technology nodes with surging random process variation which makes it hard to distinguish Trojan effect. What further complicates the problem is the large space of possibilities for Trojan exploit type and location.

An effective set of techniques pursued in this category is gate-level characterization

[99, 76, 100, 48] which works well with both delay and current measurements. This method measures the chip's delay or current for a number of test vectors. Assuming that the currents (delays) linearly add up, a linear system is then constructed from the measurement set. Solving the system of linear equations translates the side-channel characteristics to smaller gate-level structural properties. While effective, this suite of techniques does not perform well for larger chips with more gates, higher accumulated measurement noise, and more sophisticated process variation models. The existing gate-level characterization and Trojan detection techniques are evaluated on smaller benchmarks, where the performance of the method is the best. Note that, approaches based on regional testing of accumulated current which have a higher resolution only work for certain types of packaging and measurement probes [80, 79].

To mitigate the above discussed challenges, we propose to use the post-silicon power mapping framework presented in Chapter 3 to obtain very high resolution thermal and power maps. Our proposed method utilizes these very high resolution thermal and power maps in order to detect IC Trojans which results in a very high sensitivity Trojan detection technique.

5.4.3 Proposed Multimodal Trojan Detection Framework

Hardware Trojan detection is the process of detecting chips that are infected with unwanted Trojan circuitry and verifying the trustworthiness of the manufactured chips upon return to the clients. This new step requires defining a post-manufacturing step to validate the chips conformance with the original specifications, which is called *silicon design authentication*. In this work, we propose an entirely new multimodal framework for post-silicon Trojan detection using the thermal and power maps of the ICs running practical benchmarks.

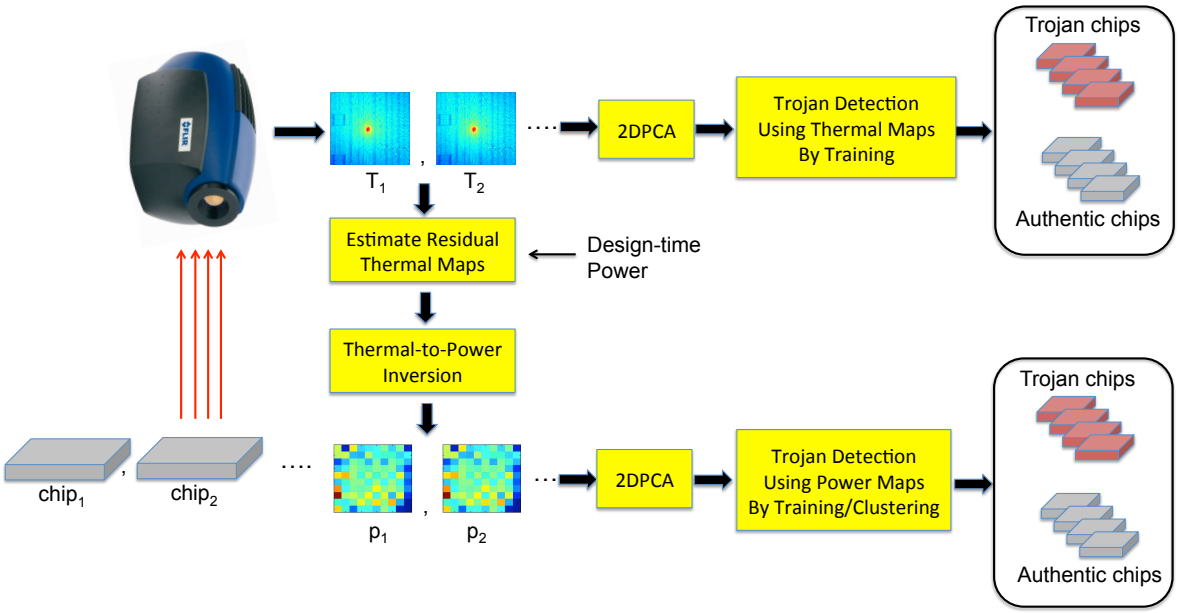

Figure 5.25: Proposed Trojan detection framework using thermal and power maps.

Figure 5.25 shows the framework of the proposed multimodal Trojan detection methods using post-silicon thermal and power characterization. The individual test modes are based on the chip temperature maps and power maps respectively. In the beginning, workloads or test patterns are applied to the integrated chips and the runtime steady-state or averaged infrared thermal maps $\mathbf{T}_1, \mathbf{T}_2, ...$ are collected under realistic loading conditions. Our purely thermal map based Trojan detection method is able to detect and locate very small size Trojans. To further increase the Trojan detection resolution, we propose to invert the thermal maps to accurately estimate detailed spatial power maps, which corresponds to our second mode and can be used to perform power map based Trojan detection.

We propose two Trojan detection techniques involving two-dimensional Principal Component Analysis (2DPCA) using the characterized thermal and power maps. Depending on the the availability of data from prior tests (training data) of chips that known to be benign, either of the two Trojan detection methods can be used on the residual power maps. The first requires a set of trainings chips to classify the Trojan infected chips and uses

thresholding techniques. If no chips for training are available, then the second technique using unsupervised clustering can be applied. Note that we perform *unsupervised clustering* only with the inverted power maps, but not with the thermal maps. Because the natural clusters created by the features of the thermal maps are not properly distinguishable, the *unsupervised clustering* method is not effective in case of the thermal map based method. We describe the details of our proposed framework components in the rest of this section.

a. Thermal Mode

In the proposed procedure, infrared imaging is used to obtain thermal maps of post-silicon chips for Trojan detection. Modern integrated circuits use flip-chip packaging, where the die is flipped over and soldered to the package substrate. By removing the package heat spreader, one can obtain optical access to every device on the die through the silicon backside. So, we can obtain optical access to the die under test through the silicon backside by removing the packages's heat spreader. Silicon is transparent in the infrared spectral region and this transparency allows the capturing of thermal infrared emissions using infrared imaging techniques [59, 17, 84, 33]. An infrared-transparent heat sink, for example, silicon window based heat sink with mineral oil, has to be used to remove heat during operation of the IC [82]. The details of infrared imaging are given in Chapters 2 and 3.

For real chips, workloads of steady nature typically takes around tens of seconds to reach the steady state. The thermal maps need to be captured after the chip temperature had reached steady state. Some workloads might not have a steady nature, in that case, the thermal maps can be captured for 30s and averaged over time.

For the purpose of this work, we first apply random vectors to the ICs and get the

estimated power trace of each block by Primetime-PX. We then use HotSpot [39] thermal simulation tools to create the steady state thermal maps of various test bench circuits as described in Section 5.4.8. We denote the steady-state thermal maps obtained using design-time simulations of the original authentic chip by $\mathbf{A}_1, \mathbf{A}_2, ...$ for each benchmark. We perform Monte-Carlo simulations of the original chip at various PV corners to get power consumption under various process variation scenarios. By using Hotspot thermal simulation we produce the thermal maps for chips under test, which is represented by $\mathbf{T}_1, \mathbf{T}_2, ...$ for each benchmark. In a real scenario, these thermal maps can be obtained by infrared imaging from the backside of the die. It is possible to use the thermal maps for Trojan detection, but the sensitivity is less than the Trojan detection using power maps. If power mapping of the thermal maps is not available, these thermal maps can be used for Trojan detection as described in Section 5.4.4. We use authentic thermal maps $\mathbf{A}_1, \mathbf{A}_2, ...$ as the training set and perform our Trojan detection methods of 2DPCA on the thermal maps under tests $\mathbf{T}_1, \mathbf{T}_2, ...$ for Trojan detection as described in Section 5.4.4.

b. Power Mode

This section describes the power characterization of the chip. The power is obtained by inverting the thermal maps using quadratic optimization framework. The chip power and temperature are related by the heat equation, which can be discretized as follows by linear matrix formulation as described by Equation (3.5) in Chapter 3,

$$\mathbf{Rp + e = T}, \tag{5.22}$$

where $\mathbf{T}$ is the thermal map that gives the measured temperatures at every pixel of the imaging system. The continuous power signal is represented by a vector **p** that gives the

Procedure: Thermal to power inversion method
Input: Design time minimum power $\mathbf{p}_{min}$, Thermal maps under test **T**, Thermal resistance matrix **R**
Output: Residual power map $\mathbf{p}_r$

Find design time minimum thermal map, $\mathbf{T}_{min} = \mathbf{R}\mathbf{p}_{min}$;

Find residual thermal map, $\mathbf{T}_r = \mathbf{T} - \mathbf{T}_{min}$;

Solve quadratic programming: $\min \|\mathbf{R}\mathbf{p}_r - \mathbf{T}_r\|_2 + \|\mathbf{p}_r\|_1$, such that $\mathbf{p}_r \geq 0$.

Return solution of the quadratic programming $\mathbf{p}_r$

Figure 5.26: Thermal to power inversion methodology.

power density at a set of discrete die locations and the vector **e** denotes measurement noise in the infrared imaging system. The matrix **R** represents the thermal resistivities between different locations and is called the modeling matrix. The formulation of the matrix **R** is given in detail in the following works [17, 84]. For each specific chip, the matrix **R** can be estimated either by analytical methods, by simulation or experimentally on the real chip. We create matrix **R** by HotSpot simulation, by dividing the chip into 10×10 blocks, and exciting each block at a time. Thermal map corresponding to one excited block represents one column in the matrix **R**, this way we estimate matrix **R** for each chip column by column basis. The lower bound of the block size is limited by the precision of infrared camera. The minimum resolution of a midwave infrared camera is 5 μm. Detection accuracy increases as the block size decreases. There is a trade-off between the size of the blocks and computation time because as the block size decreases, the number of blocks increases, and Hotspot simulation time increases. Here we make a trade-off between the resolution and accuracy based on our experiments.

Given the thermal map vector **T** and matrix **R**, the objective is to find the best power map vector **p** that minimizes the total squared error between the temperatures as computed

from the estimated power $\mathbf{p}$ and the thermal measurements. For our case, we first subtract the thermal maps $\mathbf{T}_{min}$ corresponding to minimum estimated design time power $\mathbf{p}_{min}$, from the thermal maps $\mathbf{T}$ of chips under test, where $\mathbf{T}_{min} = \mathbf{R}\mathbf{p}_{min}$, and then invert the residual thermal maps, $\mathbf{T}_r$ to get the residual power estimates $\mathbf{p}_r$. We want the estimates to be of the shape that only the blocks affected by Trojan have non-zero values while all other blocks remain zeros, which naturally leads us to finding sparse solution. Therefore, we add a regularization term in our quadratic programming to minimize the ℓ_1 norm of the power map, that is to minimize $||\mathbf{p}_r||_1$. The thermal to power inversion methodology is summarized in the algorithm given in Figure 5.26. We apply our detection technique described in the following sections on the residual power maps.

