HotSpot—A Chip and Package Compact Thermal Modeling Methodology for VLSI Design

A dissertation presented

by

Wei Huang

to

the faculty of the School of Engineering and Applied Science

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

The Charles L. Brown Department

of Electrical and Computer Engineering

University of Virginia

Charlottesville, Virginia

January 2007

Approval Sheet

The dissertation has been submitted in partial fulfilment of the
requirements for the degree of Doctor of Philosophy

_____
Wei Huang

The dissertation has been read and approved by the examining committee

_____
Dr. Mircea Stan (advisor)

_____
Dr. Lloyd Harriott (committee chair)

_____
Dr. Kevin Skadron

_____
Dr. John Lach

_____
Dr. Robert Ribando

Accepted for the School of Engineering and Applied Science

_____
Dean, School of Engineering
and Applied Science

(January, 2007)

To Shuang and my parents.

A necessary condition for a successful researcher—

"He who understands is not as good as he who likes;
He who likes is not as good as he who enjoys."
– Confucius

Acknowledgments

First of all, my greatest thanks to my wife Shuang Niu, the most important person in my life, for always being there for me with love and support throughout the years. Aside from Shuang, I am quite sure my parents are also the happiest people to know that I am now officially a Dr. Huang. I would like to thank my parents for the delicious dishes, the warm woven sweaters, the beautiful wooden swords and guns, the interesting Chinese chess and go games, etc, etc. My old buddies from Jilin, Hefei and Shanghai are also in my mind despite the half-Earth distance between us.

During my five years in the prestigious University of Virginia—Thomas Jefferson's University—I have luckily had the opportunity to meet and work with so many great people and there are so many precious moments that I will cherish for ever.

Academically, I am lucky to meet two of the best professors one can imagine. I would like to first express my sincere gratitude to my advisor, Professor Mircea Stan. He has been truly a teacher, an advisor and a friend. I am also deeply indebted to my co-advisor, Professor Kevin Skadron. Kevin has been always available to offer invaluable help and advice.

There are also colleagues who work shoulder by shoulder with me on the HotSpot project. They are Karthik Sankaranarayanan, Sivakumar Velusamy, Shougata Ghosh, and David Tarjan. It has been a wonderful experience to collaborate with these amazing guys under the guidance of Professor Skadron and Professor Stan.

Certainly I should also thank my friends from the HPLP VLSI Lab for sharing the same office, having fun together, and learning from each other all the time. I would also like to thank Dr. Zhijian Lu for exciting discussions and collaborations, and my manager Dr. Thomas Bucelot and my mentor Mr. Scott Neely for their help during my internship at

IBM Research.

Finally, my thanks also go to my PhD committee members, Professor John Lach for his caring advice on my stuttering, Professor Lloyd Harriott for his excellent and enlightening lectures on microelectronic devices and fabrication processes, and Professor Robert Ribando for his insightful suggestions on my research.

Research advisor

Mircea R. Stan

Author

Wei Huang

HotSpot—A Chip and Package Compact Thermal Modeling Methodology for VLSI Design

Abstract

To deal with the challenges of managing the ever-increasing power densities and temperatures of VLSI systems, thermal models have become more and more important for performing thermal analysis during design time. Due to their lumped thermal R-C network nature, compact thermal models are simple and efficient for thermal analysis in VLSI design, especially during early design stages, where detailed layout information is not available. Additionally, temperature predictions from compact thermal models can also be utilized to develop architecture-level run-time thermal management techniques. However, existing compact thermal modeling approaches have a number of limitations, such as not providing detailed temperature distributions, or using unrealistic assumptions for package modeling, or being restricted to particular packages, or being not parameterizable.

This dissertation presents *HotSpot*—a novel modeling methodology for developing compact thermal models based on the popular stacked-layer packaging scheme in modern VLSI systems. In addition to modeling silicon and packaging layers, HotSpot also includes a high-level on-chip interconnects self-heating power and thermal model such that the thermal impact of interconnects can also be considered during early design stages. The HotSpot compact thermal modeling approach is especially well suited for pre-RTL and pre-synthesis thermal analysis and is able to provide detailed static and transient temperature information across the die and the package, it is also computationally efficient and fully parameterizable. HotSpot compact thermal models have been validated with detailed numerical simulations, thermal test chip and temperature measurements of an FPGA-based SoC design. It is also shown to be boundary condition independent. Example applications demonstrate how our modeling approach helps in developing dynamic thermal management techniques and provides more accurate design estimations for temperature, leakage power, delay and reliability in a temperature-aware design flow. The presented HotSpot modeling method can also be easily extended to model emerging packaging schemes such as stacked chip-scale packaging (SCP) and 3-D integration.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

"Icarus, my son, I charge you to keep at a moderate height, for if you fly too low the damp will clog your wings, and if too high the HEAT of the sun will melt them."

—Daedalus, the pioneer pilot in the Greek mythology

The ever-increasing complexity and computing capability of contemporary very large scale integration (VLSI) systems have posed grand challenges to designers and manufacturers in multiple fronts. One of these challenges is to remove the heat which is an undesirable byproduct of the faster electronic devices and systems we enjoy in our daily life. The goal of heat removal is to keep the **operating temperature** *within an acceptable range where the VLSI systems can function as expected and reliably. Facing the thermal grand challenge, an efficient and accurate thermal modeling methodology is indispensable during the design process of a VLSI system in the deep submicron, or rather, nano era. This chapter motivates this dissertation research, which focuses on a chip and package compact thermal*

*modeling approach—***HotSpot**. *We begin by a brief review of the history and the current status of the thermal solutions in integrated circuit design, we then state the scope of this dissertation research.*

1.1 Historical Perspective

1.1.1 In the Past

From the inception of the microelectronic industry, heat dissipation and the negative impacts of high temperature have been a concern. In the mid-20th century, early electronic devices made of vacuum tubes [79] for high-power radio-frequency signal transmissions were often liquid cooled to remove the excessive heat dissipated at the anode of the tube. Because there were not strict space constraints at that time, providing a large room for the expensive cooling facility was normal.

With the adoption of the semiconductor devices following the invention of the first transistor [72], bipolar integrated circuits (IC) with significantly lower power consumption and higher speed became available, requiring only low-cost cooling solutions such as natural air convection. At that time, people had the false feeling that the monster of heat may have been put to sleep for ever.

This honeymoon period lasted for many years and was eventually interrupted by the insatiable desire of human beings for higher speed, which brought the power consumption, consequently heat dissipation of the bipolar ICs, to such a level that expensive cooling solutions such as liquid cooling and huge heatsink with forced convection were resurrected. Luckily, however, the thermal issue was suppressed again with the advent of CMOS tech-

nology in the mid-eighties. CMOS has much lower power consumption than its bipolar counterpart at the same level of performance. This is one of the major reasons that CMOS continues to replace the bipolar technology nowadays [41].

1.1.2 Current Status

Currently, CMOS is the dominant technology in the semiconductor industry. Historically, CMOS has the advantage of low-power operation compared to bipolar and earlier technologies. However, since the early nineties when personal computers have become wide-spread and the Internet can be easily accessed, speed and performance have been the main target of the VLSI designs. In order to achieve the ever-increasing demand for higher performance and more functionality, Moore's Law [58] has been followed by continued technology scaling, leading to higher and higher power consumption and as a result more and more heat generated from the silicon chip.

While technology scaling reduces power consumption of individual devices, the power consumption and power density of an entire chip keep increasing as a result of non-ideal technology scaling. Although the new transition to multiprocessor systems, e.g. chip multiprocessor (CMP), could somewhat alleviate the growth rate in power consumption by utilizing parallelism to achieve same performance and makes the power distribution more uniform across the silicon die, this trend is likely to continue as long as the desire for speed and performance does not diminish. Nowadays, high-performance microprocessors have already consume well above 100 Watts of total power [60], and local hot spot power density[1] has been greater than $10^6 W/m^2$, a number close to the power density of a nuclear

[1] In IC design, power density usually refers to the heat flux or the power dissipation over a unit surface area, thus has the unit of W/m^2.

reactor [7]. All these trends have rendered temperature a serious concern, if not a show stopper, for the sustained scaling of CMOS technology into the nanometer regime. Therefore, a proper methodology of modeling the detailed operating temperature distribution for silicon and its package is indispensable for contemporary and future VLSI designs.

1.2 Dissertation Scope

1.2.1 Scope

This dissertation presents a compact thermal modeling methodology—HotSpot, which simplifies the 3-D partial differential heat diffusion equation with a compact network of thermal resistances and thermal capacitances representing the heat transfer paths through both the silicon die and the package components. HotSpot is an accurate and efficient by-construction thermal modeling approach which is fully parameterized. HotSpot takes into account the heat spreading effect by modeling lateral heat transfer using lateral thermal resistances. It also takes care of the transient evolution of temperature and the time-domain temperature filtering effect by using thermal capacitances to model the transient responses. The spatial temperature filtering effect is taken care of in HotSpot by the spatial granularity analysis, which accurately model the spatial temperature gradient. All these topics are discussed in details in the coming chapters of this dissertation.

1.2.2 Research Theses

In this dissertation research, we state that

1. A parameterized compact thermal modeling method with proper considerations of

the heat transfer in both the silicon die and the thermal package is needed in order to deal with the thermal challenges in modern VLSI design. A modeling methodology—*HotSpot*—is proposed for generating compact thermal models that can be used for VLSI design, especially early design stages where detailed layout is not available. With this method, reasonably accurate spatial and temporal temperature variations of the silicon die as well as the package can be quickly obtained to help efficient design decisions during the entire design process. The modeling method is based on the stacked-layer packaging configuration that is predominant in the modern VLSI packages and is potentially extendable to future packaging schemes.

2. The accuracy and efficiency of the thermal modeling methodology is greatly affected by the thermal modeling granularity. In this dissertation, we analytically investigate the relationship between the spatial modeling granularity in the compact thermal model and the accuracy of the model. For thermal analysis during the design process, it is important to find the right granularity in order to achieve faster computation without sacrificing accuracy. Additionally, it also helps to find the proper granularity of power consumption. For example, for an accurate thermal analysis, do we need to model the power of a single cache line, or the entire cache block, or the entire chip?

3. For early analysis of thermal impacts on interconnect-related performance, power grid IR drop, and electromigration, an early-stage on-chip interconnect self-heating power and thermal model is needed. This model should be used together with the silicon and package thermal models to achieve useful thermal information for on-chip interconnects.

4. A chip and package compact thermal model should have useful applications in general research areas such as microarchitecture, high-performance and low-power circuit design, package design, and IC reliability.

1.3 Dissertation Outline

The remainder of this dissertation is organized as follows.

Chapter 2 lays the necessary background to understand the topics discussed in this dissertation. It also points out the thermal challenges facing the designers in the deep submicron and nanometer CMOS era.

In Chapter 3, we look through the literature for existing works related to the scope of this dissertation and identify the novelties and contributions of this dissertation research.

Following that, Chapter 4 describes the details of different parts of the HotSpot models, including the primary and secondary heat transfer paths, and the early-stage interconnect thermal model. It also compares HotSpot with existing thermal modeling approaches, and listed the advantages and disadvantages of HotSpot.

Chapter 5 discusses the important topic of the parametrization of compact thermal models, which distinguishes HotSpot from the existing compact package thermal models.

Chapter 6 describes two approaches for the spatial granularity issue of thermal modeling. Chapters 4-6 are the core parts of this dissertation.

Chapter 7 presents a number of validation efforts of HotSpot. The HotSpot chip and package thermal models have been validated against a finite-element package, a thermal testing chip and an FPGA platform. It is also shown to be boundary condition independent (BCI) by comparing with existing BCI compact thermal models. The early-stage intercon-

nect thermal model is also validated against results from a published finite element model.

Chapter 8 presents several successful applications of the HotSpot thermal models that are directly related to this dissertation research in various research areas, such as temperature-aware design, dynamic thermal management, temperature-aware reliability analysis, thermal package design, and thermally self-consistent leakage power calculations, etc.

Finally, Chapter 9 concludes this dissertation and points out possible future improvements and applications of the HotSpot thermal modeling method.

Chapter 2

Thermal Fundamentals and Challenges

Some reviews of the two fundamental physical mechanisms, heat generation and heat transfer, are needed in order to fully understand the topics in this dissertation. This chapter offers the related background knowledge, which also leads to the thermal challenges that are nowadays faced by circuit designers, computer architects and package designers.

2.1 Thermal Fundamentals

2.1.1 Temperature and Heat Transfer

Let us start with the definition of temperature. In thermodynamics, from a *macroscopic* viewpoint, we can define temperature as the property that "... is shared by two systems, initially at different states, after they have been placed in thermal contact and allowed to come to thermal equilibrium" [87]. Another more satisfying definition of temperature from the *microscopic* point of view for the special case of an ideal gas is: "Temperature is directly proportional to the square of the mean molecular speed. Higher temperature means faster

moving molecules" [87]. From these definitions, we can see that temperature is a state variable reflecting the level of the internal energy possessed by a system.

Another fundamental concept is the difference between heat transfer and thermodynamics. "Thermodynamics deals with systems in equilibrium and can be used to determine the energy required to change a system from one equilibrium state to another. Thermodynamics cannot, however, enable us to determine the *rate* at which the change occurs". On the contrary, heat transfer analysis can tell us the time-dependent process of the state change [27].

2.1.2 Heat Generation in CMOS Integrated Circuits

In this dissertation, we are investigating thermal issues that are caused by the heat generated in a chip of integrated-circuits. The basic construction unit of CMOS integrated circuits is a MOS transistor. Each transistor can be turned on or off like a switch depending on the voltage difference between the gate terminal and the source terminal. At a higher abstraction level, a number of transistors are connected together in a particular topology to form a logic gate. Each gate has its own input signals and output signals. Computation is the process to get the correct output voltage level for a given binary combination of the input signals. Let us use the CMOS inverter gate in Fig. 2.1 as an example. If the voltage level at Node A has a transition from $HIGH$ to LOW (or logic 1 to 0), and Node Y is LOW initially, for the inverter to perform the correct computation, the voltage level at Y has to transit from LOW to $HIGH$ (or logic 0 to 1) after some amount of time delay. During the low-to-high transition at Node Y, the load capacitance C_L has to be charged from 0 to V_{dd}. The charging path is from the power supply (V_{dd}) through the turned-on PMOS

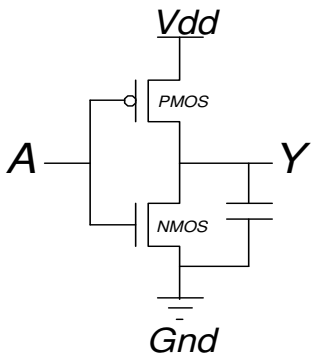

Fig. 2.1: A CMOS inverter gate.

transistor to Node Y, whereas the NMOS transistor is turned of because its gate-to-source voltage is zero. To the first order, the PMOS can be modeled as a resistor when it is turned on. Thus the inverter can be simplified as a first-order R-C circuit as shown in Fig.2.2, where R_P is the equivalent resistance of the turned-on PMOS.

From basic circuit theory, we know that the total energy drawn from the power supply for this voltage transition is $C_L V_{dd}^2$. But the energy actually stored in the capacitor is $\frac{1}{2} C_L V_{dd}^2$, only half of the total energy. Where does the other half of the total energy go? It is dissipated in the form of *Joule heat* in the resistor R_P. Later, Node Y may be discharged to ground if the input signal A makes a low-to-high transition, thus the remaining half of the total energy is dissipated as heat in the resistor R_N, as shown in Fig. 2.3. Therefore, every switching event as a result of computation draws some amount of energy from the power supply, and this energy is eventually transformed into heat dissipation.