We show an example of a thermal and power maps running workloads in Advanced Encryption Standard (AES) cipher chip with 40% process variation and 59.3 μW Trojan in Figure 5.27, which are used for Trojan detection. Figure 5.27(a) shows thermal map generated by HotSpot, Figure 5.27(b) shows the residual thermal map after subtracting the design-time minimum thermal map. We divide the chip into 10x10 blocks and estimate the residual spatial power maps using optimization formulation. The chip dimension is 163 $\times$ 163 μm^2 and each block size is 265.7 μm^2. Figure 5.27(c) shows residual power map estimated with previous thermal to power inversion method [37] and 5.27(d) shows

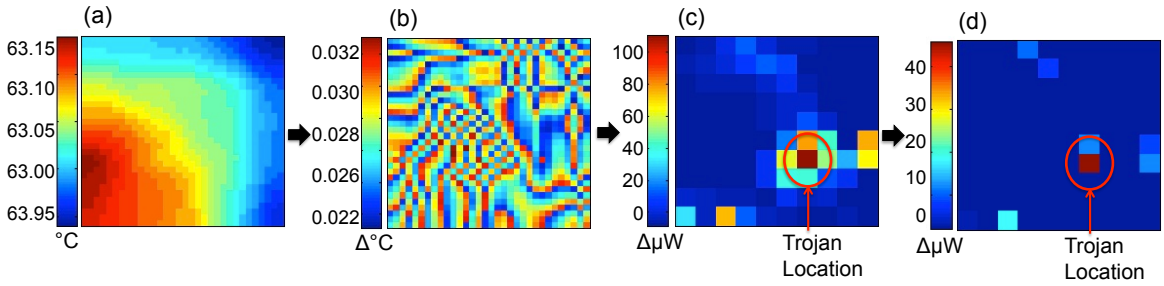

Figure 5.27: AES cipher thermal map (°C) and estimated residual power map (μW) a) Thermal map with Trojan, b) Residual thermal map, c) Residual power map without ℓ_1 regularization, and d) Residual power map with ℓ_1 regularization.

the estimated residual power map using the proposed optimization formulation with ℓ_1 regularization. The Trojan location is shown in both the power maps. We can see that the power map become more sparse after using ℓ_1 regularization, which makes it easier to detect and locate the Trojan.

5.4.4 Two-dimensional Principal Component Analysis

Principal Component Analysis (PCA) is a classical feature extraction and data representation technique widely used in the areas of pattern recognition and computer vision. PCA is mathematically defined as an orthogonal linear transformation that translates the data to a new coordinate system such that the greatest variance by any projection of the data comes to lie on the first coordinate (called the first principal component), the second greatest variance on the second coordinate, and so on. Two-dimensional principal component analysis (2DPCA) developed by J. Yang is an image projection technique that makes use of the spatial correlation information to achieve better performance than conventional one-dimensional PCA [104]. The basic idea of 2DPCA is to project image $\mathbf{A}$, an $m \times n$ random matrix, onto a projection vector $\mathbf{x}$ by the following linear transformation:

$$\mathbf{y} = \mathbf{A}\mathbf{x} \tag{5.23}$$

The discriminatory power of $\mathbf{x}$ is evaluated by the total scatter of the projected samples where the following criterion is adopted:

$$J(\mathbf{x}) = tr(\mathbf{S}_\mathbf{x}) \tag{5.24}$$

$\mathbf{S_x}$ is the covariance matrix of the projected feature vectors of the training samples and $tr(\mathbf{S_x})$ is the trace of $\mathbf{S_x}$. The covariance matrix $\mathbf{S_x}$ is given by the following equation:

$$\begin{aligned} \mathbf{S_x} &= E[(\mathbf{y} - E\mathbf{y})(\mathbf{y} - E\mathbf{y})^T] \\ &= E[((\mathbf{A} - E\mathbf{A})\mathbf{x})((\mathbf{A} - E\mathbf{A})\mathbf{x})^T] \end{aligned} \tag{5.25}$$

So,

$$tr(\mathbf{S_x}) = \mathbf{x}^T E[(\mathbf{A} - E(\mathbf{A}))^T(\mathbf{A} - E(\mathbf{A}))]\mathbf{x} = \mathbf{x}^T \mathbf{G}_t \mathbf{x} \tag{5.26}$$

where $\mathbf{G}_t$ is the image covariance (scatter) matrix. Suppose there are totally M image samples for training, then

$$\mathbf{G}_t = \frac{1}{M} \sum_{j=1}^{M} (\mathbf{A}_j - \bar{\mathbf{A}})^T (\mathbf{A}_j - \bar{\mathbf{A}}) \tag{5.27}$$

The optimal projection axes, $\mathbf{x}_{opt,1}, \mathbf{x}_{opt,2}, ..., \mathbf{x}_{opt,d}$, are the eigenvectors of $\mathbf{G}_t$ corresponding to the largest d eigenvalues.

Feature Extraction and Identification: In our experiment, for purely thermal based detection, 1000 thermal maps, $\mathbf{A}_1, \mathbf{A}_2, ..., \mathbf{A}_{1000}$, of authentic chips are used to evaluate the optimal projection axes $\mathbf{x}_{opt,1}, \mathbf{x}_{opt,2}, ..., \mathbf{x}_{opt,d}$. Then the extracted thermal feature matrix $\mathbf{B}$ is defined by

$$\mathbf{B} = [\bar{\mathbf{A}}\mathbf{x}_{opt,1}, \bar{\mathbf{A}}\mathbf{x}_{opt,1}, ..., \bar{\mathbf{A}}\mathbf{x}_{opt,d}] \tag{5.28}$$

For a given set of testing ICs, a feature matrix $\mathbf{B}_i$ is obtained for each IC after the transfor-

mation by 2DPCA. For power based detection method, instead of using authentic thermal maps to build the feature matrix, the inverted power maps from authentic chips are used.

5.4.5 Trojan Detection Methods

We use the obtained feature matrix $\mathbf{B}$, by the 2DPCA analysis for our Trojan Detection in two different ways. First, if trusted data from known chips are available for training, a supervised thresholding method is used. Second, if no prior known data are available for training, an unsupervised clustering technique is applied. Both the techniques are described in following section.

a. Supervised Thresholding method

The distance between the testing feature matrix $\mathbf{B}_i$ and the authentic feature matrix $\mathbf{B}$ is calculated by,

$$d(\mathbf{B}, \mathbf{B}_i) = \|\mathbf{B}_i - \mathbf{B}\|_2 \tag{5.29}$$

where $\|\mathbf{B}_i - \mathbf{B}\|_2$ is the Euclidean distance between $\mathbf{B}_i$ and $\mathbf{B}$. If the distance is larger than a certain threshold, the testing IC is identified as Trojan inserted. The threshold is related to false positive rates and obtained by applying the method to a set of authentic chips. For example, if we have 1000 authentic chips as training sets whose distances to golden chip are $d_1, \ldots, d_{1000}$, and we want to make the false positive rate within 1%, then the estimate of threshold is the value dividing $\{d_1, \ldots, d_{1000}\}$ into two sets, one has more than 990 chips and the other has less than 10 chips. The *supervised thresholding* method is applied to both thermal based Trojan detection and power based Trojan detection [2].

[2]The *supervised thresholding* method is formulated by the co-author Kangqiao Hu.

b. Unsupervised Clustering Method

Clustering is the most important unsupervised learning problem; it finds a natural grouping in a collection of unlabeled data and organizes objects into groups whose members are similar in some way [24]. As described in Section 5.4.3, we construct the residual power maps of the chips under test from the thermal maps. These detailed power maps with residual powers for each block have spatial groupings which can be used to distinguish chips under two different clusters, with and without Trojan. Since it is unsupervised, it means there is no learning step, and the algorithm does not need any prior knowledge, other than inputs which are the detailed power maps. This approach is suitable when we do not have a set of training chips in hand.

Appropriate Feature Selection: For clustering, it is very critical to choose an appropriate feature to be used for the partitioning. It influences the shape of the clusters as some elements may be close to one another according to one feature metric and farther away according to another. We have explored possible metrics to be used as a means for clustering our chips into two clusters of authentic chips and Trojan injected chips. We have the detailed residual power maps with $m \times n$ blocks of the chips under test, which we use for Trojan detection. To get a high resolution, we divide the chip layout into numerous blocks, which results in a high dimensional data set. One approach to cope with the problem of excessive dimensionality is to reduce the dimensionality by combining features. Some of the metrics or features that we have explored for clustering are maximum block power, variance among the block power, and spatial gradients among the block powers. The reason to use these features to distinguish between authentic and Trojan chips is that if there is a Trojan in the chip, then the maximum block power and variance among the block powers increases. Likewise, spatial gradients in the power maps also increase, which can help in detecting the Trojans. Another useful approach to tackle high dimensional data

is to perform principal component analysis. We have explored the principal components derived by the 2DPCA analysis earlier and used the norm of the feature vectors found in the feature matrix, which yields the most accurate Trojan detection rate.