During each switching event, because it takes finite amount of time for the input transition, for a short while both the pull-up and the pull-down networks are partially turned on. For example, in the case of an inverter, during the transition of the input, both the PMOS and the NMOS are partially turned on, causing so-called “short-circuit” current flowing

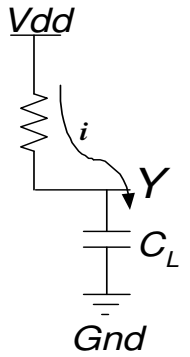

Fig. 2.2: The low-to-high transition at the output node—load capacitor is charged, half energy is dissipated as heat in R_P.

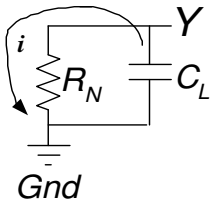

Fig. 2.3: The high-to-low transition at the output node—load capacitor is discharged, the remaining half energy is dissipated as heat in R_N.

from V_{dd} to Gnd through the transistors. This short-circuit current also causes some heat generation in the CMOS circuits.

Another source of heat generation in VLSI systems is due to the leakage energy. CMOS transistors are not ideal switches. They still conduct some small amount of current even if they are supposed to be off, this current is called leakage current. Like the active switching current, the leakage current also moves charges between power supply and ground, thus it also draws energy from the power supply. These charges are wasted without performing useful computation and the associated energy moving them is dissipated as heat through the resistance in their flow path.

In addition to the heat generated inside the transistors, Joule heating, expressed by $P = I^2R$, also occurs when electrical current flows through on-chip metal interconnects that connect the transistors. This is because the interconnects are not ideal electrical conductors and have finite amount of resistance R and hence joule heating.

In summary, heat is generated from the silicon active surface due to two factors—active switching and leakage. All the energy consumed by the integrated circuit is first dissipated in the form of heat in the transistors and interconnects, and are eventually removed to the environment by heat transfer. While power, P, is the rate of energy consumption, heat generation rate (or heat dissipation rate), Q, is the amount of heat generated or dissipated in unit amount of time. In this dissertation, sometimes when we use "heat", what we really mean is the heat generation rate. The actual meaning usually can be told from the context.

2.1.3 Heat Transfer Theory

All the heat generated in the integrated circuits must be removed, or transferred, to the ambient environment. Otherwise, the operating temperature will accumulate and cause malfunction and eventually the destruction of the system.

"Heat transfer is the transport of thermal energy from one region to another. In order for heat transfer to occur, there must be a *temperature difference* between the two regions" [27]. From the first law of thermodynamics, i.e. the conservation of energy, we know that the heat given by the hot region has to be equal to the heat absorbed by the cold region. In addition, the second law of thermodynamics states that heat must be transferred from hot region to cold region, i.e. heat flows in the direction of decreasing temperature.

There are three modes of heat transfer—conduction, convection and radiation [27].

1. *Conduction* is the heat transfer in solids. The actual mechanism includes the change of momentum and energy between molecules or, in the case of metal, the free electrons. This is the major heat transfer mode considered in this dissertation.

2. *Convection* is the heat transfer mode in which heat is transferred between a solid surface and a moving fluid touching the solid surface. If the fluid is forced to flow over the solid surface, it is called forced convection. Otherwise, it is called natural convection. In this dissertation, convection is only considered at the interface of the package and the ambient air or other fluids.

3. *Radiation* is the heat transfer mode in which thermal energy is transported by electromagnetic waves. Radiation does not need a medium because an electromagnetic wave does not need one. An example of radiation is the heat we have been receiving continuously from the sun, which is the main source of energy near the surface of our planet Earth. Because the amount of heat transfer due to radiation is extremely small for the interested temperature range and the volume of the silicon chip and package that we investigate, we neglect radiation in this dissertation.

The governing equation of heat conduction is the Fourier's Law:

$$q = -k\frac{dT}{dx} \tag{2.1}$$

Eq. (2.1) is the one-dimensional form of the Fourier's Law, where q is the heat flux (in W/m^2), k is the thermal conductivity of the material (in $W/(m{\cdot}K)$). What Eq. (2.1) says is that the heat flux, q (the flow of heat per unit area and per unit time), at a point in a medium is directly proportional to the temperature gradient at that point. The minus sign indicates

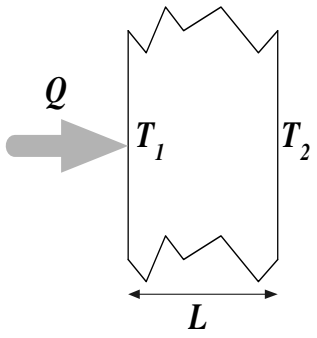

Fig. 2.4: One-dimensional heat conduction—the Fourier Law.

that heat flows in the direction of decreasing temperature. If we write $q = Q/A$, where Q is the heat transfer rate, A is the heat conducting area, and refer to Fig. 2.4, Eq. (2.1) becomes

$$Q = -kA\frac{T_2 - T_1}{L} \tag{2.2}$$

and, if we define *thermal resistance* $R_{th} = (T_1 - T_2)/Q$, i.e. the temperature drop divided by the heat transfer rate, we get

$$R_{th} = (T_1 - T_2)/Q = \frac{1}{k}\frac{L}{A} \tag{2.3}$$

Here, we discover a resemblance between Eq. (2.3) and the well-known Ohm's Law in the electrical circuit theory:

$$R = (V_1 - V_2)/I = \rho\frac{L}{A} \tag{2.4}$$

Therefore, we have the following interesting duality between electrical and thermal phenomena—thermal resistivity ($1/k$) vs. electrical resistivity (ρ); temperature difference (ΔT) vs. voltage difference (ΔV); heat transfer rate (Q or power P) vs. electrical current (I), and finally thermal resistance (R_{th}) vs. electrical resistance (R).

Heat conduction is also a transient process, a more general equation which also consid-

ers time is the heat diffusion equation:

$$\rho c_p \frac{\partial T(x,y,z,t)}{\partial t} = \nabla \cdot [k(x,y,z,T) \nabla T(x,y,z,t)] + g(x,y,z,t) \tag{2.5}$$

where ρ is the density of the material (kg/m^3), not the electrical resistivity, and g is the volume power density of the heat source(s) (W/m^3), c_p is the specific heat ($J/(kg{\cdot}^\circ C)$). While thermal conductivity k actually is a function of location and temperature, we can assume it is isotropic and temperature independent, as this is mostly true for the materials and the temperature range that we are interested in this dissertation.

For the steady-state case, the ($\partial T/\partial t$) term in Eq. (2.5) becomes zero. We can verify that at steady state, the one-dimensional form of the heat diffusion equation can be reduced to the Fourier's Law, Eq. (2.1).

If we assume both g and k are constant, Eq. (2.5) can be rearranged to the following form if we write it in one-dimensional form and integrate both sides by the integral variable x from 0 to L (the length of the material), and notice that $g{\cdot}L = Q/A = q$, the heat flux, we get

$$(\rho c_p AL) \frac{dT(t)}{dt} = kA \frac{\Delta T(t)}{L} + Q \tag{2.6}$$

The first term of the right-hand side of Eq. (2.6) is the heat transferred through the thermal resistance R_{th}, similar to that in Eq. (2.3). Note that $\Delta T = T_2 - T_1$, and move this term to the other side of the equation, we get

$$C_{th} \frac{dT(t)}{dt} + \frac{T_1 - T_2}{R_{th}} = Q \tag{2.7}$$

where $C_{th} = \rho c_p AL = c_p \rho V$ is defined as *thermal capacitance* or thermal mass, V is the volume of the material.

From the electrical circuit theories, we know that $C\frac{dV(t)}{dt} = i_c(t)$, meaning the current flow through an electrical capacitor equals the product of its capacitance and the first derivative of the voltage difference across it. This exactly resembles the first term on the left-hand side of Eq. (2.7). This is the reason we define C_{th} as thermal capacitance. Thermal capacitance describes the heat absorbing capability of a material, while electrical capacitance describes the ability of accumulating electrical charges of a material. Equation (2.7) states that the heat flowing through the thermal capacitance (the AC component) plus the heat flowing through the thermal resistance (the DC component) equals the total heat flowing through the material.

Table 2.1 summarizes the duality between thermal and electrical phenomena. We use this duality to derive the compact thermal resistance and capacitance network for HotSpot in later chapters.

As a side note, the above heat conduction equations are only valid in the *macroscopic* world. When the dimension of interest comes to the nanoscale, such as the phonon-phonon mean free path (about 300nm for silicon), quantum effects kick in. In that case, the phonon Boltzmann Transport Equation (BTE) should be used instead [65] for thermal modeling and analysis at the nanoscale transistor level.

Thermal simulations for individual nanoscale transistors using BTE suggest that the maximum temperature occurs near the drain terminal of a nanoscale transistor [65]. This phenomenon cannot be accurately observed by traditional macroscopic equations such as the Fourier's Law and the heat diffusion equation. However, using these macroscopic equations is enough for modeling the average thermal impacts on properties such as average carrier mobility, threshold voltage and metal resistivity for individual transistors and inter-

Thermal quantity	unit	Electrical quantity	unit
Q, Heat transfer rate, power	W	I, Current	A
T, Temperature difference	K	V, Voltage difference	V
R_{th}, Thermal resistance	K/W	R, Electrical resistance	Ω
C_{th}, Thermal capacitance	J/K	C, Electrical capacitance	F

Table 2.1: Duality between thermal and electrical quantities.

connects. In Chapter 6, we will also see that the granularity needed for accurate *macroscopic* thermal analysis at the chip level is much greater than the range where the quantum effect governs. Therefore, we do not consider the quantum effects in this dissertation.

2.2 Thermal Challenges

Having reviewed the necessary background, we now examine the thermal challenges for deep submicron and nanoscale CMOS technologies.

2.2.1 Thermal Design Power

Along with technology scaling and following Moore's Law, the overall power dissipation has been increasing steadily. An example is the reported Thermal Design Power (TDP) for generations of Intel's microprocessors [57] and many other sources, such as ITRS predictions [2]. For the circuits to function correctly at the specified frequency and for the silicon and package to survive the elevated thermal stress, the temperature at the active silicon surface, the "junction temperature" T_j, has to be less than a certain specified value

T_{j0}. Notice that the temperature spatial distribution across the silicon die and the package are ignored in the definition of junction temperature, which is a simplification of reality. The main task for a thermal engineer is to make sure that T_{j0} is not exceeded for a specified TDP, thus the lumped thermal resistance from the silicon active surface to the ambient should be at most

$$R_{th} = (T_{j0} - T_a)/\text{TDP} \tag{2.8}$$

where T_a is the ambient temperature.

From Eq. (2.8), it is obvious that with an increasing TDP, in order not to exceed T_{j0}, the thermal resistance of the thermal package has to be lower, requiring a more advanced and of course more expensive thermal package. This is definitely a big concern in term of cost increase of the entire system. In addition, the newly-developed advanced cooling system are usually less reliable, thus posing more reliability issues to the system. Furthermore, if the system is designed without proper thermal considerations, the TDP can be so high that even the best available packaging solutions do not meet the thermal requirement.

Simply meeting the thermal requirement posed by the TDP may not be enough to guarantee the reliable operation of a system. A common definition of TDP is the maximum amount of power that the thermal solution in a computer system is required to dissipate[1]. However, this definition is unclear—does "maximum" mean the worst-case power (or heat) that a system can potentially encounter or the "typical" maximum power that the designers have tested. In the extreme case, a vicious programmer can always write some "thermal virus" that dissipates more power than the TDP listed in the datasheet [19].

[1] In this dissertation, we will focus on a single VLSI chip and its thermal package, whereas a computer system usually also includes other components such as Multi-chip Modules (MCM), other ICs on the printed-circuit board (PCB), the chassis, the fan and the box frame, etc.

Different IC companies have slightly different definitions of TDP. For example, Intel's definition is: "TDP is defined as the worst-case power dissipated by the processor while executing publicly available software under normal operating conditions, at nominal voltages that meet the load line specifications"[2], which accounts for the maximum design power for *typical* workloads. Another interpretation of Intel's TDP is "... the maximum sustained power dissipated by the microprocessor, across a set of realistic applications" [57]. The word "sustained" is necessary because if the worst-case power (greater than TDP) is of a burst nature and lasts for a brief period that is shorter than the thermal time constant of silicon, i.e. it is not sustained, the temperature will begin to fall off before it reaches its potential maximum value. This is the time-domain temperature "filtering" effect that helps in thermal design to allow burst activities whose power oscillation period is much less than the thermal time constant of the silicon die [57, 82]. Characterizing the burst activities of a design usually falls into the the responsibilities of chip architects and circuit designers.

In summary, the total system power, either the thermal design power or the absolute worst-case total power, together with their transient impact on the junction temperature have become a challenge not only to the thermal package designers, but also to the VLSI designers such as circuit designers and computer architects.

2.2.2 Power Density

Another thermal challenge is the ever-increasing power density due to the non-ideal threshold voltage scaling of the CMOS technologies. Nowadays, local power density of high-performance microprocessors has exceeded that of a nuclear reactor ($10^6 W/m^2$), and

[2]http://download.intel.com/design/celect/854/D18741.pdf

is approaching the power density of a rocket nozzle ($10^7 W/m^2$) in the foreseeable future [7]. The reason we consider power density as a thermal challenge is that temperature is proportional to power density, not just power. Therefore, to reduce temperature, we can either reduce power or increase area or both [82]. One example of increasing area to reduce power density is the common use of heat spreader to deal with thermal effects. However, power density itself cannot be used as a proxy for temperature, i.e. power density modeling cannot replace temperature modeling, which will be discussed in Section 2.3. The direct consequence of the sky-rocketing power density is the local on-chip hot spots.

2.2.3 Local Hot Spots

Historically, for chip and package thermal modeling and design, thermal engineers used total power dissipation of a chip and use a single "junction temperature", T_j, to model silicon temperature. Although this approach is still being followed during thermal design for low-power ICs, it is not sufficient for high-performance or power-constrained designs. Increasingly non-uniform power dissipation across the chip leads to local hot spots and elevated temperature gradients across the silicon die. For example, in a 90nm Intel Itanium processor, even after stringent thermal management, local temperature can still be as high as 88$^\circ C$, while other parts of the die are relatively cool (61$^\circ C$) [64]. Therefore, a single value for the total power or the junction temperature without considering the spatial temperature distribution is certainly not enough. A thermal package designed for total power and average die temperature inevitably misses the local hot spots, resulting in reliability and performance degradations and potential thermal hazards. In addition, local hot spot temperature is expected to increase as a side-effect of technology scaling, which again poses

challenges to thermal engineers, circuit designers, as well as computer architects.

Aside from local hot spots in silicon, increasing current density and hence self-heating power density in on-chip interconnects also create local hot spots in the metal layers. Hot interconnect is becoming a serious problem as a result of the continued scaling and the introduction of low-κ dielectric in the back-end process technology, together with the elevated underlying silicon temperature. This is true because the rate of interconnect electromigration is dependent on temperature exponentially [17]. Additionally, we will see more constraints on power supply voltage IR drop due to the linearly increased metal resistivity at higher temperatures. According to [39], the interconnect temperature will be approaching 400-900$^{\circ}C$ at the future 22nm technology node if the thermal challenge on interconnects is not properly dealt with along the road.