In Figure 5.28 we plot the spatial gradients of the power maps in vertical and horizontal direction for process variation of 20% and 40%. We can see that natural clustering patterns are prevalent in the power maps, and gradients is a suitable feature to distinguish between the authentic chips and Trojan chips when process variation is 20%. We can observe that as the process variation increases, the two clusters start to overlap, as a result some of the chips become unidentifiable which makes Trojan detection harder. This problem,

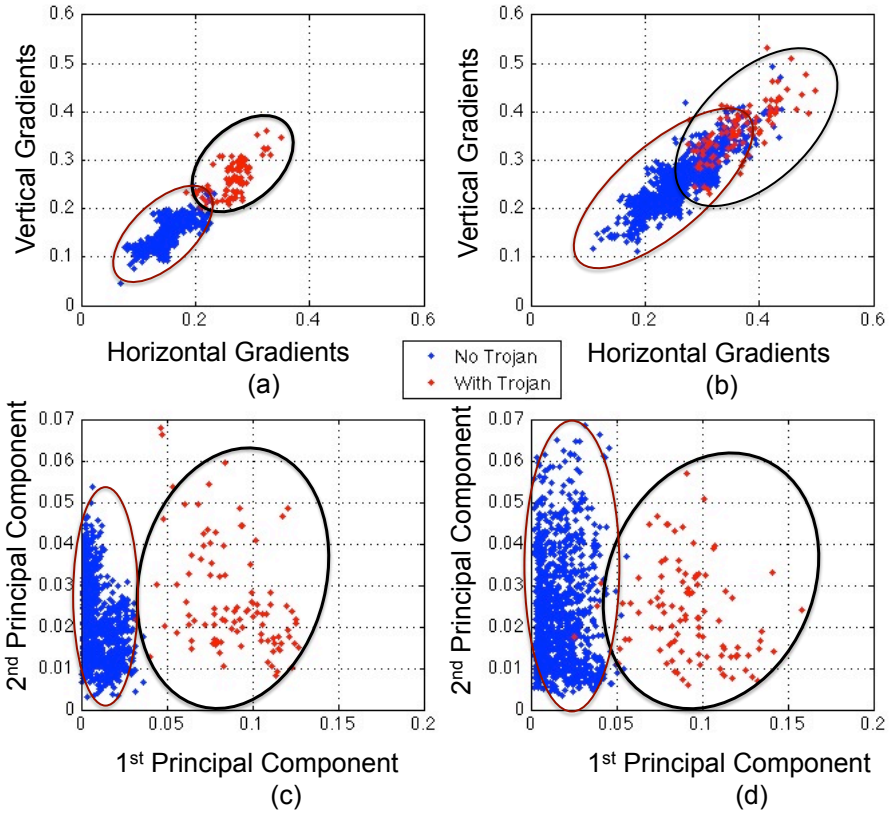

Figure 5.28: a) Gradients of power maps with PV 20%, b) Gradients of power maps with PV 40%, c) 1st and 2nd component of feature matrix with PV 20%, and d) 1st and 2nd component of feature matrix with PV 40%.

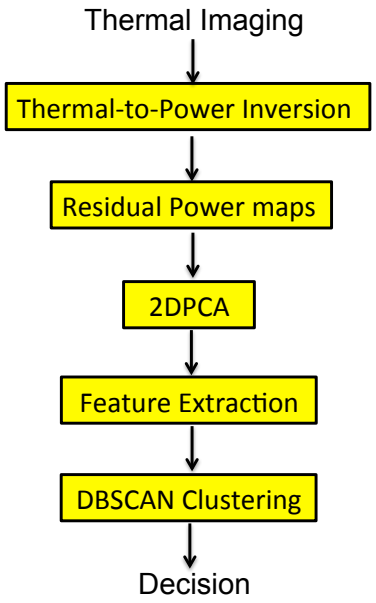

Figure 5.29: Unsupervised clustering flow.

which arises from the process variation, can be overcome by using 2DPCA. Using the norm of the feature vectors found in the feature matrix $\mathbf{B}$ in Section 5.4.4, we observe that even with 40% process variation, two clusters are properly distinguishable. We have used one trojan per chip and at one location for the purpose of this experiment. In case where there are more than one trojan, and different trojan locations, the clusters will be more distinguishable because there would be more variations in the feature matrices. As a result, detection rate will be higher. We use the norm of the feature matrix obtained by doing 2DPCA on the residual power maps for our clustering feature. Our clustering process is shown step-by-step in Figure 5.29. We propose to use spatial density based clustering DBSCAN method to distinguish chips infected with hardware Trojans.

Density-based spatial clustering: (DBSCAN) is a data clustering algorithm proposed by Ester *et al.* [26]. It is a density-based clustering algorithm because it finds a number of clusters starting from the estimated density distribution of corresponding nodes. The advantage of DBSCAN is that unlike other popular clustering algorithms, such as *k-means*

clustering, the accuracy of the clustering is not effected by the shape of the clusters. The spatial distribution of the clusters from authentic and Trojan chips can have any arbitrary shape making DBSCAN suitable for our case.

DBSCAN requires two parameters: *eps* and the minimum number of points required to form a cluster *minPts*. The *eps* is estimated from the data set, which is the geometric mean of the data. The ϵ-neighborhood is defined as the region that is covered with the given *eps*. The *minPts* is the minimum number of members that a cluster can have. It starts with an

---

Procedure: DBSCAN based Trojan Detection
Input: Infrared-based residual power estimates for each block for chips under test, $\mathbf{P}_r$ as $m \times n_x \times n_y$ matrix, where m is the number of chips and $(n_x \times n_y)$ is the number of blocks
Output: Return Trojan infected chips

Perform 2DPCA on $\mathbf{P}_r$ to get feature matrix $\mathbf{B}$, which is $m \times n_x \times n_y$ matrix

For $i = 1, 2...m$:

a. Find ℓ_2 norm of $\mathbf{B_i}$, $\mathbf{P}_m(i) = ||\mathbf{B_i}||_2$;

Estimate ϵ and minPts from $\mathbf{P}_m$

Mark all points in $\mathbf{P}_m$ as unvisited

For each unvisited point p:

a. Mark p as visited;
b. Find ϵ-points, all neighbourhood points;
c. **if** ϵ-points $<$ minPts: mark p as outlier;
d. **else if** p already in a cluster, add ϵ-points to the cluster;
e. **elseif** p not already in a cluster, start a new cluster and add ϵ-points to the cluster;

Return the outliers as Trojan chip

Figure 5.30: Unsupervised DBSCAN clustering for Trojan detection.

arbitrary starting point that has not been visited. This point's ϵ-neighborhood is retrieved, and if it contains sufficiently many points, a cluster is started. Otherwise, the point is labeled as noise. If a point is found to be a dense part of a cluster, its ϵ-neighborhood is also part of that cluster. Hence, all points that are found within the ϵ-neighborhood are added, as is their own ϵ-neighborhood when they are also dense. This process continues until the density-connected cluster is completely found. Then, a new unvisited point is retrieved and processed, leading to the discovery of a further cluster or noise. DBSCAN algorithm for clustering is a well-known method, and there are various implements of the algorithm. We follow the DBSCAN routine formulated by Daszykowski *et al.* [20]. Our Trojan detection procedure is described in Figure 5.30.

5.4.6 Trojan Localization

The inherent low-pass filter of heat conduction function makes it hard to accurately locate the Trojan, since most of the high frequency components are lost in the thermal maps [17]. With the detailed spatial power characterization technique these frequency components are well recovered in the power maps. We use the estimated residual power maps to locate the Trojans in the chip by finding the maximum power location in the Trojan detected chips:

$$\arg\max_{i,j} \mathbf{p}_{r(i,j)} \tag{5.30}$$

where $\mathbf{p}_r$ is the estimated residential power map and (i, j) is the grid position index. For more generalized cases, such as ICs with multi-Trojan, we detect local maxima points as the possible positions of the Trojans.

5.4.7 Impact of Thermal Imaging Noise on Trojan Detection

In this section, we discuss various noise sources that are present in the thermal imaging system and how it effects our proposed Trojan Detection methods. We describe methods to mitigate the effect of noise to improve Trojan detection results.

The main sources of noise in the infrared imaging system as discussed in Chapter 4 are: (1) thermal noise, (2) digitization noise, (3) dark noise, and (4) flicker noise [36]. Thermal noise is caused by agitation of charge carriers and is present in all electronic devices. The analog-to-digital converter in the infrared camera causes the digitization noise. Dark noise is due to the random generation of electron-hole pairs in the quantum detectors which is usually present in photosensitive devices. Flicker noise is related to the trapping and detrapping fluctuations of charge carries at the transistor interfaces. The first three noise sources fall under the category of frequency-independent white noise which is the major source of noise in thermal imaging [10]. The presence of noise in the thermal images can deteriorate the Trojan detection by hiding Trojans or by creating false alarms. We mainly focus on mitigating the effect of white noise in the thermal images in this work.

By using a larger integration time, the effect of white noise can be reduced dramatically. Integration time is defined as the time for which the thermal images for one single test are collected and averaged. As described in Section 5.4.3, the thermal maps in vector form is denoted by $\mathbf{t}$. Let $\mathbf{T_i(s)}$ denote the temperature of pixel i at time s as recorded by the thermal imaging system, and let I_p denote the integration period of the measurements. Then the *temperature magnitude*, $\mathbf{T_i}$, of pixel i is given by

$$\mathbf{T_i} = \frac{\mathbf{1}}{\mathbf{I_p}} \int_{\mathbf{0}}^{\mathbf{I_p}} \mathbf{T_i(s)ds} \tag{5.31}$$

If there is no noise in the measurements, then we expect $\mathbf{T_i}$ to exhibit no stochastic behavior. However, noise in the measurements lead to a stochastic process where $\mathbf{T_i}$ is a random variable. Since white noise has a Gaussian distribution, by increasing the integration time the standard deviation of the thermal signal which is responsible for the amplitude of noise can be reduced. From the central limit theorem, we know that the standard deviation of the average of a number of samples of random variable has $\frac{1}{\sqrt{s_n}}$ dependency on the number of samples, s_n. Thus, by increasing the integration time, we can reduce white noise proportionally to the square root of integration time [10].

5.4.8 Experimental Setup and Results

To test our proposed Trojan detection methods, we provide sophisticated simulation results which mimic a realistic experiment setup with process variation and then test four different benchmarks. We vary Trojan sizes and locations across the chips[3]. We provide the experimental results of two approaches. First, the thermal map based method which is more efficient in terms of computation time but less accurate; this does not require the thermal to power inversion procedure which increases detection time. Second, the detailed spatial power based method is used which is very accurate and can detect and locate very small Trojan; this requires the residual thermal to power inversion with l_1 regularization procedure.

[3]The addition of process variation, using four benchmarks for creating thermal maps, Trojan design and insertion to chip layout was done by the co-author Kangqiao Hu.

a. Experimental Setup

Process Variation: To characterize real-world ICs accurately, we add $20-40\%$ Process Variation (PV) to the gates' parameters. We use multi-level quad-tree approach to model the spatial within-die PV [3]. Higher levels of the quad-tree structure reflect the spatial correlations in larger scale while lower levels reflect the spatial correlations in smaller scale [3]. The effect of PV on dynamic power is neglected in our experiment since it is insignificant compared to the effect of PV on leakage power. Since I_{sub} is the dominant component of leakage power, we assume that the leakage current is equal to sub-threshold current. We add PV to gates' length, gates' width and gates' oxide thickness as [94]. In our experiment we set 5 different PV levels with variation of 20%, 25%, 30%, 35% and 40%, which introduces $\pm 0.5\%$ to $\pm 3\%$ variation to total power consumption.

IC Benchmarks: Four benchmarks from Opencores that are developed with Hardware Description Language (HDL) are used in our analysis: 1) 128-bit Advanced Encryption Standard (AES) cipher, 2) 32-bit MIPS Processor (MIPS), 3) Reed-Solomon (RS) Decoder and 4) Joint Photographic Experts Group (JPEG). Table 5.3 gives the basic information of benchmarks including number of gates, core size and total power consumption with standard voltage $1.1V$ at 1 GHz. We used Design Compiler synthesis tool from Synopsys to map the benchmarks to Nangate 45nm library and used Primetime-PX from Synopsys to estimate the average power consumption during a certain period with random vectors. We used Cadence SoC Encounter RTL compiler for floor planning, placing and routing, and Hotspot [39] for IC temperature simulation.

Trojan Design and Insertion: We have designed Trojans modules with power consumption varying from 0.05% to 0.2% total IC power consumption. Our Trojans do not have any specific functional modules but certain power consumption triggered by the test patterns that are used to evaluate the minimum size of Trojan that can be detected. Despite

Test bench	No. of Gates	Core Size (μm^2)	Nominal Power (W)
AES	10610	163×163	0.0732
MIPS	8661	195×195	0.0494
RS Decoder	23224	394×394	0.12
JPEG Encoder	269970	1094×1094	1.4675

Table 5.3: Test benches

the Trojan type, sequential or combinational, the power consumption ratio of the Trojan circuit and the IC is the only factor that impact our detection results. The Trojan circuits are implemented using the same standard cells as the ICs with a constant core utilization of approximate 70%. We divide the IC area into 10×10 blocks and insert one Trojan per chip into the blank space within these blocks. The impact of core utilization will be studied in the future work. For each benchmark at different PV levels, 10000 chips with one Trojan per chip of different sizes inserted in different locations are generated. With different PV levels, different Trojan sizes, different Trojan locations 100,000 chips of each benchmark are generated for testing.

b. Results

We conduct and report the results of five experiments.

1. In the first experiment, we perform our *supervised thresholding* Trojan detection technique on high resolution thermal maps, and report results for four different benchmarks. We also analyze the effect of false positive rate on Trojan detection rate.

2. In the second experiment, we present results of two different Trojan detection methods using residual power maps. We assess our detection results with four different

benchmarks and five different process variations.

3. We create Trojan infected chips with ten different Trojan locations. We compare our Trojan localization method under various benchmarks and process variation.

4. The fourth experiment evaluates the effect of thermal noise in an infrared imaging system on the detection results. We use different integration times and compare the accuracy of the Trojan detection.

5. In the fifth experiment, we increase the voltage of the chip from 1.1V to 1.2V, and assess the effect of increasing voltage on the Trojan detection results.

Experiment 1. In the first experiment we perform Trojan detection on high-resolution thermal maps. Based on the method proposed in Section 5.4.4, we first calculate the optimal projection vectors for each benchmark. All the thermal maps are simulated by HotSpot in $2^n \times 2^n$ grids. n depends on the die size and the resolution of infrared camera, $5 \times 5\mu m^2$. Thus, the thermal resolution of MIPS and AES is 32×32 grids, RS Decoder is 64×64 grids and JPEG is 128×128 grids. The thermal maps with resolution $2^n \times 2^n$ have 2^n eigenvectors in total. The number of eigenvectors that are used for feature extraction is determined by the magnitude of corresponding eigenvalues. Here we use benchmark AES as an example. We select eigenvectors corresponding to the first 10 largest eigenvalues as the optimal projection axes. Then the average thermal map of 1000 authentic chips are

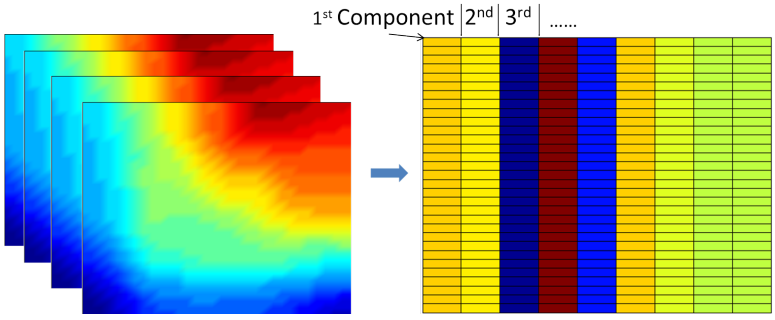

Figure 5.31: Golden feature matrix extraction.

used to extract the golden feature matrix $\mathbf{B}$ as shown in Figure 5.31. For each chip under test, the distance of its feature matrix and the golden feature matrix is computed.

As we mention in Section 5.4.4, the testing IC instance is identified as an authentic chip or a Trojan infected chip by a certain threshold that is associated with detecting false positives. Based on the distance histogram, we apply a kernel function to estimate the empirical probability distribution function (pdf) $f(d)$ for the authentic instances, where d denotes the distance from the golden feature matrix. Therefore, for a certain threshold d_{th}, the false positive is $\alpha = 1 - F(d_{th})$. By this, we fix the false positive to a certain value and observe how the false negative changes. Figure 5.32 shows that as the false positive increases, the detection rate increases while the false negative decreases. The controllability of the threshold helps us to easily adjust the algorithm to trade off false alarm and detection rate according to different detection requirements.

Detection Results under Different PV Level: The impact of PV is the most important factor that affects the performance of Trojan detection methods. Table 5.4 shows that with the fixed false positive rate, as the magnitude of PV increases, the detection rate decreases.

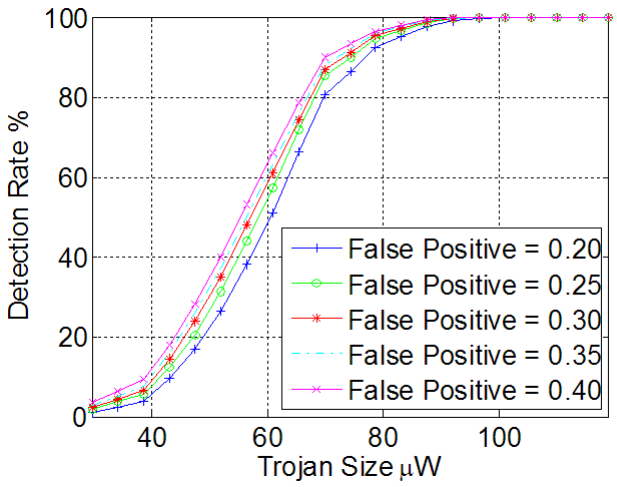

Figure 5.32: With fixed PV (0.2) and nominal voltage value (1.1V), the detection rate of AES under different false positive.

		Detection Rate %				
Benchmark	Trojan Size (uW)\PV(%)	20	25	30	35	40
AES	39.50	80.0	2.0	2.2	2.7	3.0
	59.30	26.5	19.1	15.4	16.3	11.4
	79.00	86.5	67.2	54.4	43.5	31.2
	98.80	99.7	96.7	89.9	77.6	64.9
	119.00	100.0	100.0	100.0	96.5	90.8
MIPS	58.60	32.0	19.7	11.7	8.0	5.6
	87.80	60.6	39.0	23.3	16.5	10.8
	117.00	81.8	62.6	40.1	26.0	17.7
	146.00	91.1	77.8	58.2	39.4	27.6
	176.00	92.4	88.5	73.7	55.9	39.3
RS Decoder	96.80	5.4	5.0	2.7	2.3	2.3
	145.00	9.8	7.5	4.3	4.8	3.7
	194.00	21.2	13.4	9.1	7.3	4.8
	242.00	36.2	22.2	12.5	9.3	6.1
	290.00	57.0	38.1	20.9	12.8	10.1
JPEG	43.90	3.2	3.1	2.0	2.0	1.5
	176.00	8.0	7.8	7.5	8.0	3.3
	307.00	15.0	13.0	11.1	10.5	4.5
	439.00	25.8	21.2	20.6	15.5	10.3
	571.00	50.1	45.6	30.5	18.5	11.1

Table 5.4: Trojan detection rate using thermal maps

The detection rate decreases in the following order: AES, MIPS, RS Decoder, JPEG. The main difference among these three benchmarks are the total power and the core size. If we define *power density*, as $\rho = \frac{P}{S_{core}}$, where P is total power and S_{core} is the size of the core, we notice that ρ decreases in the same order as performance, which means $\rho_{AES} > \rho_{MIPS} > \rho_{RS} > \rho_{JPEG}$. The chip with higher power density will generate more heat during the same period. Thus, a larger temperature gradient is formed, which makes the region with Trojan more prominent.

Experiment 2. In this experiment, we apply our *supervised thresholding* and *unsupervised clustering* techniques proposed in Section 5.4.5 on detailed residual power maps. These high resolution power maps results in a very high sensitive Trojan detection. Overall, the power mapping approach has a much higher sensitivity than the thermal mapping approach. The main reason for the dramatic improvement from thermal mapping to power

mapping is the proper recovery of the high frequency components of the power map. We plot detection rates for four benchmarks with tens of Trojan sizes and under five different process variations in Figure 5.33. We define detection rate and false positive rate as follows,

$$DetectionRate = N_{TD}/N_T \tag{5.32}$$

$$FalsePositiveRate = N_{FD}/N_F \tag{5.33}$$

where N_{TD} is the total number of detected Trojan chips, N_T is the total number of Trojan chips, N_{FD} is the number of authentic chips that is detected as Trojan and N_F is number of authentic chips.

We present detection results from our two different approaches, the *supervised thresholding* in Figure 5.33a and *unsupervised clustering* method in Figure 5.33b. We observe that the *supervised thresholding* has a higher detection rate than the unsupervised method,

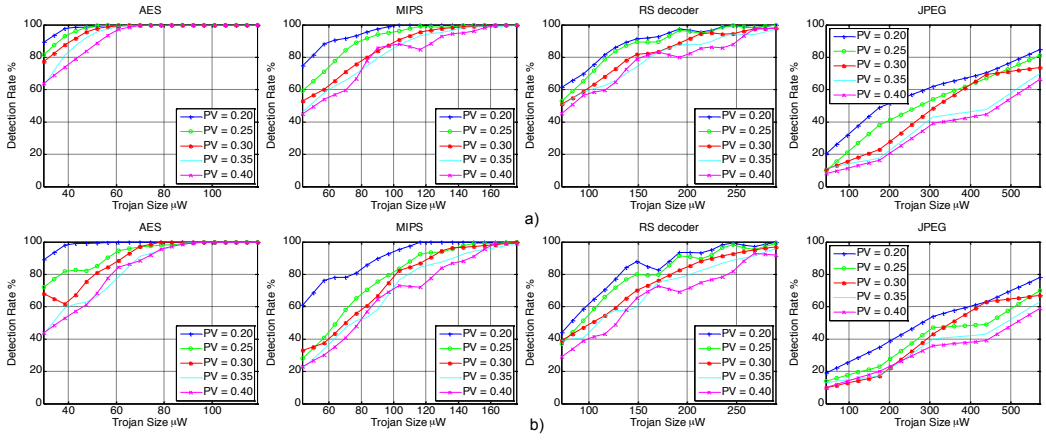

Figure 5.33: Detection rates for four benchmarks using power maps (1.1V), under different process variation level, a) *Supervised thresholding* technique and b) *Unsupervised clustering* technique.