2.2.4 Economical Constraints

A direct impact of the above-mentioned thermal challenges is the increasing cooling cost to computer systems and other microelectronic products. For high-performance processors, cooling solutions are rising at $1-3 or more per watt of heat dissipated [7, 26], not to mention that this may not include the extra cooling cost to deal with local hot spots. This exponential rising in cooling cost threatens the computer industry's ability to deploy new systems [78].

2.3 The Need to Directly Model Temperature

Since operating temperature is related to power and power density, why don't we use power, or more precisely, power density as a proxy of temperature, according to Eq. (2.3)? Why do we need to model temperature directly?

The fundamental reason is that temperature fluctuation is not simply proportional to the power consumption, neither the power density. There are other factors that significantly impact temperature distribution in space and time that are also needed to be taken care of. These factors include heat spreading and temporal and spatial temperature filtering effects. None of them is accounted for in Eq. (2.3). Therefore, temperature must be modeled directly in order to perform accurate thermal analysis during the design process.

Heat spreading happens when heat transferred from a small surface area to a larger one. Temperature filtering happens in the time domain where the long thermal time constant of silicon and package tends to filter out fast changes (high frequency component) in power and power density. Temperature filtering can also happen spatially where the power and power density change over a small dimension (high spatial frequency). All these effects can be modeled by solving the heat diffusion equation, Eq. (2.5). But directly solving the three-dimensional partial differential equation in Eq. (2.5) is a daunting task, if not impossible, without simplifications and numerical techniques.

The HotSpot thermal modeling methodology provides an efficient and accurate way to construct compact thermal R-C networks to simplify the heat diffusion equation.

Chapter 3

Compact Thermal Modeling and Related Work

There have been some existing work on thermal modeling at either the chip level or the package level. Package thermal models are usually compact thermal R-C networks with values extracted from detailed finite-element simulations, thus are not parameterizable. The existing chip-level thermal models are usually not very compact, therefore, they are computationally intensive. In addition, existing chip-level thermal models always neglect or over-simplify the thermal packages, leading to significant errors. In addition, none of the previous studies has analytically dealt with the thermal and power modeling granularity issue. In this chapter, we review the related work in the literature and point out the novelty and contributions of the HotSpot chip and package thermal modeling method.

3.1 Overview

As stated in earlier chapters, an unfortunate side effect of miniaturization and the continued scaling of CMOS technology is the ever-increasing power densities. The resulting difficulties in managing temperatures, especially local hot spots, have become one of the major challenges for designers at all design levels. High temperatures have several significant impacts on VLSI systems. First, the carrier mobility is degraded at higher temperature, resulting in slower devices for contemporary CMOS technologies. Second, leakage power is escalated due to the exponential increase of sub-threshold current with temperature. Third, the interconnect resistivity increases with temperature, leading to worse power grid IR drops and longer interconnect RC delays, hence causing performance loss and complicating timing and noise analysis. Finally, elevated temperatures can shorten interconnect and device lifetimes and package reliability can be severely affected by local hot spots and higher temperature gradients. For all these reasons, in order to fully account for the thermal effects, it is important to model temperature for VLSI systems in an accurate but also efficient way. For example, knowing the across-die temperature distribution at design time permits thermally self-consistent leakage power calculations in an iterative manner as shown in Fig. 3.1(a) [3, 28, 32]. Similarly, an efficient thermal model can also help to close the loop for temperature-aware performance and reliability analysis, as suggested in Fig. 3.1(b). In particular, it is crucial to take thermal effects into account as early as possible in the design flow, because optimal early and high-level thermally-related design decisions can significantly improve design efficiency and reduce design cost.

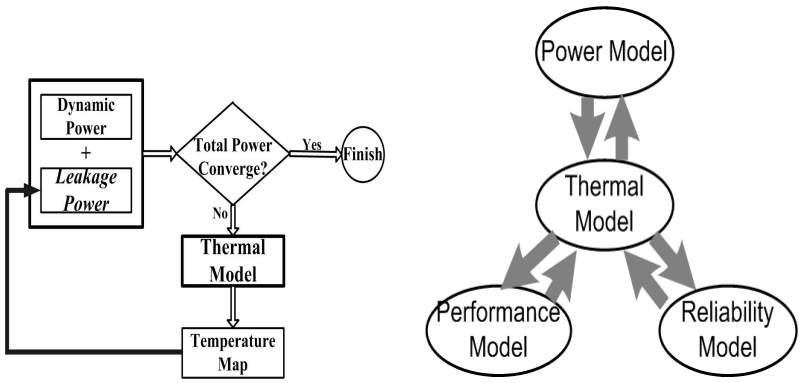

(a) (b)

Fig. 3.1: (a) A thermal model closes the loop for leakage power calculation [32]. (b) The role of temperature in power, performance and reliability models [36].

3.2 Compact Thermal Modeling

Obviously, it is impractical to accurately analyze thermal effects and model temperature distribution of a system together with the environment in their full details. Using numerical thermal analysis methods, such as the finite element method (FEM) and finite volume method (FVM), is a time-consuming process not suitable for design-time and run-time thermal analysis. In order to gain more insights of the thermal effects during early IC design stages, the trade-off solution is to build compact thermal models (CTMs) that give reasonably accurate temperature predictions with little computational effort at desired levels of abstraction [70].

Based on the well-known duality between thermal and electrical phenomena (Table 2.1), a compact thermal model is a lumped thermal R-C network, with heat dissipation modeled as current sources. The resulting thermal R-C networks are typically relatively small, can be solved for temperature very efficiently and introduce little computational overhead. Due to this computational efficiency, it is desirable to have compact thermal models for both

temperature-aware design and fast simulations of architecture-level dynamic thermal management techniques. Here, temperature-aware design refers to a design methodology which uses temperature as a guideline throughout the design flow. The resulting design can thus be thermally optimized, as it takes into account potential thermal limitations [36].

A top-down hierarchy of compact thermal models would be helpful for designers at different design levels [70]. There are several desired features that increase the usefulness of such compact thermal models at a particular design level.

1. *Detailed temperature distribution*: A compact thermal model should provide enough thermal information at the desired design level. For example, for package-level compact thermal models, previous studies [6, 67] have shown that the information of temperature distribution across the package is required. Using only a single junction-to-case thermal resistance will lead to an inferior package design; instead, multiple nodes are needed on the package surfaces. Similarly, a compact thermal model at the silicon die level should consist of enough nodes that give detailed temperature distribution information across the die. In addition, both static and transient temperatures should be modeled.

2. *Granularity*: A compact thermal model should model just at the granularity that is needed and hide the details of the lower levels, so that the compact thermal model itself is no more complex than necessary. Modeling at finer granularity introduces unnecessary details and makes the computation slower. For example, package-level compact thermal models, such as the DELPHI models [43, 47, 68], hide the lower level details of the package structures, including the die, the thermal attach, the solder balls and so on, mainly because these details are intellectual properties of the ven-

dors, but also because they would just increase the complexity of the model without significantly improving the simulation accuracy. Similarly, a compact thermal model at the die level should also hide the lower level details of the die, such as the actual circuit structures and physical layout.

3. *Parametrization*: A fully parameterized compact thermal model allows package designers, circuit designers and computer architects to explore new design alternatives and evaluate different thermally-related design trade-offs at their corresponding design levels before the actual physical designs are available. More importantly, with the aid of the parameterized compact thermal models, designers at different design levels can have more productive interactions and collaborations at early design stages of a microelectronic system. This leads to early discovery and considerations of potential thermal hazards of the system. True parametrization requires that the models be constructed based solely on design geometries and material properties.

4. *Boundary condition independence (BCI)*: A crucial feature of compact thermal models is boundary condition independence (BCI). By achieving BCI, the variation of the environment does not affect the actual model. The package-level compact thermal models in [47, 68] achieve BCI by finding a thermal resistance network with minimum overall error when applied to different boundary conditions [46, 47, 91]. In Section 7.5 we show that our physical compact thermal model is BCI with reasonable accuracy.

5. *Computational speed*: The structure of a compact thermal model needs to be relatively simple. Additionally, efficient algorithms should be developed to simulate the

compact thermal model with little computational overhead, so that thermal analysis at all design levels can be carried out efficiently.

3.3 Related Work

Existing work related to thermal modeling of VLSI can be divided into two major categories: (1) die-level full-chip thermal models and (2) package-level compact thermal models.

3.3.1 Die-Level Full-Chip Thermal Models

In recent years, there have been several published efforts in full-chip thermal modeling and compact thermal modeling for microelectronics systems. Wang et al. [92] present a detailed and stable die-level transient thermal model based on full-chip layout, solving temperatures for a large number of nodes with an efficient numerical method. The die-level thermal models by Su et al. in [84] and Li et al. in [51] also provide the detailed temperature distribution across the silicon die and can be solved efficiently, but with no information about the transient behavior. An earlier detailed full-chip thermal model by Cheng et al. [14] has an accurate 3-D model for the silicon and 1-D model for the package. A significant limitation of the above modeling approaches is the over-simplified thermal package model. For example, the thermal interface material and heat spreader that greatly affect the die temperature distribution are either not included or not properly modeled. The bottom surface of the silicon substrate is treated as isothermal by the above previous works, which significantly deviates from reality and therefore introduces errors. Additionally, these mod-

els are not quite suitable for early design stages, since their computation effort is non-trivial while fine-grained thermal analysis is not necessary when detailed layout information is not available. Finally, except [14], none of these models has shown validation from simulations with detailed numerical models or measurements from real designs. There are also some thermal models used in industry. However, to our best knowledge, most of them are either not publicly available or solely based on detailed layouts, empirically derived from measurements, and/or detailed numerical simulations.

3.3.2 Package-Level Compact Thermal Models

On the other hand, there have been abundant existing package-level compact thermal modeling methods in the literature. For example, the authors of [47, 48, 68] propose the DELPHI approach and introduce the important concept of boundary condition independence (BCI). The DELPHI approach extracts and optimizes a thermal resistance network from detailed model simulations under a set of standard boundary conditions. In [8] and [70], the authors independently propose alternative compact thermal modeling methods considering non-uniform boundary conditions and non-uniform boundary heat flux, resulting in much less optimization efforts compared to the DELPHI approach. There are also other modeling methods using model reduction techniques [18, 25, 71] or extracting transient thermal R-C network from real package temperature measurements [66] and detailed model simulations [62]. All these models are accurate to characterize existing designs, but are not fully parameterized to perform efficient explorations of new package design alternatives. Another limitation of these models is that they have only one or a few junction nodes representing die temperature distributions, thus are not suitable for die-level thermal

analysis during the design process.

3.4 This Work—Chip and Package Thermal Modeling

HotSpot belongs to the third category of thermal models, which model both die-level and package-level temperature distributions. HotSpot also improves our group's previous works such as [73] and [75] by modeling lateral heat spreading in silicon and package and eliminating fitting factors to make HotSpot fully parameterizable. Modeling detailed package temperature distribution is an important attribute for a thermal model to be more useful. This is because IC package components, especially the thermal interface material, heat spreader and heat sink, can greatly affect the die temperature and temperature distribution. Without modeling these components, a full-chip thermal model could lead to inaccurate temperature estimations, hence incorrect design decisions. In the published papers resulting from this dissertation, various aspects and related applications of the HotSpot thermal modeling approach, which is based on a block- or grid-like lumped thermal R-C network, are presented [31, 35, 36, 74, 77, 83]. This model can provide localized die-level full-chip temperature details as well as ways to model the temperature distributions of different packaging components for both the primary and secondary heat transfer paths. This thermal model can also be further extended to be flexible enough to model emerging packaging schemes such as stacked chip-scale packaging (SCP) [1] and 3D IC [5].

The HotSpot compact thermal modeling approach takes a structured assembly approach of constructing a physical compact thermal model by first modeling the silicon die and other packaging components as a collection of simple 3-D shapes and then assembling them into more complex compact thermal models according to the overall structure. This modeling

approach is fully parameterized and satisfy most of the above desirable features mentioned earlier in this chapter.

The HotSpot models also include a novel early-stage interconnect thermal and power model, which is suitable for pre-RTL and pre-synthesis thermal analysis of interconnect to detect potential problems such as exacerbated power supply IR drop, interconnect electromigration, etc, and to aid the architecture-level interconnect-related decisions.

The important issue of thermal modeling spatial granularity is answered analytically in this work. This analysis is the guidance for choosing proper level of abstraction in thermal and power modeling for the purpose of thermal analysis in VLSI design.

As example applications of the HotSpot thermal modeling method, in Chapter 8, we explicitly demonstrate the significant impact of detailed temperature distribution on the accuracy of leakage power estimations by an example of thermally self-consistent leakage calculations for an IBM POWER4-like microprocessor design.

We also, for the first time, show the circuit design and computer architecture communities that modeling package details is an indispensable part for a die-level thermal model to be really useful. With several example HotSpot thermal analysis regarding different properties of the thermal interface material (TIM), we show that omitting a package component in the thermal model can lead to significant errors in die temperature estimations.

Naturally, a full-chip and package thermal model can also act as a convenient medium for enhanced collaborations among circuit, architecture and package designers. This implies a design flow leading to early design evaluations from a thermal point of view. If potential thermal hazards are discovered early in the design process, different design tradeoffs can be carried out in an efficient way.

In the following chapters, we will present details of all the aspects of the HotSpot modeling method.

Chapter 4

HotSpot Modeling Details

In this chapter, we describe the details of the HotSpot modeling method, and itemize the differences of HotSpot from existing compact package thermal models and full-chip thermal models. We also summarize the advantages and disadvantages of the proposed HotSpot models.

4.1 Brief Outline of the Modeling Methodology

Let us first briefly outline the modeling methodology presented in this chapter:

1. We first present our compact thermal R-C modeling methodology which models multi-layered package structures. Typical components that are modeled include silicon die, heat spreader, thermal interface material, heat sink, etc. Each layer can be modeled with different levels of details, e.g. functional units or regular grid cells.

2. In Section 4.4, we present our interconnect self-heating power and thermal model that can be used during early design stages. This part includes:

– Interconnect self-heating power model. Three steps are needed: (a) average interconnect length and number of interconnects for each metal layer. (b) average interconnect RMS self-heating current for each metal layer. (c) total interconnect self-heating power in each metal layer. Signal interconnects and power supply network were considered separately in each step.

– Equivalent thermal resistance of wires and vias for each layer. In this part, wires and their surrounding dielectrics are modeled separately from the vias.

– With self-heating power model applied to the interconnect and via thermal resistance network, and adding the underlying thermal resistance network modeling silicon and package, the entire compact thermal model for a VLSI system is constructed.

4.2 Overview

A compact thermal modeling approach must have several features for it to be useful. First, it should provide detailed temperature distribution at the desired level of abstraction. For example, a single node representing the die temperature is unacceptable for thermal modeling at the die level. In addition, both static and transient thermal behavior should be modeled. Second, a compact thermal model should model just at the needed accuracy and hide the details of lower levels, so that the model itself is no more complex than necessary. Third, the model structure should be kept as simple as possible and should introduce little computational overhead. Fourth, it is also desirable to develop a parameterized compact thermal model for designers at all design levels. By doing this, explorations of the design space can be easily achieved without building physical prototypes or detailed thermal mod-

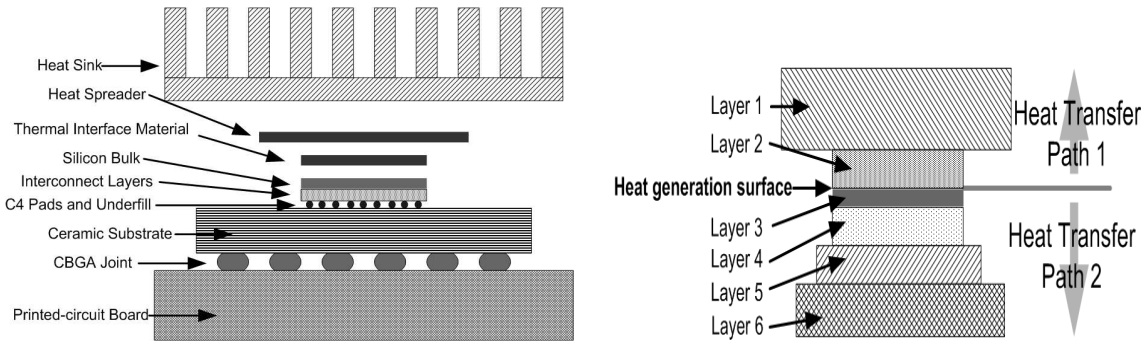

(a) (b)

Fig. 4.1: (a) The stacked layers in a typical CBGA package [63]. (b) The abstract stacked-layer structure used in HotSpot [34].

els, this will be shown in Chapter 5. Fifth, a compact thermal model has to be boundary condition independent (BCI). Therefore, the variation of the environment should not affect the internal structure of the compact thermal model. In Chapter 7, we will show that the HotSpot models are reasonably BCI [34, 35].

4.3 Primary and Secondary Heat Transfer Paths

Most modern VLSI systems have a package consisting of several stacked layers made of different materials, as shown in Fig 4.1(a). This is also the package scheme adopted for the HotSpot thermal models used in this dissertation. Typical layers are: heat sink, heat spreader, thermal interface material (thermal paste), silicon substrate, on-chip interconnect layers, C4 pads, ceramic packaging substrate, solder balls, etc. The recently proposed stacked chip-scale packaging (SCP) [1] and 3D IC designs [5] are also stacked-layer structures and can be easily modeled as extensions of the generic stack structure in Fig. 4.1(b).

When deriving a compact thermal model in HotSpot, the different layers, their positions

and adjacency are first identified. Each layer is then divided into a number of blocks. For example, in Fig. 4.2(c), the silicon substrate layer is divided according to architecture-level units or into regular grid cells, depending on what the die-level design requires. Note that only three blocks are shown in Fig. 4.2(c) for simplicity. Other layers that greatly affect across-die temperature distribution (e.g. thermal interface material) can be modeled similar to the silicon substrate. For the analysis of the needed size of regular grid cells, see Chapter 6.

For other layers that require less detailed thermal information (such as heat spreader and heat sink), we simply divide that layer as illustrated in Fig. 4.2(a). The center shaded part in a layer shown by Fig. 4.2(a) is the area covered by another adjacent layer such as the one shown in Fig. 4.2(c). This center part can have the same number of nodes as its smaller neighbor layer, or can collapse those nodes into fewer nodes, depending on the accuracy and computation speed requirements. The remaining peripheral part in Fig. 4.2(a) is then divided into four trapezoidal blocks, each assigned to one node.

Every block or grid cell in each layer has one vertical thermal resistance connected to next layer and several lateral resistances to its neighbors in the same layer. Fig. 4.2(b) shows a side view of one block with both the lateral and the vertical thermal resistances. The vertical thermal resistance is calculated by $R_{vertical} = t/(k \cdot A)$, where t is the thickness of that layer, k is the thermal conductivity of the material of that layer, and A is the cross-sectional area of the block. We see that each layer is not further divided into multiple thinner layers in the vertical direction, i.e., our modeling method is not fully 3-D. This is a reasonable approximation for early design stages since each layer is relatively thin (a millimeter or less), further discretization in the vertical direction would induce more

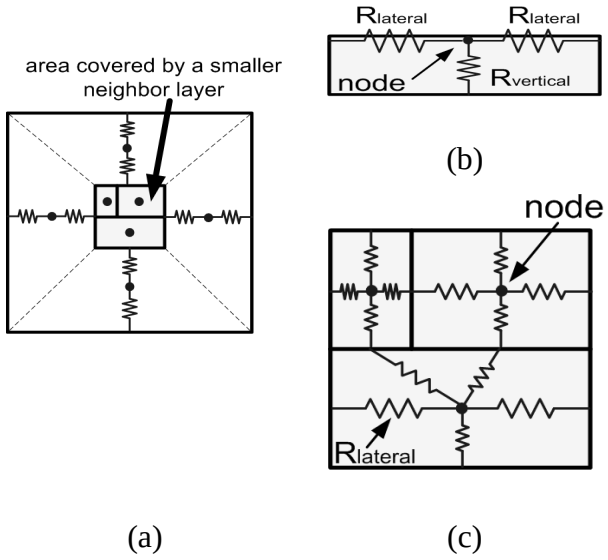

Fig. 4.2: (a) Partitioning of large area layers (top view). (b) one block with its lateral and vertical thermal resistances (side view). (c) a layer, for example, the silicon die, can be divided into arbitrary number of blocks if detail thermal information is needed (top view). [34, 35]

computation while not improving accuracy significantly.

Calculating lateral thermal resistance is not as straightforward as the vertical resistance. This is because heat spreading or constriction in the lateral directions must be accounted for. Basically, the lateral thermal resistance on one side of a block can be considered as the spreading/constriction thermal resistance of the neighboring part within a layer to that specific block. Lateral thermal resistances are normally much greater than their vertical counterparts due to the fact that the lateral heat-transfer cross-sectional areas are usually much less than vertical ones. To clarify how spreading/constriction resistances are computed, consider the two adjacent blocks, Block 1 and Block 2, in Figure 4.3. The lengths are L_1 and L_2 respectively. The chip thickness is t. Now our target is to calculate the lateral thermal resistance R_{21}, which is the thermal resistance from the center of Block 2 to the shared edge of Blocks 1 and 2. In this case, we can consider the heat is constricted

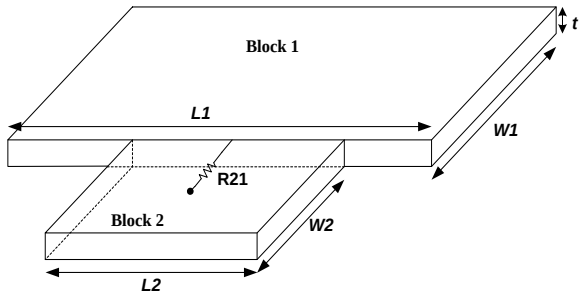

Fig. 4.3: Example to illustrate the calculation of spreading/constriction thermal resistance.

from Block 1 to Block 2 via the surface areas defined by $L_1 \cdot t$ and $L_2 \cdot t$. The constriction thermal resistance can be calculated by assuming the heat source area to be $L_1 \cdot t$, the silicon bulk area that accepts the heat is $L_2 \cdot t$, the thickness of the bulk is $W_2/2$. With these values found, we calculate the spreading/constriction resistance based on the formulas given in [50]. The resistance is a spreading one if the lateral area of the source is smaller than the bulk lateral area, and it is a constriction one otherwise.

For each node, there is also a thermal capacitance $C_{th} = \alpha \cdot c_p \cdot \rho \cdot t \cdot A$, connected to ground, where c_p and ρ are the specific heat and density of the material, respectively. The factor $\alpha \approx 0.5$ is a scaling factor accounting for lumped vs. distributed thermal RC time constants. [1]

Finally, the heatsink-to-air convection thermal resistance can be modeled as $R_{convection} = 1/(h \cdot A)$, where A is the convection surface area, and h is the heat transfer coefficient that is boundary condition dependent. For a first-order approximation, this is adequate for thermal analysis during early design stages. Typical values of h for typical heat sinks under different convection conditions usually be found in the heat sink datasheets.

[1] There should be no surprise that the same factor also appears in the analysis of distributed R-C electrical interconnect lines. This approximation is legitimate since the lateral thermal resistances are usually much greater and make negligible contribution to the thermal RC time constants compared to the vertical thermal resistances. For most accurate transient thermal simulations, multi-ladder R-C thermal circuits should be used instead.

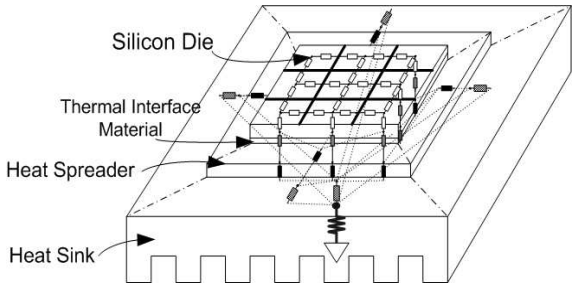

Fig. 4.4: Example compact thermal model with 3x3 grid cells for silicon die. For clarity of drawing, the structure is upside-down compared to the primary heat transfer path in Fig. 4.1(a). Thermal capacitors and heat sources are also omitted for clarity. [36]

Fig. 4.4 shows a model for the primary heat transfer path from the silicon through heat sink to the ambient air for the package in Fig. 4.1(a).

From the above descriptions, it can be seen that our method can model relatively detailed static and transient temperature variations for the silicon die. In particular, different packaging components can be modeled with more detailed temperature distribution information, which is not available in existing works such as [51, 84, 92]. Additionally, because the thermal models are built as lumped thermal R-C networks, the computational overhead for solving the temperatures is small. Therefore, the modeling method is suitable for developing compact thermal models used during design stages.

Compact thermal models developed from this method are also parameterizable and boundary condition independent (BCI). The models are parameterizable because they are built only using the physical geometries and material properties. The models are also BCI because the entire internal R-C network is built independent of boundary conditions. More discussions on parametrization and BCI of the HotSpot models can be found in Chapter 5 and Chapter 7 [34, 35].

4.4 Interconnect Self-Heating Power and Thermal Modeling

As mentioned before, there are two major heat transfer paths inside an IC package [63]—a primary heat transfer path (silicon substrate, heat spreader, heat sink) and a secondary heat transfer path (silicon substrate, on-chip interconnect layers, C4 pads, ceramic packaging substrate, solder balls, printed circuit board). We have already seen an example compact thermal model of the primary heat transfer path in Fig 4.4. On the other hand, the secondary heat transfer path usually removes a non-negligible amount of total generated heat (up to 30%). Neglecting the secondary heat transfer path can lead to inaccurate temperature predictions. In addition, as part of the secondary heat transfer path, the on-chip interconnect layers are of particular interest, because interconnect temperature information allows designers to perform more accurate electromigration, wire delay and IR drop analysis. Until now, a high-level interconnect self-heating model has been unavailable for early design stages. Most existing interconnect self-heating power and thermal models are either based on analysis of only a few wires [4] or need full-chip detailed layout information that is not available during early design stages [93]. We have presented preliminary results on this problem in [36]. Here we extend the discussion in [36] and present more details.

There are two aspects to be considered in the interconnect model: 1) *the average self-heating power* of interconnects in each metal layer, and 2) *the equivalent thermal resistance* for metal wires and their surrounding inter-layer dielectric. Vias also play an important role in heat transfer among different metal layers, and therefore need to be included as well.

4.4.1 Related Work

There have been a number of previous works on thermal modeling of on-chip interconnects and vias. For example, Chen et al. [13] present an interconnect thermal model that closely considers thermal coupling phenomenon between nearby interconnects. This model is accurate but it is on a per interconnect basis and is not extended to model multi-level structure at a higher design level. Chiang et al. [15] describe an analytical multi-level interconnect thermal model with considerations of via effects. This model copes with the thermal effect of vias by lumping the heat transferred through the vias into an equivalent thermal conductivity for the inter-layer dielectrics (ILD). However, to make interconnect thermal analysis complete, self-heating power also needs to be modeled. Unfortunately, to our best knowledge, we have not seen any previous works providing models on the multi-layer interconnect self-heating power at a higher design level. Our interconnect model is able to provide early estimations for the self-heating power for each metal layer. It also separately models the heat transfer through ILD and vias, which is a different modeling approach to [15]. In addition, our interconnect thermal model is constructed on top of the underlying silicon and package thermal models, therefore, the resulting interconnect temperature estimations have already taken the temperature distribution of the silicon into account.

Next, we first present the details of the interconnect self-heating power model, then the interconnect thermal model.

4.4.2 Interconnect Self-Heating Power Model

The self-heating power of a metal wire can be written as

$$P_{self} = I^2 \cdot R = I^2 \cdot \rho_m \cdot l / A_m \tag{4.1}$$

where I is the root-mean-square (RMS) current flowing through the wire, $R = \rho_m \cdot l / A_m$ is the electrical resistance, ρ_m is the metal resistivity (which is temperature dependent), l and A_m are the length and cross-sectional area of the individual wire. Because the model needs to *predict* wire temperatures before physical layout is available, first it has to be able to predict the average wire length and the self-heating current (RMS current) for wires in each metal layer. It is also important to notice that because the routing schemes are significantly different for the *signal interconnects* and the *power distribution network*, the methods of predicting average wire length and self-heating current are also different for signal and power supply wires, therefore, we treat them separately.

Average interconnect length in each metal layer for signal interconnects

We predict the average signal interconnect length in each metal layer by adopting and extending the statistical *a priori* wire-length distribution model presented by Davis et al. in [21], which improves the pioneer wire-length distribution model by Donath [23]. It is important to note that an interconnect thermal model at high levels of abstraction strongly depends on the *a priori* wire-length distribution model, and hence is limited by the accuracy and efficiency of the wire-length distribution model.

The model in [21] is based on the well-known Rent's Rule:

$$T = kN^p \tag{4.2}$$

where k and p are Rent's Rule parameters, N is the number of gates in a circuit, T is the predicted number of I/O terminals in the circuit. If the circuit block of interest is of a heterogeneous nature, i.e. there are different Rent's Rule parameters for different sub-circuit blocks, then equivalent Rent's Rule parameters can be found using the heterogeneous Rent's Rule proposed by Zarkesh-Ha et al. [94]:

$$k_{eq} = N_{G_{eq}} \cdot \sqrt{\Pi_{i=1}^{n} k_i^{N_{G_i}}} \quad \text{and} \quad p_{eq} = \frac{\Sigma_{i=1}^{n} p_i N_{G_i}}{N_{G_{eq}}} \tag{4.3}$$

where k_i and p_i are the homogeneous Rent's Rule parameters for each sub-circuit block, N_{G_i} is the number of gate in each sub-circuit block and $N_{G_{eq}} = \Sigma_{i=1}^{n} N_{G_i}$.

Table 4.1, which is extracted from [96], shows typical values of N, k and p for a RISC microprocessor [88]. For a given microprocessor family, the Rent's Rule parameters of each circuit block (a mega-cell or macro) tend to remain the same over generations due to the recursive application of Rent's Rule throughout the entire monolithic circuit block [21]. Therefore, we can assume that the same parameters, k and p, can be used in a design at future technology nodes for the same microprocessor family, although the number of gates N increases. Sometimes, wire-length distribution is needed at a higher abstraction level, e.g. at die level, thus equivalent Rent's Rule parameters need to be calculated at that level. Using the method in [95], one can calculate the equivalent Rent's Rule parameters for the whole processor and the core of the processor (excluding the on-chip cache memories), as shown in the last row of Table 4.1.