Detection Method		Supervised Thresholding					Unsupervised Clustering									
		Detection Rate %					Detection Rate %					False Positive Rate %				
Benchmark	Trojan Size(uW)\PV(%)	20	25	30	35	40	20	25	30	35	40	20	25	30	35	40
AES	29.6	85	81	71	59	53	89	72	68	43	44	5	11	12	6	8
	39.5	98	93	86	80	66	99	83	61	61	54	5	10	4	6	7
	49.4	98	99	97	91	81	99	82	79	63	63	5	6	8	4	6
	59.3	100	100	99	99	97	100	94	87	76	84	5	4	5	4	5
	69.2	100	100	100	100	99	100	97	97	91	88	5	1	4	4	5
	79.0	100	100	100	100	100	100	98	100	95	96	5	1	4	2	5
	88.9	100	100	100	100	100	100	99	100	100	99	5	1	4	3	5
	98.8	100	100	100	100	100	100	100	100	100	100	5	1	4	3	5
	109.0	100	100	100	100	100	100	100	100	100	100	5	1	2	3	4
	119.0	100	100	100	100	100	100	100	100	100	100	5	1	2	3	4
MIPS	43.9	69	51	44	33	33	61	28	33	22	23	6	3	5	4	4
	58.6	85	65	51	49	49	78	42	38	35	31	6	3	4	5	4
	73.2	90	81	67	62	50	78	62	53	50	43	5	3	4	5	4
	87.8	96	90	76	72	77	89	74	64	55	63	5	2	4	4	4
	102.0	100	93	84	84	84	95	83	82	76	73	4	2	4	3	4
	117.0	100	99	93	90	80	100	93	87	85	72	5	3	4	3	3
	132.0	100	97	98	93	94	100	94	96	88	86	4	2	3	3	4
	146.0	100	100	97	98	94	100	99	97	94	89	4	2	3	3	4
	161.0	100	100	100	98	100	100	99	98	96	99	4	2	3	3	4
	176.0	100	100	100	99	100	100	100	100	99	99	4	2	3	3	4
RS Decoder	72.6	59	48	41	43	35	44	37	39	39	29	3	4	6	8	7
	96.8	70	66	57	53	47	60	53	48	41	40	3	4	4	6	7
	121.0	81	78	70	57	58	73	69	56	58	44	2	4	3	7	5
	145.0	91	91	80	71	74	89	80	69	57	64	3	3	3	4	5
	169.0	91	86	82	82	81	82	79	76	75	73	2	3	3	3	3
	194.0	98	97	88	86	76	94	92	83	78	69	1	2	3	4	4
	218.0	95	94	94	85	84	93	89	89	83	76	1	0	2	2	3
	242.0	100	100	94	94	82	100	99	92	88	79	1	0	2	2	3
	266.0	98	98	98	93	97	97	95	95	91	93	1	0	2	2	3
	290.0	100	100	97	97	98	100	99	97	96	92	1	0	1	2	2
JPEG	43.9	20	10	11	11	8	19	14	10	13	10	9	13	8	11	11
	176.0	49	38	23	18	16	35	23	17	17	20	4	8	7	11	11
	307.0	62	54	48	43	39	54	47	43	40	36	5	5	8	11	11
	439.0	70	67	69	48	45	63	49	63	43	39	3	3	6	7	7
	571.0	85	81	74	70	66	78	70	67	63	59	3	3	4	6	6

Table 5.5: Trojan detection results using power maps

which is expected. The advantage of the unsupervised is that we do not require any prior data set for the training purpose. One can observe the trend that detection rate increases as the Trojan sizes increase, and decreases with process variation as expected. From the table, we see that we are able to detect Trojan as small as 29.6 μW. The detection rate for each benchmark follows the same order as we have seen in the thermal map results in Experiment 1. The detection rate performance follows the same order as the power density of the chip, detection rate of AES $>$ MIPS $>$ RS_DEC $>$ JPEG.

For the *supervised thresholding*, we are able to fix the false positive rate to 1%. Since the second method is unsupervised, there is no way to fix the false positive rate in that case. We add the corresponding false positive rate for the clustering method in Table 5.5 for comparison purposes. The false positive rate increases as the Trojan size decreases. For

smaller Trojan sizes, we also see that false positive rate increases as the process variation goes up. Table 5.5 lists all the experimental results with residual power mapping.

Experiment 3. In this experiment, we present results of our Trojan localization method under various benchmarks and process variations. The sparsification process makes it very easy to localize the Trojan. Once a chip is identified as an infected chip, we simply use the estimated residual power map to locate the Trojan by finding the maximum power location. We compute the Euclidean distance of the estimated location and the real location. By normalizing the distance to the chip core dimension we get the normalized localization error. Figure 5.34 shows the normalized localization error for four different benchmarks, AES, MIPS, RS Decoder and JPEG with five different process variations and five different

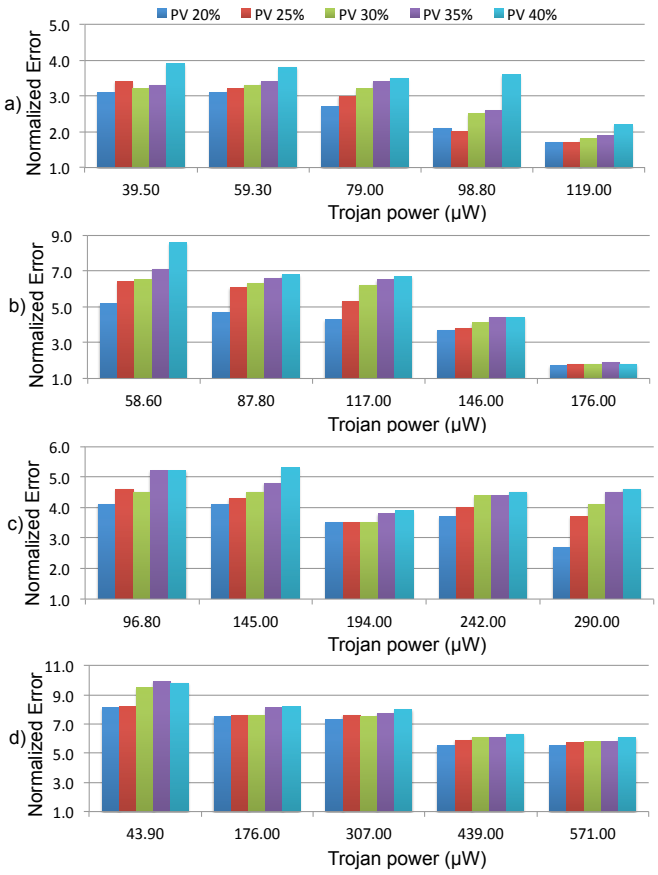

Figure 5.34: Normalized localization error with five different PVs and four benchmarks, a) AES, b) MIPS, c) RS Decoder and d) JPEG.

Trojan sizes. We can see the Trojan localization error increases with increasing process variation. Also, as the Trojan size decreases, it becomes more difficult to localize Trojan under the same process variation.

Experiment 4. In this experiment, we add white Gaussian noise to our thermal maps to mimic a real infrared imaging system which has standard deviation of 10mK. We select benchmark MIPS with process variation of 30% for this experiment. We add Gaussian noise to all our thermal maps, and then change our integration times. As discussed in Section 5.4.7, the white noise in infrared imaging is indirectly proportional to the square root of the integration time (IT). We then perform our residual power mapping as described in Section 5.4.3 on the integrated thermal maps. We apply the Trojan detection method with clustering as presented in Section 5.4.5 on the residual power maps. Figure 5.35 shows the detection results. 'With Noise(IT-1)' is the case where noise has been added to the thermal maps, and integrated over one frame, (IT-100) is integrated over 100 frames, (IT-1000) is integrated over 1000 frames. If the camera frame rate is 100 Hz, then 100 frames correspond to integration over 1s, 1000 frames is integration over 10s. 'No Noise(IT-1)' stands for where no thermal noise has been added to the thermal maps. We see that as the integration time increases, the detection rate with noisy thermal maps approaches the

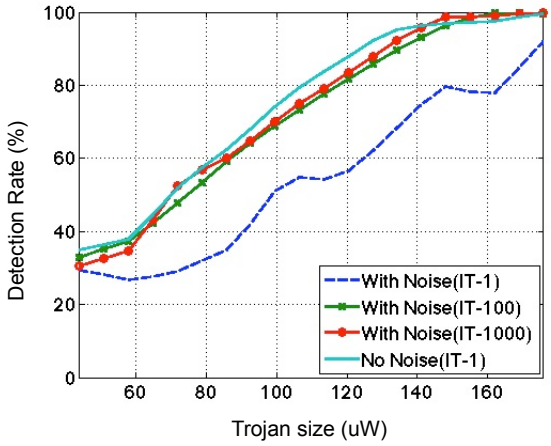

Figure 5.35: Detection rate for MIPS with PV 30%.

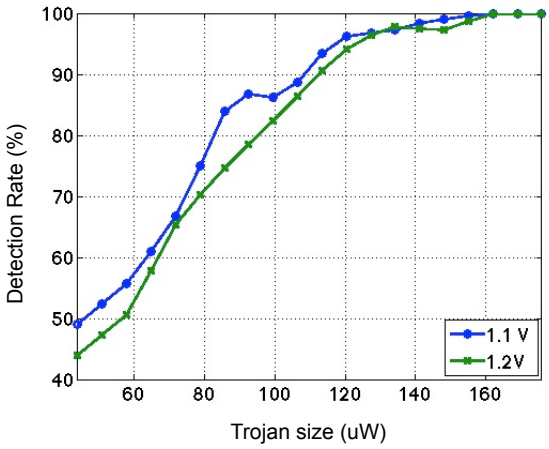

Figure 5.36: Detection rate for MIPS with 1.1 and 1.2 V (PV 30%).

detection rate without noise. We conclude from this experiment that, the white noise that is present in the thermal maps can be compensated with increasing the integration times.