Note that the choice of proper Rent's Rule parameters significantly affects the accuracy of the wire-length distribution model, hence the interconnect self-heating power model presented here. Recently, limitations of the conventional Rent's Rule has been investigated by Christie et al [16] and improvements on conventional Rent's Rule have been proposed in

Megacell's Name	k_r	N	p_r	Megacell's Name	k_r	N	p_r
I$	4.12	380000	0.20	Instr. Fetch Addr.	3.20	16500	0.60
I$ Tags	3.80	18000	0.47	Instr. Fetch Datapath	3.20	13800	0.60
D$	4.12	350000	0.20	Instr. Fetch Ctrl	3.20	9500	0.60
D$ Tags	3.80	25500	0.47	Addr. Queue	3.20	22000	0.60
TLB	3.80	22400	0.35	Instr. Decode	3.20	45300	0.60
L2 Cache Ctrl.	3.20	15700	0.60	Integer Datapath	3.20	43800	0.60
Exter. I/F	3.20	18400	0.60	Integer Queue	3.20	19700	0.60
System I/F Buf.	3.20	22600	0.60	FP Datapath	3.20	32600	0.60
Free List	3.20	9800	0.60	FP Queue	3.20	51000	0.60
Graduation Unit	3.20	26300	0.60	FP Multiplier	3.20	19300	0.60
Die-level Equiv.	3.79	1162200	0.34	Logic Core Equiv.	3.23	388700	0.58

Table 4.1: Rent's Rule parameters for a RISC microprocessor [88], data are extracted from [96]. The last row shows the equivalent Rent's Rule parameters of the whole die and the logic core (excluding the cache memories).

works such as [11, 16, 20]. More accurate wire-length distribution model can be obtained with these improvements on Rent's Rule.

Three wire-length regions are considered in [21]—local, semi-global and global. The model predicts the number of wires of any specific length, which is called the interconnect density function $i(l)$, where l is the wire length in gate pitches. Fig. 4.5 shows an example wire-length distribution based on ITRS data [2] for high-performance designs at the 45nm technology node, where L_{loc}, L_{semi}, L_{glob} are maximum local, semi-global and global wire lengths, respectively.

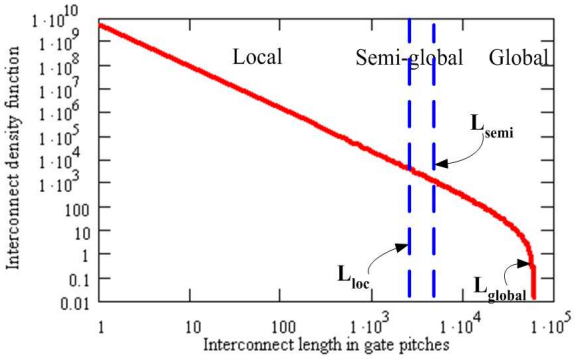

Fig. 4.5: An example of wire-length distribution at 45nm technology node, with three regions (local, semi-global and global) [36].

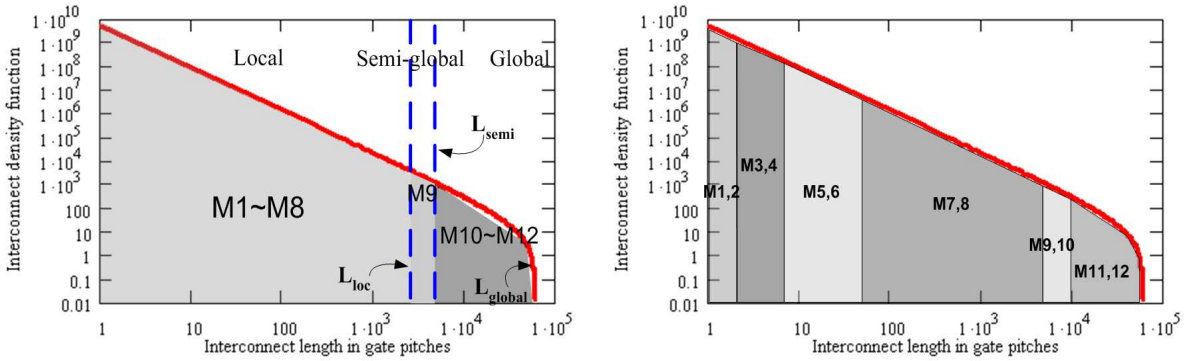

(a) (b)

Fig. 4.6: (a)—Metal layer assignment by calculating number of metal layers needed for each of the three regions. (b)—Metal layer assignment by filling every two metal layers with signal wires, starting from Metal 1 and Metal 2. (b) is superior to (a) by providing more detailed metal layer assignment information [36].

Using the interconnect density function $i(l)$, one can calculate the average length and number of wiring nets for each region. For example, for the semi-global region:

$$l_{semi} = \chi \text{ f.o.} \frac{\int_{L_{loc}}^{L_{semi}} i(l) \cdot l \, dl}{\int_{L_{loc}}^{L_{semi}} i(l) \, dl} \tag{4.4}$$

$$n_{semi} = \frac{1}{\text{f.o.}} \int_{L_{loc}}^{L_{semi}} i(l) \, dl \tag{4.5}$$

where χ is the correction factor that converts the point-to-point interconnect length to

wiring net length (using a linear net model $\chi = 4/(\text{f.o.} + 3)$), and f.o. is the average number of fan-outs per wiring net. More details can be found in [21].

However, there is no wire-length distribution information regarding each metal layer when using this three-region division method in [21]. For the interconnect compact thermal model, we need the wire-length distribution predictions of every metal layer. Because of the predominant usage of Manhattan routing, in general two metal layers are needed to route one wiring net—one layer for horizontal routing, the other for vertical routing. In this dissertation, we estimate the pair of metal layers where each wiring net is routed by filling every two metal layers with wiring nets, starting from the shortest wiring nets. We thus assume that the shortest wiring nets of the wire-length distribution in Fig. 4.5 are assigned to Metal 1 and 2. Once the first two metal layers are filled, we proceed to Metal 3 and Metal 4, so on and so forth, until all the wiring nets are assigned to their corresponding pair of metal layers. Although this is a simplification we expect it to be representative of an actual routing strategy. A useful byproduct of our approach is that we are also able to estimate the total number of metal layers needed for a design. As illustrated in Fig. 4.7, assume the length of the shortest and longest point-to-point interconnects that can be assigned to a pair of metal layers are L_{min} and L_{max} in gate pitches; we can then find the average length and total number of wiring nets within a pair of metal layers by

$$l_{avg} = \chi \,\text{f.o.} \frac{\int_{L_{min}}^{L_{max}} i(l) \cdot l \, dl}{\int_{L_{min}}^{L_{max}} i(l) \, dl} \tag{4.6}$$

$$n_{total} = \frac{1}{\text{f.o.}} \int_{L_{min}}^{L_{max}} i(l) \, dl \tag{4.7}$$

Furthermore, by assuming the routing structure of Fig. 4.7, where M is the number of signal wires between two power rails and Sp is ratio of the space between every two signal

wires to l_{avg} (both M and Sp are design parameters and are tunable by the designer), we get the following relation:

$$n_{total} \cdot (Sp + 1) \cdot l_{avg} \cdot \frac{M + 1}{M} \cdot p = 2 \cdot Area \tag{4.8}$$

where p is the wire pitch of a metal layer, and $2 \cdot Area$ is the available routing area for the pair of metal layers under consideration. Using this relationship, and starting at Metal 1 and Metal 2 with $L_{min} = 1$, we are able to solve for L_{max} and L_{min} for each pair of metal layers. An example metal layer assignment for the interconnect distribution of Fig. 4.5 is shown in Fig. 4.6(b).

Another way to assign signal wiring nets to different layers is to calculate the number of metal layers needed for each of the three regions, namely, local, semi-global and global region as in [21]. The resulting metal layer assignment is shown in Fig. 4.6(a). As can be seen, the results in Fig. 4.6(a) and (b) are similar, but Fig. 4.6(b) provides detailed metal layer assignment estimations for every two metal layers without considering the three regions, while the information provided in Fig. 4.6(a) is coarser. Therefore, we prefer the approach used in Fig. 4.6(b). On the other hand, if the total number of metal layers is fixed, the parameters Sp and M can be adjusted accordingly to fit all the signal interconnects into the metal layers.

Average interconnect length in each metal layer for power and ground

So far, we have considered the average *signal* interconnect length in each metal layer. We also need to find the average wire length for the *power* and *ground* networks, which are usually grid-like. This is relatively simple: we only need to find the length of the power grid section in each metal layer. The assumption here is that the power grid for each metal

layer is uniformly distributed, which is a reasonable assumption for early high-level design stages.

With this we are done with estimating wire length, next we need to use this information to estimate interconnect self-heating power.

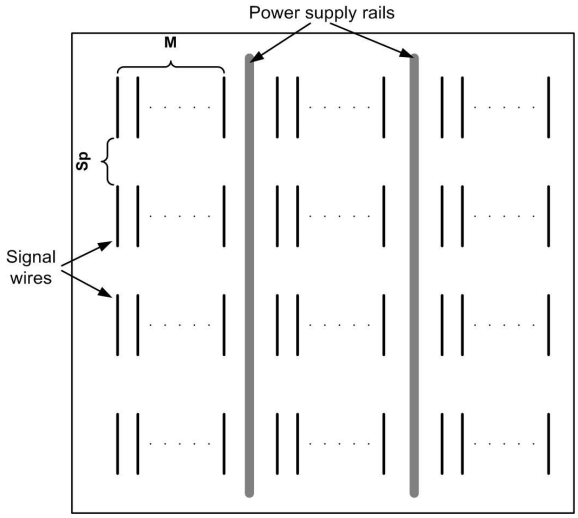

Fig. 4.7: A scheme to assign signal interconnects to metal layers. M is the number of signal wires between two power rails and Sp is ratio of the space between every two signal wires to average signal wire length of that metal layer [31].

Average interconnect RMS self-heating current in each metal layer for signal interconnects

For each switching event, half of the energy drawn from the power supply is dissipated in the form of heat on the charging/discharging transistor and on the output signal interconnect. The average current flow through the interconnect during a switching event can be solved from the following equation:

$$I_{RMS}^2(R_{tr} + R_{wire})t_d = \frac{1}{2}\alpha\, C_L V_{dd}^2 \tag{4.9}$$

where I_{RMS} is the self-heating current per wire in each metal layer. R_{tr} is the on-resistance of the transistor, R_{wire} is the wire resistance, α is the switching activity factor, C_L is the load capacitance, and t_d is the delay of the switching event. For long interconnects, repeaters are inserted in order to achieve optimum delay, and these need to be also taken into account. The critical wire-length between repeaters (L_{crit}), the delay for one section of buffered interconnect (τ_{crit}), the optimal number of repeaters (Nr_{crit}) and the optimal size of repeaters (s_{crit}) for interconnects in each region can be found using the repeater insertion model proposed in [61]. The calculations of R_{tr}, R_{wire}, C_L and t_d are different for wires with or without inserted repeaters—the wire length is either the total wiring net length or the length of a wire section between repeaters; the driving and load gates are either gates with average transistor size or repeaters with size of s_{crit}. Finally, the delay of the switching event, t_d, can be approximated as τ_{crit} for interconnects with repeaters, or as $clock\_cycle\_time/logic\_depth$ for interconnects without repeaters.

Average interconnect RMS self-heating current in each metal layer for power and ground

To calculate average RMS currents for power supply grid sections we can use one of two methods.

The first method is to build a grid-like resistive network model for V_{dd} and GND, somewhat resembling the grid-like die-level thermal model as in Fig. 4.4. Each resistor connecting two nodes in the same metal layer is now the electrical resistance of one power supply grid section. Resistors connecting power grid nodes of different metal layers represent the vias. The topology of the network is obtained by knowing the pitch between power

rails in each metal layer, average length and number of power grid sections between power grid. Next, by applying voltage source to the top-layer nodes where the C4 pads resides and current sources at the M1-M2 layer nodes where power are drawn by transistors, the resistive network is solved to find the average self-heating current of the power grid in each metal layer. The other method to calculate average RMS self-heating current of power grid section in a metal layer is quite straightforward—we can simply divide the total current delivered to a metal layer by the number of power grid sections. This method is suitable for high-level design stages, but is not as accurate as the first method.

Total interconnect self-heating power in each metal layer

With all the above information of average interconnect length and RMS self-heating current in each layer (for both signal interconnects and power grid sections), we calculate the average self-heating power per interconnect in each metal layer:

$$P_{wire} = I_{RMS}^2 \cdot R_{wire} = I_{RMS}^2 \cdot \rho_m \frac{l_{wire}}{A_{wire}} \tag{4.10}$$

where A_{wire} and l_{wire} are the cross-sectional area and the average length of signal interconnects or power grid sections in each metal layer, respectively.

Finally, we calculate the self-heating power for each metal layer. For example, we calculate the self-heating power of the ith metal layer as:

$$P_{self\_i} = P_{wire\_sig\_i} \cdot n_{sig\_i} + P_{wire\_pwr\_i} \cdot n_{pwr\_i} \tag{4.11}$$

where $P_{wire\_sig\_i}$ and $P_{self\_pwr\_i}$ are the self-heating power of each individual signal interconnect and power supply wire for the ith metal layer, respectively. n_{sig_i} and n_{pwr_i} are number of signal interconnects and power supply sections in the ith metal layer.

So far, we are done with the first aspect of interconnect thermal modeling—self-heating power calculation of metal layers. Next, we need to calculate the equivalent thermal resistance of wires and the surrounding dielectric, together with the thermal resistance of vias.

4.4.3 Equivalent Thermal Resistance of Wires/Dielectric and Vias

In order to derive a model we consider the case in Fig. 4.8, where two wires (Wire1 and Wire2) are adjacent to each other. On top of and beneath them are orthogonal wires in neighboring metal layers. All wires are surrounded by inter-layer dielectrics (ILD). Remember that we want to find the equivalent thermal resistance (R_0) from Wire1 to $d/2$ above Wire1, where d is the thickness of the inter-layer dielectric between two metal layers. The other half of d belongs to the metal layer above Wire1, and is considered when calculating equivalent thermal resistance for wires in that layer. Since we have assumed all the signal wires (or power supply wires) in the same metal layer are the same, Wire1 and Wire2 are two identical wires dissipating the same power at the same time. Wire1 and Wire2 also have the same temperature. We approximate the isothermal surface by the outer dashed area in Fig. 4.8. This isothermal surface is used for the calculation of R_0 and is $d/2$ away from the wires. Also, it does not overlap with similar isothermal surfaces for the perpendicular wires in neighboring layers. The effective heat conducting angle used for the calculation of R_0 can be approximated by $\theta = 2 \cdot tan^{-1}(D/(d+H))$, as shown in the figure.

There are also a lateral thermal resistance between Wire1 and Wire2—R_{lat}. However, because Wire1 and Wire2 are identical and have the same temperature, there is no heat transfer in the lateral direction and R_{lat} can be removed.

For the calculation of R_0, we first calculate the thermal resistance of the dark slice of

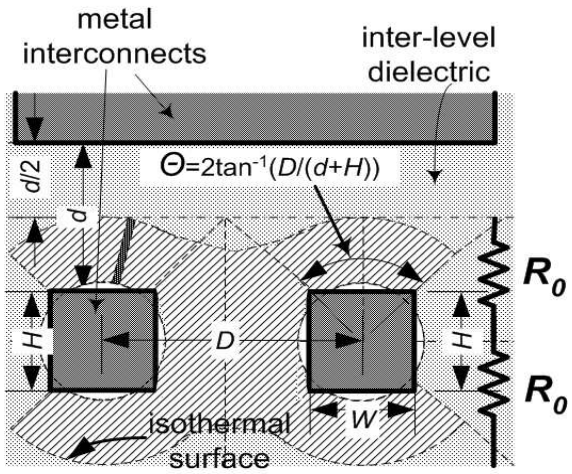

Fig. 4.8: Interconnect structures for calculating equivalent thermal resistance of wires with surrounding dielectric [36].

ILD shown in Fig. 4.8, which can be written in the form of integral

$$dR_0 = \int_0^{d/2} \frac{1}{k_{ins}} \frac{dx}{(r+x)d\phi \cdot l} = \frac{1}{k_{ins} \cdot l \cdot d\phi} \ln(\frac{d+2r}{2r}) \tag{4.12}$$

where x is the integral variable, k_{ins} is the thermal conductivity of ILD, ϕ is the angle of the slice, $r = \sqrt{WH/\pi}$ is the equivalent radius of the wire, and l is the length of the wire.

If we define thermal conductance G_0 as the reciprocal of thermal resistance R_0, we have

$$dG_0 = \frac{k_{ins} \cdot l \cdot d\phi}{\ln(\frac{d+2r}{2r})} \tag{4.13}$$

Therefore, we have

$$G_0 = \int_0^{\theta} dG_0 = \frac{\theta \cdot k_{ins} \cdot l}{\ln(\frac{d+2r}{2r})} \tag{4.14}$$

so the total equivalent thermal resistance is

$$R_0 = \frac{1}{G_0} = \ln(\frac{d+2r}{2r})/(\theta \cdot k_{ins} \cdot l). \tag{4.15}$$

Inter-layer heat transfer also happens through vias. A simplistic approximation of the number of vias for signal interconnect is to assume that each wiring net has two vias, one

connected to the upper metal layer, and another one connected to the lower metal layer. A more accurate approximation is to assume that each wiring net has $(2\cdot\text{f.o.} + 2)$ vias, where f.o. is the average fan out number of each gate. As illustrated in Fig. 4.9, $(\text{f.o.} + 2)$ vias are at the ends of the wiring net and connecting the wiring net to lower metal layers and eventually to the device layer at the silicon surface. The other vias are used to aid the routing of the wiring net between the pair of metal layers in which the wiring net resides. For the power supply grid, in order to increase the reliability and because the wires are typically wider than minimum size, designers usually use multiple vias at the intersection of two power rails between different metal layers. As illustrated in Fig. 4.10, the number of vias at an intersection of power rails can be estimated by

$$\frac{1}{4}\left(\frac{W_{wire}}{W_{via}} - 1\right)^2 \tag{4.16}$$

where W_{wire} and W_{via} are the widths of the power wire and the via, respectively. The thermal resistance of each via is approximately calculated as $R_{via} = t_v/(k_v A_v)$, where k_v is thermal conductivity of via-filling material, t_v and A_v are the thickness and cross-sectional area of the via.

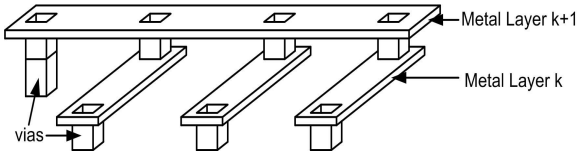

Fig. 4.9: Estimating the number of vias for signal interconnects — A wiring net with fan out 3 is shown in this figure. The number of vias is $(2\cdot\text{f.o.} + 2)$ [31].

All thermal resistors of wires and vias between two metal layers can be considered parallel to each other. Thus, combining all the thermal resistors between two metal layers, we obtain the total equivalent thermal resistance between two metal layers.

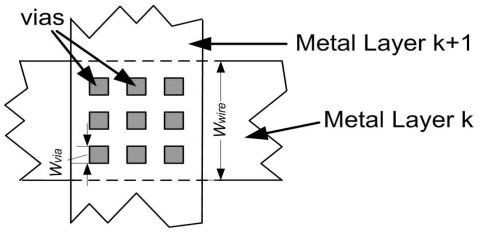

Fig. 4.10: Estimating the number of vias for power supply wires — An array of vias are put in the intersection of power wires at two metal layers. W_{wire} and W_{via} are the widths of the power wire and the via, respectively [31].

Now we are almost done with the interconnect thermal modeling. One last step is to stack the thermal resistances for each layer to construct the whole thermal circuit for all interconnect layers. Thermal capacitances can also be calculated for each metal layer and the inter-layer dielectric based on dimensions and material properties using an equation similar to the one in Section 4.3.

4.4.4 Interconnect Power and Thermal Model Granularity

Although the above interconnect thermal modeling approach was presented at the entire die level, in principle, it is also applicable at other granularities. For example, Rent's Rule can also be applied at the functional unit level to estimate intra- and inter-functional-unit wire-length distribution for metal layers above each functional unit. The total self-heating power and the equivalent metal layer thermal resistance can then be calculated for each functional unit using similar methods as described before. As part of our future work, the power and temperature estimations for each metal layer at the functional unit level or other abstraction levels will be investigated.

4.4.5 Accuracy Concerns about the Interconnect Model

From the above descriptions of the proposed interconnect power and thermal model for early design stages, one might raise concerns such as (1) Why would one need such a interconnect thermal and power model in early design stages? (2) Why would the estimations from the model be meaningful given the empirical nature of Rent's Rule? (3) What about the accuracy of the current loading calculation method in the model? The answers to these concerns are listed as follows:

1. Because interconnect layers usually have much higher absolute temperatures and greater temperature differences than silicon, reliability issues such as thermo-mechanical stress between metal layers, thermo-electromigration of long wires, are increasingly more important. This has been the major reason that we need to estimate wire temperature during early-stage design. If reasonably accurate early-stage wire temperature estimations are available, they would be helpful for the designers to discover and deal with such thermally-related reliability hazards early in the design flow, hence greatly expediting the design convergence process.

   To illustrate, consider the micro-architecture level design stage as an example. Architectural choices, such as number of cores, core-to-core interconnect organization, or within a core whether to multi-port a structure or partition it, are major determinants of later system properties, today also including thermal and reliability properties. This means that the architect needs a way to reason about these properties at the pre-RTL architecture determination stage. These kinds of choices don't necessarily need high degrees of precision, it is enough only to know what combination of choices might be problematic. Then alternative organizations can be explored if necessary: if

a much safer choice is found that exhibits tolerable overhead, this may be preferable to a detailed implementation and the *risk* of re-design or thermal mitigation—even if the original design (due to the margin of error in the high-level modeling) would have met all the requirement.

2. About the accuracy concern of Rent's Rule: It is true that arbitrarily chosen chip-level Rent's Rule parameters would greatly affect the accuracy of the presented model. However, for a mature circuit design style of a specific functional unit along a microprocessor family, Rent's Rule parameters derived from ancestor designs can be used to predict future designs' wire-length distributions with good accuracy, as indicated by Rent's Rule validation data presented in previous works on Rent's Rule such as [11, 16, 20, 21, 94]. Our understanding of the Rent's Rule accuracy is that Rent's rule is indeed inaccurate for any *individual* wires, but can be accurate about aggregate average wire behavior for mature circuit design styles. This is also true about other applications of Rent's Rule. Therefore, it is reasonable that we only use Rent's Rule to determine average aggregate interconnect temperatures. Here, it is important to notice that Rent's Rule estimations are not used to perform detailed design rule check, for which accurate individual wire information is required, but just for aggregate temperature estimations to be used in early-stage design decisions. Also, for pre-RTL architecture studies, interconnect thermal estimations only need to be accurate to within a similar coarse precision as modeling of other metrics, such as microarchitecture power modeling from Wattch [10].

3. About the concern of current loading accuracy: We think that reasonably accurate average/RMS current estimations for typical signal wires (e.g. at the functional block

granularity for micro-architecture level designs) are achievable as presented earlier in this section. This is because power estimations at this level (dynamic power with switching factors and static power) are available from tools such as Wattch. Average current loading in the power/gound network can also be roughly estimated by solving a coarse Vdd/GND mesh (similar to the regular-grid-cell thermal resistive network in HotSpot) without loss of much accuracy. Of course, by modeling aggregate average wire-length and average/RMS current, important worst cases would not be available, but this is acceptable for early-design-stage analysis and design explorations since other inputs to the model such as power estimations are not as detailed, either. It is also obvious that this kind of approach is consistent with the needs of pre-RTL architectural modeling.

In summary, the early-stage interconnect self-heating and thermal model we propose could save design efforts significantly by discovering potential thermal-related interconnect reliability violations in early design stages. The modeling approach is reasonably accurate given the acceptable accuracy of using (a) validated and more mature Rent's Rule characteristics, (b) available widely-adopted high-level power estimation tools, and (c) simplified but early-stage-suitable current loading model.

4.5 Computation Speed of HotSpot Models

The computation speed of HotSpot thermal models to obtain steady-state and transient solutions for several different simulated time intervals at different granularities are in the order of mili-seconds to minutes, depending upon the number of blocks/grids, number of

material layers, and the simulated transient time interval. Table 4.2 shows the CPU time used to simulate a HotSpot model with 40×40 grid cells.

simulated time interval	CPU time
0.1 ms	20 ms
20 ms	100 ms
2 sec	7.9 sec
20 sec	78.5 sec
steady-state	14.9 sec

Table 4.2: Computation speed of a HotSpot model, running on a dual-processor (AMD MP 1.5GHz) system. (Converging method for transient solutions is different from that for steady-state solutions.)

The small overhead is due to the relatively small and manageable number of nodes in the lumped thermal R-C circuit, together with the use of first-order difference equations to iteratively solve the R-C network. The computational efficiency of HotSpot models means there is little computation overhead for existing design methodologies to incorporate the compact thermal models for temperature-aware design or dynamic thermal management simulations.

4.6 Discussions about HotSpot Modeling Method

4.6.1 Functional Units vs. Regular Grids

As mentioned before, the silicon die, the interface material and the center part of the heat spreader can either be divided naturally according to functional units or be divided

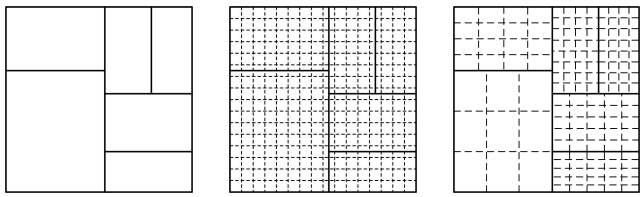

(a) (b) (c)

Fig. 4.11: Modeling at the granularity of (a) functional blocks (b) uniform grid cells (c) hybrid-sized grid cells [28].

into regular grid cells, depending on the needs of the designer (Fig. 4.11(a) and (b)). For example, a computer architect may only need to estimate average temperatures for each functional unit, a thermal model at the functional unit granularity being best in that case. However, for a package designer, the temperature gradients across silicon die and other package layers are important metrics to evaluate the reliability of the package. In this case, a grid-like thermal model is more suitable, since it provides more detailed estimations of maximum and minimum temperatures, whereas the functional-unit-level model does not. Our method offers a choice between a possibly irregular functional-unit partitioning and a regular grid of variable granularity. Ideally, we would prefer a hybrid grid scheme that combines both the per-function unit model and the uniform-size grid model as in Fig. 4.11(c) [28]. By doing this, we could still get detailed thermal information for particular blocks under consideration while saving computation effort by introducing fewer nodes inside other blocks. It is clear that the desired accuracy determines the minimum grid cell size needed, i.e. the temperature difference across one grid cell should be less than a certain percentage of the maximum temperature difference across the die. In what follows we show an analytical method to derive the proper size of a grid cell.

When modeling the die with regular grid cells, our method is essentially a combination

of simplified finite difference method (for the silicon die and the interface material) and finite element method (for other package components). On the other hand, if the die is modeled at the architectural functional unit level, our modeling method can be considered as a very simple form of FEM approach. Our modeling method achieves low computational overhead as well as reasonable accuracy by taking advantage of the established architecture-level and package-level information during the simplified finite element and finite difference processes.

4.6.2 Differences to Existing Compact Thermal Models

From the above it can be seen that our modeling approach is different from other existing compact thermal modeling methods:

1. Existing compact thermal models are mainly at the package or system level, hence with only a few nodes for the silicon. This is adequate for the package vendors and system-level designers, but not for circuit designers or architects. Our compact thermal models uses more nodes for the die structures, which is convenient for the silicon-level designers.

2. Existing compact thermal models hide the packaging details due to the requirements of package vendors, while our models need somewhat detailed package-level and silicon-level information.

3. The thermal resistances in the existing compact thermal models are extracted from detailed model simulations or real package temperature measurements, while the thermal resistances of our models are calculated based on dimensions and physical

properties of materials.

All these differences lead to different applications for existing compact thermal models and our compact thermal models. Existing models are good for thermal analysis and characterization of existing package designs without revealing details of the package, while our models are better for explorations of new silicon-level and package-level designs. Also, our models are intrinsically parameterized and reasonably BCI, as discussed in [34, 35].

4.6.3 Advantages and Limitations

Advantages

So far, the modeling method and major characteristics of our compact thermal modeling approach have been presented. Despite a few limitations, our modeling approach has several significant advantages; the advantages are mainly due to the fact that it is parameterized and BCI.

1. Parametrization is useful because a variety of design explorations can be carried out by only changing the dimensions and material parameters without reconstructing the whole compact thermal model through detailed simulations. For example, using our models, one can easily find the optimum die thickness by simply sweeping the die thickness parameter and keeping all the other parameters constant. Another example would be investigating the effect of different types of heat spreaders or heat sinks. One can easily add/change the layers of heat spreader or heat sink by following our presented modeling method earlier in this chapter.

2. It is also important to notice that our modeling method can be used to study hypothet-

ical systems for which physical implementations and thermal measurements cannot yet be obtained. One example is the investigation of emerging 3-D integrated circuits. Prototyping 3-D ICs would introduce prohibitive cost due to the drastic change of existing fabrication process. But with our models, designers can easily model heat transfer and temperature rise in 3-D IC structures, thus evaluate the feasibility of this new design paradigm from a thermal point of view.

3. Because our modeling approach can be validated to be reasonably BCI (see Section 7.5), designers can focus more on their design efforts without worrying about the thermal model's validity under different boundary conditions.

4. There have been several successful applications of our modeling approach in different design areas (see Chapter 8. For example, it has been used to build compact thermal models in research areas such as dynamic thermal management (DTM) techniques for microprocessors [74] and die-level thermal-aware computer-aided designs [36]. Similarly, we expect the presented modeling method will stimulate more research collaborations among package designers, VLSI circuit designers and computer architects.

Limitations

Our modeling approach also has a few limitations compared to existing package-level compact thermal models that are extracted from detailed thermal simulations or direct measurements, such as [47, 68].

1. It is not as "compact" as other existing package-level compact thermal models, but the number of nodes are still within a manageable amount, and the computational

overhead is also negligible compared to detailed numerical models.

2. When it comes to analyze or release a fixed compact thermal model for an existing design or a final product, our model's user interface is not as concise as some existing package-level compact thermal models. This is due to the complexity of our model and the revealing of package design details.

3. At the same level of complexity, our model is not as accurate as other existing package-level compact models. This is because those existing models are extracted from detailed model simulations or real package temperature measurements, which are still the most accurate ways to model thermal effects, while our model is essentially a simplified version of the detailed model, therefore can not achieve the same level of accuracy as the detailed simulations, or the derived DELPHI-like models. In addition, some lumped thermal resistances (e.g. the ones in the peripheral parts of heat spreader and heat sink) do not fully account for all the possible heat transfer paths, thus not really representing the exact thermal resistance according to the analysis in [42, 68].

4. Our modeling approach is not as BCI as the DELPHI and other existing modeling approaches. This is because the surface area division method used by our model is not exactly the optimal one [44], although it is proved to be a reasonable one as shown in Section 7.5 and [8]. It is also partly due to the fact that heat spreading in the package cannot be as well considered in our modeling approach as in the detailed models.