Experiment 5. We selected benchmark MIPS with PV 30% for this experiment and performed thresholding detection method. Figure 5.36 shows the results for two voltage cases, 1.1V and 1.2V. The plots show that there is not much difference since by increasing the voltage, the power of the Trojan and the power of variation caused by PV are both increased.

5.4.9 Conclusions

We described a novel multimodal post-silicon framework using spatial thermal and power maps in order to detect and locate Trojans in modern ICs as the third application for post-silicon power mapping. We developed two different Trojan detection methods, *supervised thresholding* and *unsupervised clustering* technique utilizing post-silicon power maps. These power maps can also reveal the Trojan location very accurately. We introduced ℓ_1 regularization in the thermal to power inversion to exploit the sparsity of the residual power maps. Using proposed multimodal methods, we are able to detect Trojans

which consume power as small as 29.6 μW. To create realistic chips, we added 20-40% process variations, our results show Trojan detection is inversely proportional to the process variations as it can hide Trojans. We added Gaussian noise to our thermal maps, and showed that the effect of infrared imaging noise on the detection rates can be mitigated by increasing the image integration time. We also compared detection results for different chip voltage settings.

5.5 Summary of the Applications of Post-silicon Power Mapping

In Section 5.2, we used our proposed DC power mapping methodology to estimate the spatial power distribution of an embedded soft processor and a real quad-core processor. We performed thermal to power inversion of the embedded processor with different settings while running practical applications. For the quad-core processor, we have proposed techniques to model leakage power through the use of thermal conditioning. These leakage power models were used to yield fine-resolution leakage power maps and within-die variability trends. Various workload scenarios from the SPEC CPU2006 benchmarks are used to estimate and analyze runtime spatial power maps. The basic applicability of our AC-based techniques is also presented using a real dual-core processor.

In Section 5.3 we proposed new directions for post-silicon power mapping methodologies where thermal maps reconstructed from sparse thermal sensor measurements are used in place of infrared thermal maps. The proposed method provide good power mapping accuracy for real processors while reducing the cost drastically. We explore various frequency domain techniques for the full thermal characterization, such as, `k-LSE`, `CS-L1_MIN`, and the `CS-STOMP`. Using a sophisticated experimental setup, we evaluated our techniques using various computing devices; dual-core processor, FPGA test chip and a quad-core processor with various applications.

Section 5.4 presents a novel multimodal framework utilizing post-silicon spatial thermal and power maps in order to detect and locate Trojans in modern ICs. We proposed two different Trojan detection methods, *supervised thresholding* and *unsupervised clustering* technique utilizing post-silicon power maps. These power maps can also reveal the Trojan location very accurately. To exploit the sparsity of the residual power maps we added ℓ_1

regularization in the thermal to power inversion procedure. Using proposed multimodal methods, we are able to detect Trojans which consume power as small as 29.6 μW. We added 20-40$\%$ process variations to gates' length, gates' width and gates' oxide thickness which results into 0.5-3$\%$ variation in total chip power consumption. These variations, which are present in real modern ICs, hide Trojans and makes Trojan detection more difficult. To mimic real infrared imaging, we added Gaussian noise to our thermal maps, and showed that the effect of infrared imaging noise on the detection rates can be mitigated by increasing the image integration time. Detection results for different chip voltage settings are also compared. One limitation to our approach is that we assume the Trojan circuitry is consuming a certain amount of total power. In case the Trojan is not activated during infrared imaging by specific runtime applications it will not consume any significant amount of dynamic power, and Trojan can remain undetected. But the Trojan will contribute to more leakage power which is dependent on temperature. By increasing chip temperature, the effect of Trojan leakage power can be detected, which can aid in the Trojan detection.

Chapter 6

Summary of Dissertation and Possible Future Extensions

In this dissertation, we proposed techniques for post-silicon power characterization of various computing devices, such as, FPGAs and multi-core processors. We addressed main challenges in a post-silicon power characterization process, and proposed methods to overcome such challenges. Utilizing the proposed framework, we presented three different applications of post-silicon power mapping. We verified all our proposed methods through extensive set of experimental results using various test chips. In this chapter, we summarize our contributions by highlighting the main results, and discussing possible future research directions.

6.1 Summary of Results

In Chapter 3, we presented a novel framework for spatial post-silicon power characterization using the infrared emissions from the back side of the silicon die. Various challenges of thermal to power inversion procedure are investigated. We demonstrated mathematically and empirically how challenges such as inherent lowpass filtering, discretization, and measurement errors could all compromise the accuracy of power estimation. We proposed techniques from regularization theory to improve the accuracy of post-silicon thermal to power inversion. Furthermore, we have provided experimental techniques to compensate for the varying emissivity of different chip materials and to measure the thermal resistance model 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 provided deep insights on the challenges that are involved in post-silicon power characterization and confirmed that our methodology works very well in practice providing power estimation improvement of 30% over previous approaches. Our results also quantified the effect of the two main challenges of post-silicon power mapping, noise and spatial filtering.

In Chapter 4, we investigated the use of AC thermography techniques for quantitive power mapping that is relevant to design-related applications. We demonstrated analytically and experimentally that using AC excitation reduces measurement noise and spatial heat diffusion. We devised techniques for realistic estimation of the parameters of the thermal to power model matrix, and we devised numerical techniques to invert the thermal emissions into power estimates. We designed a test chip to scientifically evaluate the accuracy of AC inversion techniques compared to DC techniques. Our test chip enabled us to create any desired spatial power maps with different excitation frequency. Using a

number of constructed intricate power patterns, we demonstrated that our technique can dramatically improve post-silicon power mapping. The average error reduced from 40% at DC to about 8.5% using proposed AC methods. To quantify the noise in our system, we analyzed and quantified the signal-to-noise ratio of infrared thermal maps. We showed the impact of AC excitation frequency on noise and spatial filtering of the thermal maps. We also quantized the effect of integration time on noise. We analyzed the power mapping results for different AC excitation frequencies and integration times, and linked these results to the SNR analysis.

In Chapter 5, we presented detailed methods and experimental results for three different applications of our power mapping framework. Firstly, we implemented our proposed post-silicon power mapping techniques on FPGAs and multi-core processors. For power characterization of FPGAs, we embedded a soft processor in FPGA test chip, and used our power mapping methodology to estimate the spatial power distribution of the soft processor during operation. For multi-core processors, we devised accurate finite-element models that relate power consumption to temperatures. By using thermal conditioning methods, we proposed to model leakage power through infrared thermal measurements. These leakage power models were used to yield fine-resolution leakage power maps and within-die variability trends for multi-core processors. We analyzed the power consumption of different blocks of a quad-core processor under different workload scenarios from the SPEC CPU2006 benchmarks. Our results revealed a number of insights into the make-up and scalability of power consumption in modern processors. We showed the impact of increasing number of applications on the power consumption of different blocks. By using a dual-core processor, we demonstrated the applicability of our AC techniques in a practical scenario.

Secondly, we proposed new frequency domain based methods for full thermal characterization of real processors. We introduced reconstruction methods such as `k-LSE`,

CS−L1_MIN, and the CS−STOMP, which enable us to reconstruct fine-grain thermal maps from sparse sensor measurements accurately. By utilizing the full reconstructed thermal maps, we proposed an alternate method to the infrared imaging for our post-silicon power mapping procedure. This approach can reduce the cost of post-silicon power mapping drastically. We demonstrated experimental results on a FPGA test chip with several artificially created power patterns and relate the power mapping accuracy to the spatial frequency of the power maps. We utilized a multi-core processor with real benchmark applications to show that our proposed thermal sensor based method shows relatively good accuracy compare to the infrared based power mapping.

Thirdly, we investigated the use of multimodal post-silicon spatial thermal and power maps in order to detect and locate Trojans in modern ICs. We developed two different Trojan detection methods, *supervised thresholding* and *unsupervised clustering* technique. Through an extensive set of benchmarks and experiments, we demonstrated that using high resolution thermal maps increases the Trojan detection sensitivity. To improve the sensitivity further, we inverted the residual thermal maps to detailed spatial power maps which are then utilized for Trojan detection. To exploit the sparsity of the residual thermal maps, we added ℓ_1 regularization in our power mapping procedure which improves detection rate. These power maps can also reveal the Trojan location very accurately. Using proposed multimodal methods, we were able to detect Trojans which consume power as small as 29.6 μW. We have demonstrated that detection rate is directly proportional to the power density of the chip, and the Trojan size. Process variation changes power profile of the chip and complicates the Trojan detection. To account for die process variation, we added 20-40$\%$ Gaussian random variations to gates' length, gates' width and gates' oxide thickness. Our results show Trojan detection is inversely proportional to process variations. In real infrared imaging setup, thermal maps have various measurement noise. Since we utilized thermal tool Hotspot for creating all our thermal maps, we added

Gaussian noise to our thermal maps to mimic real infrared setup. We have shown that by increasing integration time the white noise can be minimized to improve detection rate. We also compared different chip voltage settings, which did not show much effect on the detection results.

6.2 Possible Research Extensions

This dissertation can lead to the following possible research directions:

One possible extension to the current power mapping framework is to handle transient analysis. The relationship between temperature and power in transient analysis can be described using state space models. In particular, the temperature vector $\mathbf{T}[k+1]$ at time $k+1$ is linked to the temperature vector $\mathbf{T}[k]$ at time k and the power consumption $\mathbf{p}[k]$ at time k by $\mathbf{T}[k+1] = \mathbf{A}\mathbf{T}[k] + \mathbf{B}\mathbf{p}[k]$. Transient analysis can follow a similar approach to the one described in this work, where $\mathbf{T}[k+1]$ and $\mathbf{T}[k]$ are measured using the camera, and $\mathbf{A}$ and $\mathbf{B}$ are learned in a similar way to the approach used to learn $\mathbf{R}$.

Another possible future extension of the post-silicon power mapping is applying the power mapping procedure to 3D ICs. 3D circuits are lucrative future directions in the IC industry due to smaller footprint, lower potential cost, heterogeneous integration and shorter interconnects. But heat build up within the stack could be a major constraints in the design of 3D structures. Post-silicon power mapping can aid in the thermal and power design issues. The modeling matrix can be built using Finite Element Modeling procedure in a similar way as described in Chapter 5. Once the modeling matrix is build, power mapping procedure with $\mathbf{Rp} = \mathbf{T}$ as described in Chapter 3 altered according to the 3D design, can be used. In Chapter 5 we proposed an alternative way for post-silicon power

mapping using reconstructed thermal maps from sparse thermal sensor measurements. This alternative way can be also be used to extend the 2D power mapping framework to the 3D power mapping in future. In the case of 3D IC power mapping, thermal sensor readings can be used to get access to thermal status within the 3D stacks. The reconstructed thermal maps from the thermal sensor measurements can be used for 3D post-silicon thermal to power inversion.