5. Our modeling approach needs to be improved to account for more different types of

packages. For example, if one layer of the package contains more than one material or the package itself is a multi-chip module, our current thermal models are not flexible enough to cover these cases. This is just a limitation of our current implementation since our general modeling method can be easily extended to account for such more complicated structures.

Chapter 5

Parametrization

A fully parameterized compact thermal model (CTM) allows package designers, circuit designers and computer architects to explore new design alternatives and evaluate different thermally-related design trade-offs at their corresponding design levels before the actual physical designs are available. More importantly, with the aid of the parameterized compact thermal models, designers at different design levels can have more productive interactions and collaborations at early design stages of a microelectronic system. This leads to early discovery and considerations of potential thermal hazards of the system. True parametrization requires that the models be constructed based solely on design geometries and material properties. In this chapter, we show that HotSpot is parameterized. Several examples demonstrate the usefulness and flexibility of a parameterized compact thermal model.

5.1 Importance of Parametrization and Related Work

Parametrization of compact thermal models is desirable and has drawn attention from researchers. In [44, 91], the authors point out that achieving a sensible parametrization of compact thermal models is next to impossible for the chosen simple structure of the some of the existing models, such as the DELPHI ones [47, 48, 68]. This is because the DELPHI model structure consists of only a few thermal resistances which makes it impossible to parameterize the actual very complex package structure, together with the variations of thermal conductivities and the heat spreading/constriction effects within the die and the package. On the other hand, our modeling approach can be better parameterized due to its physically-based nature. The cost for parametrization is that our models are usually more complex than, and not as accurate as, the DELPHI-like and other existing models.

The importance of achieving full parametrization of compact thermal models is obvious. Fully parameterized compact thermal models allow designers at all design levels to freely explore all the possible thermally-related design spaces. For example, the heat spreader and heat sink are two important package components for high-performance VLSI designs. While a large heat spreader and heat sink made from high thermal conductivity materials can reduce the temperature of silicon die, they also significantly increase the total price of the system. Therefore, exploring design trade-offs between hot spot temperatures and package cost is crucial. With parameterized compact thermal models, this exploration can be done easily and efficiently by simply sweeping the size of the heat spreader and heat sink with different material properties to achieve the desired package design point. On the other hand, building package prototypes or detailed thermal models greatly slows the design process and increases the design cost.

In addition, fully parameterized compact thermal models can also provide a more productive communication channel among designers at different design levels. Therefore, potential thermal hazards in the design can be discovered and dealt with in early design stages. For example, a typical design scenario would be: circuit designers come up with estimations of power consumptions of each circuit block, computer architects come up with a floorplan, while package designers come up with a proposed package. Using a parameterized compact thermal model, these designers can together easily evaluate the combined system design from a thermal point of view. If it appears that there are unacceptable hot spots on silicon, or the temperature difference across the package is too high and degrades the reliability of the design, different design decisions can be made in this early design stage—circuit designers may need to develop novel circuits with lower power consumption, computer architects may apply different dynamic thermal management techniques and re-arrange the floorplan for better across-silicon temperature distribution, or the package designer can improve the proposed package design by using a more advanced heat spreader or heat sink, etc.

In fact, in a previous work [25], the authors have proposed compact thermal models which include the parametrization of certain package parameters such as the PCB size, PCB thermal properties, heat sink dimensions and thermal vias. However, full parametrization was not achieved in [25] because the die itself is not parameterized. On the other hand, our modeling approach can be considered as fully parameterized, including the silicon die itself. The parametrization of the die is a crucial requirement to explore different preliminary die-level designs before the prototypes are available.

5.2 Some Examples

For illustration, we first present an example analysis to show the strength of using parameterized compact thermal models to efficiently investigate the impact of thermal interface material on across-silicon temperature differences. Fig. 5.1 shows the relationship between the thickness of the thermal interface material (TIM), which glues the silicon die to the heat spreader. We plot the temperature readings from a compact thermal model similar to Fig. 4.4 with 40×40 grid cells on silicon. This analysis is based on an Alpha 21364-like microprocessor floorplan. Average silicon die temperature is also plotted in Fig. 8.4 for reference. The total heat generated from the silicon surface is 40.2W, the die size is 15.9mm×15.9mm×0.5mm, and the thermal conductivity of the thermal interface material is 1.33W/(m-K).

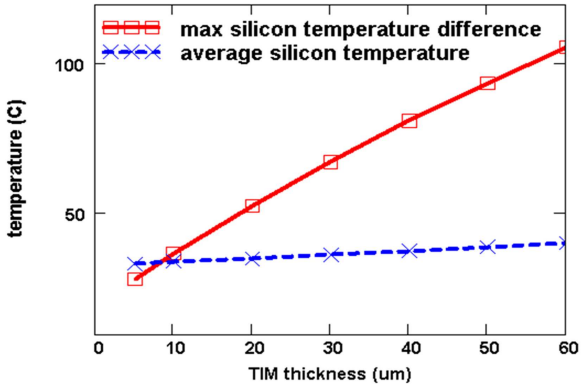

Fig. 5.1: The impact of thermal interface material (TIM) thickness to silicon die temperature difference. Average silicon die temperature is also plotted as reference [35].

As can be observed from Fig. 5.1, although the TIM thickness doesn't have obvious impact on the average die temperature, thicker TIM results in poor heat spreading which leads to large temperature differences across the die. Such large temperature differences may be disastrous to circuit performance and die/package reliability. Using a better heat

sink will only lower the average silicon temperature but will not help to reduce the temperature difference. From this analysis, which has been easily performed by our parameterized model, we can reach the conclusion that using as thin as possible thermal interface material is one of the key issues for package designers to consider. In some recent work [47, 45], the importance of measuring and modeling thermal interface's impacts on the entire package has been also discussed, although not at the same silicon die level as the above example shows.

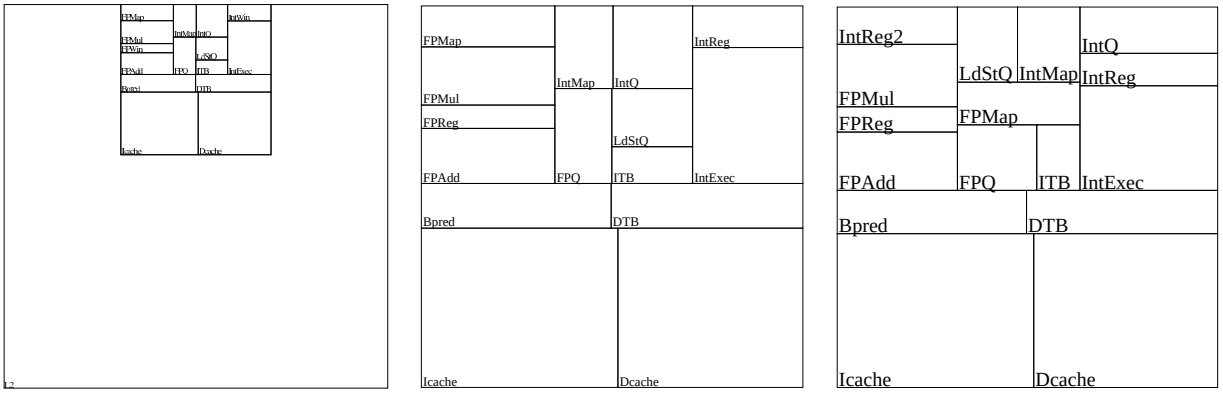

(a) (b) (c)

Fig. 5.2: (a) Approximated floorplan of Alpha 21364 microprocessor. (b) Close-up look of the 21364 core, with only one register file on the top-right corner. (c) 21364 core with an extra copy of register file at the top-left corner [74].

Another example of utilizing our parameterized compact thermal model is the investigation of a dynamic thermal management technique (DTM) called migrating computation (MC), at the micro-architecture level [74]. It is obvious that two silicon-level functional units that run hot by themselves will tend to run even hotter when adjacent. On the other hand, separating them will introduce additional communication latency that is incurred regardless of operating temperature. This suggests the use of spare units located in cold areas of the chip, to which computation can migrate only when the primary units overheat.

In [74], the floorplan of the Alpha 21364 core is carefully changed to include an extra copy of integer register file (see Fig. 5.2), which is usually the hottest spot on the silicon die for this design. We also model the fact that accessing the secondary integer register file entails extra power and extra access time due to the longer distance. With this new floorplan, we can shift the workload back and forth between the two register files when the one in use overheats, with a little performance overhead (11.2% slower). The changes in silicon floorplan can be easily adapted into corresponding parameterized compact thermal models, thus the temperatures of functional units can be analyzed efficiently to investigate the usefulness of the migrating computation DTM technique. By doing this, packaging complexity and cost can be significantly reduced and we can still get almost the same operating temperature and performance as if we use a much more expensive and complicated thermal package.

Chapter 6

Thermal Modeling Spatial Granularity

This chapter offers two different first-order analytical approaches to the spatial modeling granularity issue. The first approach is more intuitive by utilizing the well-known duality between thermal and electrical circuits. The second approach performs a more rigorous analysis in the spatial frequency domain. We also discuss some design implications from the spatial granularity analysis.

6.1 Overview and Related Work

As technology scales, more and more functions are put into a single silicon die, whereas the area of the die remains relatively constant according to ITRS [2]. At the same time, one of the trend in the architecture research is moving toward chip multiprocessor (CMP), which supports the continued performance scaling by taking advantage of the parallelism that the on-die multiple cores can provide. At some point in the future, the physical size of each individual core becomes so small that it is sufficient to model temperature, hence

power at the core level rather than at the microarchitecture level. Similar argument may hold for other future VLSI architectures. Therefore, to efficiently consider thermal effects in future VLSI systems such as CMPs, it is important to analyze the spatial granularity of temperature and power models. Additionally, modeling temperature at the proper granularity also reduces the complexity and computational effort of the compact thermal models while maintaining the accuracy at the same time.

To the best of our knowledge, the only existing work in the literature on the power and thermal modeling granularity issue is [24], where the authors also notice the spatial temperature low-pass filtering effect and derive a power modeling granularity with specific settings of power densities and package, by performing finite-element simulations. The discussions in [24] are not fully analytical. In this chapter, we provide two analytical approaches, which confirms and extends the findings in [24] for the granularity issues.

6.2 First Approach

Given the grid thermal model, it is clear that the desired temperature accuracy determines the minimum grid density, i.e. modeling granularity. Temperature difference across one grid cell should be less than a certain percentage of the maximum temperature across the die surface.

In order to derive the proper size of a grid cell, let us first start from a simple case. Assume there is a slab of material with unit width and infinite length. The thickness of the slab is t, and the bottom surface of the slab is isothermal. Half of the slab has a uniform power density of i, while the power density of the other semi-infinite half is 0, as shown in Fig. 6.1(a). The resulting temperature distribution of the top surface of the slab is approx-

imated in Fig. 6.1(b). The far end of the left half with power density i has a temperature of T_{max0}, and the far end of the right half with no power dissipated has a temperature of 0, for simplicity. Due to symmetry the temperature at the boundary between the two halves is $T_{max0}/2$. The temperature at point x can then be derived from the equivalent lumped thermal circuit in Fig. 6.2(a). The lateral and vertical thermal resistances of an infinitesimal portion of the slab with length of dx are

$$r_1 = \frac{dx}{k{\cdot}t} \quad \text{and} \quad r_2 = \frac{t}{k{\cdot}dx} \tag{6.1}$$

where k is the thermal conductivity of the slab material. Now, the equivalent thermal resistance for the semi-infinite half of the slab should be the same whether or not including the first vertical thermal resistance r_2, i.e.

$$R_{eq} = r_2||(r_1 + R_{eq}) \quad \text{and also} \quad R_{eq} = r_1 + (r_2||R_{eq}) \tag{6.2}$$

Solving the above two equations for R_{eq} leads to

$$R_{eq}^2 + r_1 R_{eq} - r_1 r_2 = 0 \quad \text{and} \quad R_{eq}^2 - r_1 R_{eq} - r_1 r_2 = 0 \tag{6.3}$$

When $dx{\rightarrow}0$, we have $r_1{\rightarrow}0$, also R_{eq} should have a finite value, which means we can neglect the term $r_1 R_{eq}$ in the above two equations of R_{eq}. Therefore, both equations become

$$R_{eq}^2 {\approx} r_1 r_2 = \frac{1}{k^2} \quad \text{i.e.} \quad R_{eq} \approx 1/k. \tag{6.4}$$

Next, to find the temperature $T(x)$, we consider the circuit in Fig. 6.2(b), in which heat $i(x)$ flows into node x. According to Kirchoff's Current Law, we have

$$\frac{T(x+dx)}{R_{eq}} + \frac{T(x)}{r_2} = i(x) = \frac{T(x)}{R_{eq}} \tag{6.5}$$

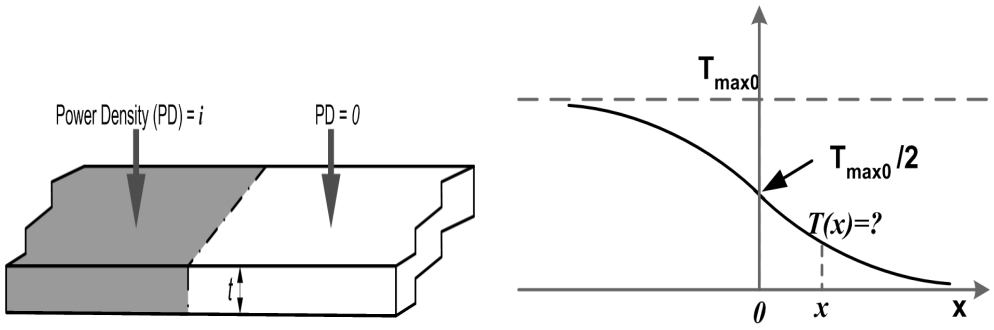

(a) (b)

Fig. 6.1: (a) A one-dimensional slab of material with heat flux on the left top surface. The bottom surface is isothermal. The right top surface and the side surfaces are adiabatic, i.e. there is no heat transfer through these surfaces. (b) Temperature distribution along the length of the slab [31].

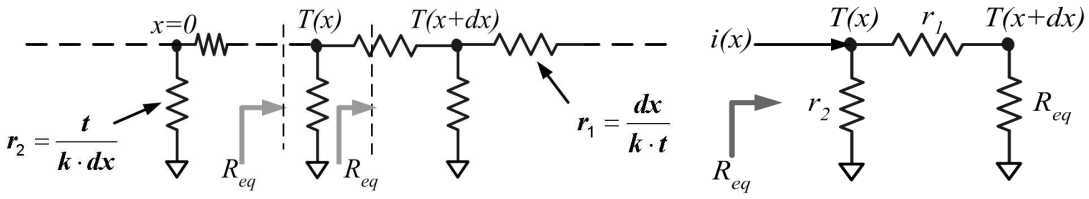

(a) (b)

Fig. 6.2: (a) Thermal resistance network for Fig. 6.1(a). (b) Thermal circuit at Node x for calculating temperature $T(x)$ [31].