For future work of AC-based framework, higher excitation frequencies can be investigated to find the limits of AC techniques. In this work, we have shown basic applicability of the AC procedure in a dual-core processor, which can be extended to a more detailed full power mapping procedure for multi-core processors. Furthermore, the AC techniques can be used in many scenarios to reduce the effect of noise and spatial lowpass filtering, such as, thermal sensor noise and process variation noise in Trojan detection. As a future extension of the hardware Trojan application, the AC thermography framework described in Chapter 4 can be incorporated to the multimodal Trojan detection framework to increase the detection accuracy.

There are several other applications for the post-silicon power mapping framework. Most importantly, the post-silicon power estimates can be used to close the loop between pre-silicon modeling and post-silicon characterization, where one can leverage the post-silicon results to improve the accuracy of different design-time thermal and power modeling tools.

Bibliography

[1] E. Acar, A. Devgan, R. Rao, Y. Liu, H. Su, S. Nassif, and J. Burns. Leakage and Leakage Sensitivity Computation for Combinational Circuits. In *International Symposium on Low-Power Electronics*, pages 96–99, 2003.

[2] A. Agarwal, S. Mukhopadhyay, A. Raychowdhury, K. Roy, and C. H. Kim. Leakage Power Analysis and Reduction for Nanoscale Circuits. *IEEE Micro*, 26(2):68–80, 2006.

[3] A. Agarwal, V. Zolotov, and D.T. Blaauw. Statistical Timing Analysis using Bounds and Selective Enumeration. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 22(9):1243–1260, 2003.

[4] D. Agrawal, S. Baktir, D. Karakoyunlu, P. Rohatgi, and B. Sunar. Trojan Detection using IC Fingerprinting. In *IEEE Symposium on Security and Privacy*, pages 296–310, 2007.

[5] E. Alpaydin. *Introduction to Machine Learning*. The MIT Press, first edition, 2004.

[6] M. Banga and M.S. Hsiao. A Region Based Approach for the Identification of Hardware Trojans. In *IEEE International Workshop on Hardware-Oriented Security and Trust*, pages 40–47, 2008.

[7] L. A. Barroso and U. Holzle. *The Datacenter as a Computer*. Morgan and Claypool Publishers, 2009.

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

[9] D. Boning and S. Nassif. Models of Process Variations in Device and Interconnect. In A. Chandrakasan, W. J. Bowhill, and F. Cox, editors, *Design of High-Performance Microporcessor Circuits*, pages 98–115. IEEE Press, 1 edition, 2001.

[10] O. Breitenstein, W. Warta, and M. Langenkamp. *Lock-In Thermography: Basics and Use for Functional Diagnostics of Electronic Components*. Springer Verlag, second edition, 2010.

[11] D. Brooks, R. Dick, R. Joseph, and L. Shang. Power, Thermal, and Reliability Modeling in Nanometer-Scale Microprocessors. *IEEE Micro*, 27(3):49 – 62, 2007.

[12] D. Brooks, M. Martonosi, J.-D. Wellman, and P. Bose. Power-Performance Modeling and Tradeoff Analysis for a High End Microprocessor. In *In Power Aware Computing Systems Workshop at ASPLOS-IX*, pages 126–136, 2000.

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

[14] E Candès. Compressive Sampling. *Proceedings of the International Congress of Mathematicians*, pages 1–20, 2006.

[15] J. Chang, A. A. Abidi, and C. R. Viswanathan. Flicker Noise in CMOS Transistors from Subthreshold to Strong Inversion at Various Temperatures. *IEEE Transactions on Electronic Devices*, 41(11):1965–1971, 1994.

[16] R. Y. Chen, M. J. Irwin, and R. S. Bajwa. Architecture-Level Power Estimation and Design Experiments. *ACM Transactions on Design Automation of Electronic Systems*, 6(1):50–66, 2001.

[17] R. Cochran, A. Nowroz, and S. Reda. Post-Silicon Power Characterization Using Thermal Infrared Emissions. In *International Symposium on Low Power Electronics and Design*, pages 331–336, 2010.

[18] R. Cochran and S. Reda. Spectral Techniques for High-Resolution Thermal Characterization with Limited Sensor Data. In *Proceedings of the Design Automation Conference*, pages 478–483, 2009.

[19] B. P. Colwell. *The Pentium Chronicles: The People, Passion, and Politics Behind Intel's Landmark Chips*. Wiley-IEEE Computer Society, 2005.

[20] M. Daszykowski, B. Walczak, and D. L. Massart. *Looking for Natural Patterns in Data. Part 1: Density Based Approach*. Chemometrics and Intelligent Laboratory Systems 56, 2001.

[21] Kapil Dev, Abdullah Nazma Nowroz, and Sherief Reda. Power Mapping and Modeling of Multi-core Processors. In *International Symposium on Low Power Electronics and Design*, 2013.

[22] D Donoho, Y Tsaig, I Drori, and J Starck. Sparse Solution of Underdetermined Linear Equations by Stagewise Orthogonal Matching Pursuit. *Stanford University Technical Report*, Jan 2006.

[23] Defense Science Board (DSB) study on high performance microchip supply, http://www.acq.osd.mil/dsb/reports/2005-02-hpms_report_final.pdf.

[24] Richard O. Duda, Peter E. Hart, and David G. Stork. *Pattern Classification*. Wiley, John & Sons, Incorporated, second edition, 2000.

[25] J. Emer, P. Ahuja, E. Borch, A. Klauser, Chi-Keung Luk, S. Manne, S.S. Mukherjee, H. Patil, S. Wallace, N. Binkert, R. Espasa, and T Juan. Asim: A Performance Model Framework. *Computer*, 35:68–76, 2002.

[26] Martin Ester, Hans peter Kriegel, Jrg S, and Xiaowei Xu. A Density-based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. pages 226–231. AAAI Press, 1996.

[27] 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.

[28] S. H. Fuller and L. I. Millett. *The Future of Computing Performance: Game Over or Next Level?* The National Academies Press, 2011.

[29] P. Gupta, A. B. Kahng, and S. Muddu. Quantifying Error in Dynamic Power Estimation of CMOS Circuits. *Analog Integrated Circuits and Signal Processing*, 42(3):253–264, 2005.

[30] S. Gupta and F. N. Najm. Power Modeling for High-Level Power Estimation. *Transactions on Very Large Scale Integration Systems*, 8(1):18–29, 2000.

[31] H. Hamann. Personal Communication, 2010.

[32] H. Hamann, A. Weger, J. Lacey, E. Cohen, and C. Atherton. Power Distribution Measurements of the Dual Core PowerPC 970MP Microprocessor. In *IEEE International Solid-State Circuits Conference*, pages 2172–2179, 2006.

[33] 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.

[34] P. C. Hansen. *Discrete Inverse Problems: Insight and Algorithms*. Society for Industrial and Applied Math, 2010.

[35] Wim Heirman, Souradip Sarkar, Trevor E. Carlson, Ibrahim Hur, and Lieven Eeckhout. Power-aware Multi-core Simulation for Early Design Stage Hardware/software Co-optimization. In *Proceedings of the 21st International Conference on Parallel architectures and compilation techniques*, PACT '12, pages 3–12, New York, NY, USA, 2012. ACM.

[36] C. Hsieh, C. Wu, F. Jih, and T. Sun. Focal-Plane-Arrays and CMOS Readout Techniques of Infrared Imaging Systems. *IEEE Transactions on Circuits and Systems for Video Technology*, 7(4):594–605, 1997.

[37] Kangqiao Hu, Abdullah Nazma Nowroz, Sherief Reda, and Farinaz Koushanfar. High-Sensitivity Hardware Trojan Detection Using Multimodal Characterization. In *Design, Automation and Test in Europe*, pages 1271–1276, Grenoble, France, 2013.

[38] 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.

[39] W. Huang, S. Ghosh, S. Velusamy, K. Sankaranarayanan, K. Skadron, and M. Stan. HotSpot: A Compact Thermal Modeling Methodology for Early-Stage VLSI Design. *IEEE Transactions onVery Large Scale Integration Systems*, 14(5):501 – 513, 2006.

[40] W. Huang, K. Skadron, S. Gurumurthi, R. J. Ribando, and Mircea R. Stan. Differentiating the Roles of IR Measurement and Simulation for Power and Temperature-Aware Design. In *Proceedings of the International Symposium on Performance Analysis of Systems and Software*, pages 1–10, 2009.

[41] S. Huth, O. Breitenstein, A. Huber, and D. Dantz. Lock-in IR-Thermography - A Novel Tool for Material and Device Characterization. *Solid State Phenomena*, 82-84:741–746, 2002.

[42] 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.

[43] Y. Jin and Y. Makris. Hardware TROJAN detection using path delay fingerprint. In *International Workshop on Hardware-Oriented Security and Trust*, pages 51–57, 2008.

[44] Z. Kamal, Q. Hassan, and Z. Mouhcine. Full On-chip CMOS Low Dropout Voltage Regulator using MOS Capacitor Compensation. In *International Conference on Multimedia Computing and Systems*, pages 1109–1114, 2012.

[45] T. Kemper, Y. Zhang, Z. Bian, and A. Shakouri. Ultrafast Temperature Profile Calculation in IC Chips. In *International Workshop on Thermal Investigations of ICs and Systems*, pages 133–137, 2006.

[46] N. S. Kim, K. Flautner, D. Blaauw, and T. Mudge. Drowsy Instruction Caches: Leakage Power Reduction Using Dynamic Voltage Scaling and Cache Sub-Bank Prediction. In *International Symposium on Microarchitecture*, pages 219–230, 2002.

[47] Wonyoung Kim, Meeta S. Gupta, Gu yeon Wei, and David M. Brooks. Enabling OnChip Switching Regulators for Multi-Core Processors using Current Staggering. In *In Proceedings of the Work. on Architectural Support for Gigascale Integration*, 2007.

[48] F. Koushanfar and A. Mirhoseini. A Unified Framework for Multimodal Submodular Integrated Circuits TROJAN Detection. *IEEE Transactions on Information Forensics and Security*, 6(1):162–174, 2011.

[49] B. Lee and D. Brooks. Accurate and Efficient Regression Modeling for Microarchitectural Performance and Power Prediction. In *Architectural Support for Programming Languages and Operating Systems*, pages 185–194, 2006.

[50] F. Li, Y. Lin, L. He, D. Chen, and J. Cong. Power Modeling and Characteristics of Field Programmable Gate Arrays. *IEEE Transactions on Computer Aided Design of Integrated Circuits*, 24(11):1712–1724, 2005.

[51] J. Li and J. Lach. At-speed Delay Characterization for IC Authentication and TROJAN Horse Detection. In *International Workshop on Hardware-Oriented Security and Trust*, pages 8–14, 2008.

[52] S.-C. Lin and K. Banerjee. Cool Chips: Opportunities and Implications for Power and Thermal Management. *IEEE Transactions on Electron Devices*, 55(1):245 – 255, 2008.

[53] J. Long, S. Memik, G. Memik, and R. Mukherjee. Thermal Monitoring Mechanisms for Chip Multiprocessors. In *ACM Transactions on Architecture and Code Optimization*, volume 5(2), pages 9:1–9:23, 2008.

[54] M. Mamidipaka and N. Dutt. eCACTI: An Enhanced Power Estimation Model for On-Chip Caches. Technical report, In Technical Report TR-04-28, CECS, UCI, 2004.

[55] E. Marin. The Role of Thermal Properties in Periodic Time-Varying Phenomena. *Eur. J. Physics*, 28(3):429–445, 2007.

[56] R. McGowen, C. Poirier, C. Bostak, and J. Ignowski. Power and Temperature Control on a 90-nm Itanium Family Processor. *IEEE Transactions on Solid-State Circuits*, 41(1):229–237, 2006.

[57] H Mehta, R.M. Owens, and M.J. Irwin. Energy Characterization based on Clustering. In *Design Automation Conference*, pages 702 – 707, 1996.

[58] S. 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*, 27(3):516–527, 2008.

[59] F. J. Mesa-Martinez, E. Ardestani, and J. Renau. Characterizing Processor Thermal Behavior. In *Architectural Support for Programming Languages and Operating Systems*, pages 193–204, 2010.

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

[61] Francisco Javier Mesa-Martinez, Joseph Nayfach-Battilana, and Jose Renau. Power Model Validation Through Thermal Measurements. In *Proceedings of the 34th annual International Symposium on Computer architecture*, pages 302–311. ACM, 2007.

[62] J. Monteiro, R. Patel, and V. Tiwari. Power Analysis and Optimization from Circuit to Register-Transfer Levels. In L. Scheffer, L. Lavagno, and G. Martin, editors, *EDA for IC Implementation, Circuit Design, and Process Technology*, volume 2. Taylor & Francis, 2006.

[63] M. Moudgill, J.-D. Wellman, and J.H. Moreno. Environment for PowerPC microarchitecture exploration. *Micro*, 19:15 – 25, 1999.

[64] R. Mukherjee and S. Memik. Systematic Temperature Sensor Allocation and Placement for Microprocessors. In *Proceedings of the Design Automation Conference*, pages 542 – 547, 2006.

[65] F. Najm. Power Estimation Techniques for Integrated Circuits. In *International Conference on Computer Aided Design*, pages 492–499, 1995.

[66] Farid N. Najm. A Survey of Power Estimation Techniques in VLSI Circuits. *IEEE Transactions on VLSI Systems*, 2:446–455, 1994.

[67] S. Narasimhan, X. Wang, D. Du, R.S. Chakraborty, and S. Bhunia. TESR: A robust Temporal Self-Referencing approach for Hardware TROJAN detection. In *International Symposium on Hardware-Oriented Security and Trust*, pages 71–74, 2011.

[68] K. Natarajan, H. Hanson, S.W. Keckler, C.R. Moore, and D. Burger. Microprocessor Pipeline Energy Analysis. In *International Symposium on Low Power Electronics and Design*, pages 282–287, 2003.

[69] A. N. Nowroz, R. Cochran, and S. Reda. Thermal Monitoring of Real Processors: Techniques for Sensor Allocation and Full Characterization. In *Proceedings of the Design Automation Conference*, pages 56–61, 2010.

[70] A. N. Nowroz and S. Reda. Thermal and Power Characterization of Field-Programmable Gate Arrays. In *Proc. ACM International Symposium on Field Programmable Gate Array*, pages 111–114, 2011.

[71] A. N. Nowroz, G. Woods, and S. Reda. Improved Post-Silicon power Modeling Using AC Lock-In Techniques. In *ACM/IEEE Design Automation Conference*, pages 101–106, 2011.

[72] A. N. Nowroz, G. Woods, and S. Reda. Power Mapping of Integrated Circuits Using AC-Based Thermography. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, PP(99):1–1, 2012.

[73] M. Orshansky and S. Nassif. *Design or Manufacturability and Statistical Design: A Constructive Approach*. Springer, 2007.

[74] 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.

[75] E. Pop, S. Sinha, and K. Goodson. Heat Generation and Transport in Nanometer-Scale Transistors. *Proceedings of the IEEE*, 94(8):1587–1601, 2006.

[76] M. Potkonjak, A. Nahapetian, M. Nelson, and T. Massey. Hardware TROJAN Horse Detection using Gate-level Characterization. In *Design Automation Conference*, pages 688–693, 2009.

[77] 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.

[78] Z. Qi, B. H. Meyer, W. Huang, R. J. Ribando, K. Skadron, and M. R. Stan. Temperature-to-Power Mapping. In *International Conference on Computer Design*, pages 384–389, 2010.

[79] R. Rad, X. Wang, M. Tehranipoor, and J. Plusquellic. Power Supply Signal Calibration Techniques for Improving Detection Resolution to Hardware Trojans. In *International Conference on Computer-Aided Design*, pages 632–639, 2008.

[80] Reza Rad, Jim Plusquellic, and Mohammad Tehranipoor. Sensitivity Analysis to Hardware TROJANS using Power Supply Transient Signals. In *International Workshop on Hardware-Oriented Security and Trust*, pages 3–7, 2008.

[81] S. Reda. Thermal and Power Characterization of Real Computing Devices. *IEEE Journal on Emerging Topics in Circuits and Systems*, 1(2), 2011.

[82] S. Reda, R. Cochran, and A. N. Nowroz. Improved Thermal Tracking for Processors Using Hard and Soft Sensor Allocation Techniques. *IEEE Transactions on Computers*, 60(6):841–851, 2011.

[83] S. Reda and S. Nassif. Analyzing the Impact of Process Variations on Parametric Measurements: Novel Models and Applications. In *IEEE Design, Automation and Test in Europe Automation*, pages 375–380, 2009.

[84] S. Reda, A. N. Nowroz, R. Cochran, and S. Angelevski. Post-Silicon Power Mapping Techniques for Integrated Circuits. *ElSevier VLSI Integration Journal*, 46(1):69–79, 2013.

[85] J. Renau. Personal Communication, 2009.

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

[87] E. Rotem, J. Hermerding, C. Aviad, and C. Harel. Temperature Measurement in the Intel Core Duo Processor. In *Proceedings of the International Workshop on Thermal Investigations of ICs*, pages 23–27, 2006.

[88] K. Roy, S. Mukhopadhyay, and H. Mahmoodi-Meimand. Leakage Current Mechanisms and Leakage Reduction Techniques in Deep-Submicrometer CMOS Circuits. *Proceedings of the IEEE*, 91(2):305–327, 2003.

[89] A. Shakouri. Nanoscale Thermal Transport and Microrefrigerators on a Chip. *Proceedings of the IEEE*, 94(8):1613–1638, 2006.

[90] L Shang, A Kaviani, and K Bathala. Dynamic Power Consumption in Virtex-II FPGA Family. In *International Symposium on Field Programmable Gate Arrays*, pages 157–164, 2002.

[91] A. Shen, A. Ghosh, S Devadas, and K. Keutzer. On Average Power Dissipation and Random Pattern Testability of CMOS Combinational Logic Networks. In *IEEE/ACM International Conference on Computer-Aided Design*, pages 402–407, 1992.

[92] P. Shivakumar, N. P. Jouppi, and P. Shivakumar. CACTI 3.0: An Integrated Cache Timing, Power, and Area Model. Technical report, 2001.

[93] G. Spirakis. Designing for 65nm and Beyond: Where is the Revolution? In *Electronic Design Process Symposium*, 2005.

[94] 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.

[95] G. Strang. *Computational Science and Engineering*. Wellesly-Cambridge Press, first edition, 2007.

[96] M. Tehranipoor and F. Koushanfar. A Survey of Hardware TROJANS: Taxonomy and Detection. *IEEE Design Test of Computers*, 27(1):10–25, 2010.

[97] M. Tehranipoor, H. Salmani, X. Zhang, X. Wang, R. Karri, J. Rajendran, and K. Rosenfeld. Trustworthy Hardware: TROJAN Detection and Design-for-Trust Challenges. *IEEE Computer Magazine*, 44(7):66–74, 2011.

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

[99] S. Wei, S. Meguerdichian, and M. Potkonjak. Gate-level Characterization: Foundations and Hardware Security Applications. In *Design Automation Conference*, pages 222 – 227, 2009.

[100] S. Wei and M. Potkonjak. Scalable Consistency-based Hardware TROJAN Detection and Diagnosis. In *International Conference on Network and System Security*, pages 176–183, 2011.

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

[102] F. Wolff, C. Papachristou, S. Bhunia, and R. Chakraborty. Towards TROJAN-Free Trusted ICs: Problem Analysis and Detection Scheme. In *Design, Automation and Test in Europe*, pages 1362–1365, 2008.

[103] B. Wong, A. Mittal, Y. Cao, and G. W. Starr. *Nano-CMOS Circuit and Physical Design*. Wiley-Interscience, 2004.

[104] Jian Yang, D Zhang, A.F Frangi, and Jing yu Yang. Two-dimensional PCA: A New Approach to Appearance-based Face Representation and Recognition. *Pattern Analysis and Machine Intelligence*, 26(1):131 – 137, 2004.

[105] W. Ye, N. Vijaykrishnan, M. Kandemir, and M. J. Irwin. The Design and Use of SimplePower: A Cycle-Accurate Energy Estimation Tool. In *Design Automation Conference*, pages 340–345, 2000.

[106] J. Zheng and M. Potkonjak. Securing Netlist-Level FPGA Design through Exploiting Process Variation and Degradation. In *International Symposium on Field-Programmable Gate Arrays*, pages 129–139, 2012.