Substituting R_{eq} with $1/k$, and r_2 with $t/(k{\cdot}dx)$, then re-arranging both sides of the equation, we reach

$$\frac{T(x+dx)-T(x)}{dx} = -\frac{T(x)}{t} \quad <=> \frac{dT}{T} = -\frac{1}{t}dx \tag{6.6}$$

Taking the integral for both sides from $T_{max0}/2$ to $T(x)$ and from 0 to x, respectively, and solving for $T(x)$, we obtain

$$T(x) = \frac{T_{max0}}{2}e^{-\frac{x}{t}} \quad (x \geq 0) \tag{6.7}$$

The above equation shows that the temperature distribution for the right half of the slab

is approximately an exponential decay curve with a "spatial" constant of t, which is the thickness from the surface under consideration to the isothermal surface. Furthermore, we can write the temperature distribution of the left half of the slab as a function of position x according to symmetrical nature of the slab structure:

$$T(x) = T_{max0} - \frac{T_{max0}}{2} e^{\frac{x}{t}} \quad (x < 0) \tag{6.8}$$

Fig. 6.3 confirms the accuracy of the above analysis by comparing with FEM simulations using FloWorks.[1]

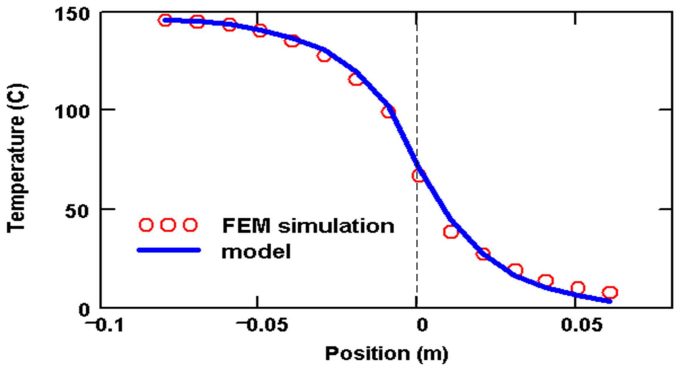

Fig. 6.3: Comparing FEM simulation result with (6.7) and (6.8) for the structure in Fig. 6.1(a). Power density=0.5W/mm^2, $t \approx 20mm$. Bottom surface of the silicon slab is approximately isothermal [31].

Next, we consider the scenario where heat is dissipated on a finite part of the slab, as shown in Fig. 6.4(a). The corresponding FEM-simulated temperature distributions within that part of the slab are shown in Fig. 6.4(b) for different block sizes ($w_2 > w_1$). It is obvious in Fig. 6.4(b) that if the size is sufficiently small, the heated part of slab does not actually reach its maximum temperature as can be seen in Fig. 6.1(b). This is due to the above mentioned spatial constant and it means that a block with small size acts

[1]FloWorks is an FEM software analyzing computational fluid dynamics and heat transfer. http://www.nika.biz/index2.htm

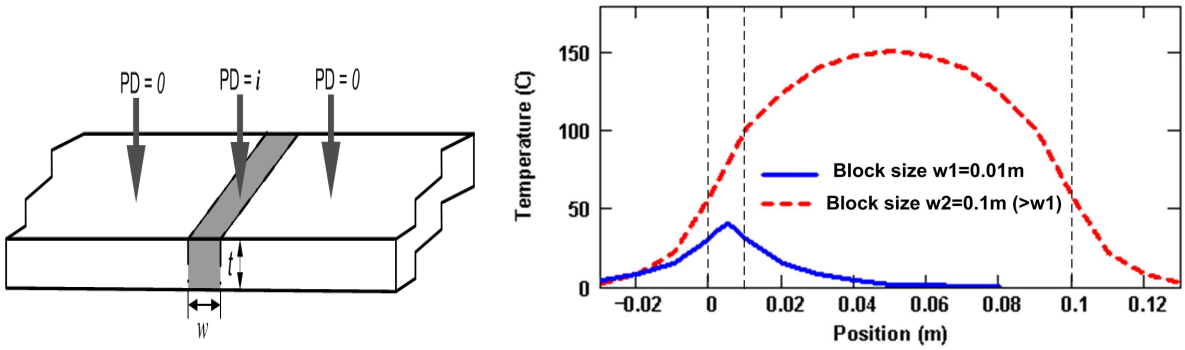

(a) (b)

Fig. 6.4: (a) Part of the slab of material dissipating power. The size of the part is w. (b) FEM simulation results—temperature distribution along the slab with different sizes dissipating power ($w_1 < w_2$). Smaller size "filters" out the temperature difference. (Silicon thickness $t = 20mm$) [31].

as a temperature spatial low-pass "filter" that prevents the temperature from reaching the maximum possible value. In contrast, with the same power density, a bigger block with its size much larger than the "spatial" constant t can have significantly higher temperature differences. The above analysis explains the "abnormal" observations that although some tiny structures such as clock buffers in a microprocessor have very high power densities, they don't necessarily cause hot spots, due to this "spatial temperature filtering" effect.

For a particular grid size w, from Eq. (6.8) and Fig. 6.4(b), the temperature difference within the grid is

$$\Delta T_{grid} \leq T\left(\frac{-w}{2}\right) - T(0). \tag{6.9}$$

By setting

$$T\left(\frac{-w}{2}\right) - T(0) = \frac{T_{max0}}{2}(1 - e^{\frac{-w}{2t}}) = p\,\%{\cdot}T_{max0} \tag{6.10}$$

where $p\,\%$ is the tolerable percentage error, and T_{max0} now represents the maximum possible temperature difference across the silicon die[2]. Solving Eq. (6.26), we get the lower

[2] T_{max0} is dependent on the power distribution across the die, and cannot be known *a priori* without

bound for w:

$$w = 2 \cdot t \cdot \ln \frac{1}{1 - 2 \cdot p\%} \tag{6.11}$$

Note that t, the thickness from the surface under consideration to the isothermal surface, needs to be calculated first. An isothermal surface is an ideal concept that is not found in real packages, but surfaces with negligible temperature differences can be considered as isothermal [3]. For instance, the thickness t for the silicon surface of the package shown in Fig. 4.1 can be found by adding up the thickness of silicon substrate, thermal interface material and heat spreader. An important detail is that, if we use the conductivity of silicon in the above equations, we need to first convert the actual thicknesses of the thermal interface material and the heat spreader to "equivalent" silicon thickness by multiplying their thickness by the ratio of their thermal conductivities to the one for silicon.

Fig. 6.5 plots the required grid size for different desired levels of precision according to (6.11), with equivalent thickness $t = 4$mm and $t = 2$mm. The horizontal axis is the ratio of ΔT_{grid} to T_{max0}, i.e. p, in percentage. For example, consider that we have a 20mm$\times$20mm silicon die, and the maximum possible temperature difference across the die is 30 degrees, from Fig. 6.5, we can find that if we desire all grid temperature error of less than 3% ($30 \times 3\% \approx 1$ degree) for $t = 4$mm, a grid size of about 0.5mm is sufficient. This corresponds to dividing the die into 40$\times$40 grid cells. Any finer grid size is unnecessary in this case[4].

performing thermal analysis. But one can always start with a reasonable guessed value for T_{max0} based on previous design experience, then solve the thermal model and iterate the analysis for a few times to get the needed grid size.

[3] One example is the bottom surface of the heat spreader since the heat spreader is usually made of materials with high thermal conductivity, such as copper.

[4] It is worth noting that the above granularity analysis is based on simplifications of classical heat transfer equations, which underestimates temperature when applied at size scales less than the phonon-phonon mean free path (about 300nm for silicon at room temperature) [65]. Thus, for granularity analysis at the transistor level, the phonon Boltzmann Transport Equation (BTE) should be used instead.

One assumption that we have made so far is that the power is uniform within each grid cell. This assumption is legitimate if the thermal analysis is performed at early design stages, because detailed layout and power information are not available yet. In later design stages, the structures that are included in one grid cell may turn out to be heterogenous. In this case, we can always first resort to finer grid cells inside which power distribution can be considered as uniform, then perform the above accuracy analysis and decide whether or not that finer grid size is necessary or not. Due to the "spatial temperature filtering effect" mentioned above, often we should find that temperature difference within a finer grid cell is negligible and we need to come back to larger grid cells, unless the power density is extremely high.

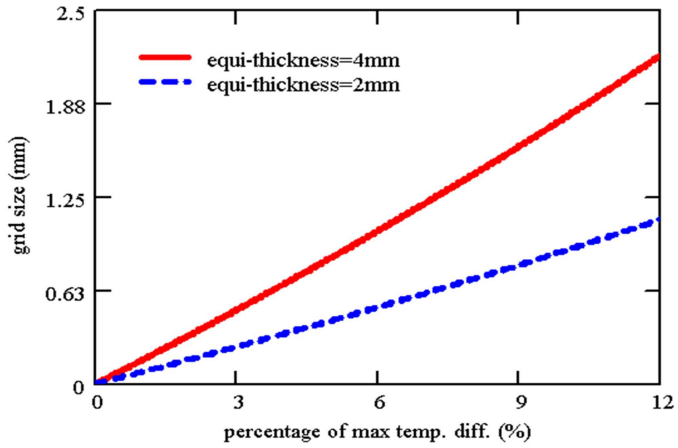

Fig. 6.5: Minimum necessary grid size for different desired levels of precision, with $t = 4$mm and $t = 2$mm, respectively. X-axis is the ratio of ΔT_{grid} to T_{max0}, i.e. p, in percentage. For example, for a system with $t = 4$mm, if 3% of temperature precision is desired, from the solid line, one finds that 0.5mm grid cell size would be enough [31].

6.3 Second Approach

As we can see, the first approach to the spatial granularity of thermal and power modeling uses the intuition of lumped vertical and lateral thermal resistance network and extend it to the integral form. In this section, we present another more rigorous theoretical approach, which is based on the spatial frequency domain analysis.

In mathematics and physics, *spatial* frequency is an attribute of any quantity that is periodic in space. It is a measure of how often that quantity is repeated per unit distance (e.g. per meter). It is defined as

$$f_s = \frac{1}{\lambda} \tag{6.12}$$

where f_s denotes the spatial frequency, λ is the period or wavelength of the repeating patten.

For the purpose of illustration, we first perform traditional temporal frequency-domain analysis for a first-order electrical R-C circuit, we then utilize the analogy between the temporal frequency (in s^{-1} or Hertz) and the spatial frequency (in m^{-1}) to extend the analysis from time to space.

For an electrical capacitor, C, assume the voltage drop between its two terminals is a sinusoidal form with an angular frequency ω, that is, $V_c(t) = V_0 cos(\omega t + \phi)$, or in the complex exponential form, $V_c = V_0 e^{j(\omega t+\phi)}$. From circuit theories, the current flow through the capacitor $I_c(t)$ is

$$I_c = C\frac{dV_c}{dt} = C\frac{d}{dt}V_0 e^{j(\omega t+\phi)} = j\omega C \cdot V_c \tag{6.13}$$

Thus, the *electrical* impedance of the capacitor is

$$Z_C = \frac{V_c}{I_c} = \frac{1}{j\omega C} \tag{6.14}$$

Now, consider the electrical circuit in Fig. 6.6, which has a resistor R, a capacitor C and a sinusoidal voltage source $V_s(t) = V_0 cos(\omega t + \phi)$. All of them are in series. We all know that this circuit is a low-pass filter, that is, the voltage drop across the capacitor tracks the input voltage $V_s(t)$ at low frequency, and is increasingly attenuated at higher frequency. In other words, the equivalent impedance of this circuit is $Z_R || Z_C = R || (\frac{1}{j\omega C})$, which is $Z_R = R$ at DC, and is approaching zero at high frequencies, thus comes the term "low-pass filter". The resistor R determines the "DC" component of the output voltage, whereas the capacitor determines the "AC" component.

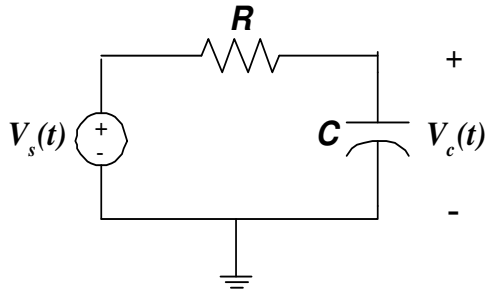

Fig. 6.6: A first-order electrical RC circuit.

In space, there is also this "low-pass filtering" effect for temperature distribution, as we have seen from the first section of this chapter. Here, we extend the temporal frequency analysis to the one-dimensional spatial frequency domain. Consider a sinusoidal heat flux $q(x)$, which causes a sinusoidal temperature distribution

$$T(x) = T_0 cos(\omega_s x + \phi) = T_0 e^{j(\omega_s x + \phi)} \tag{6.15}$$

where $\omega_s = 2\pi/\lambda$ is the spatial angular frequency (the subscript s means "spatial"), and x is the position in the 1-D space. The governing equation of heat transfer is Fourier's Law

$$q(x) = k\frac{dT(x)}{dx} = k\frac{d}{dx} T_0 e^{j(\omega_s x + \phi_s)} = j\omega_s k T(x) \tag{6.16}$$

where k is the thermal conductivity. The minus "-" sign in Fourier's Law goes away if we define $dT(x)$ as the temperature decrease (high temperature minus low temperature). Notice the similarity between Eq. (6.16) and Eq. (6.13). This leads us to some quantity analogous to the electrical capacitor in the spatial domain for heat transfer. We call it *thermal spatial capacitive impedance*, and write it as

$$Z_{Cs} = \frac{T}{q} = \frac{1}{j\omega_s k} = \frac{T}{q} = \frac{1}{j\omega_s C_s} \tag{6.17}$$

where C_s is defined as *thermal spatial capacitance* (notice that C_s is completely unrelated to the thermal capacitance C_{th} that we defined earlier in this dissertation that determines the *transient* heat transfer), and Z_{Cs} is the "thermal spatial impedance". The subscript "s" denotes the spatial nature of these definitions. The unit of both C_s and Z_{Cs} is m^2K/W, which is different from the unit of the thermal resistance we used earlier in this dissertation (in K/W). This is legitimate because we use heat flux (in W/m^2) instead of heat transfer rate (in W), i.e. the thermal impedance and resistance in this section is defined as the temperature drop divided by the heat flux (power density), not by heat transfer rate (power).

Eq. (6.17) is used when there is an AC component, with spatial frequency ω_s, in the applied heat flux. In the case where there is only DC heat flux, Fourier's Law leads to the traditional definition of thermal resistance

$$Z_{Rs} = \frac{L}{k}. \tag{6.18}$$

where L is the distance from the active silicon surface to the isotherm in the package, it is the same as t in the first section of this chapter. Also note that this DC spatial thermal impedance also has the unit of m^2K/W, which is consistent with the unit of the AC spatial thermal impedance Z_{Cs}. From the above derivation, naturally we can reach a first-order

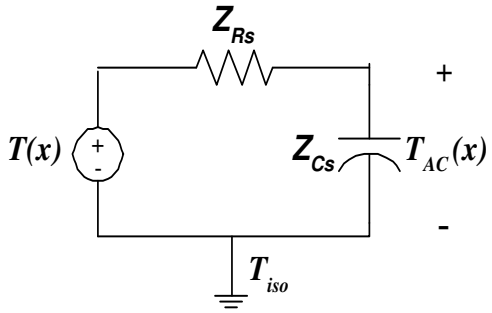

Fig. 6.7: The Thevenin equivalent first-order thermal spatial "RC" circuit.

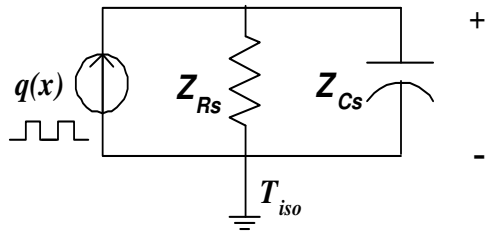

Fig. 6.8: The Norton equivalent first-order thermal spatial "RC" circuit.

spatial thermal RC circuit as shown in Fig. 6.7. To make it more comprehensible, Fig. 6.8 shows a more intuitive Norton equivalent circuit of Fig. 6.7. The heat flux generated by the active silicon layer is written as $q(x)$, which models the non-uniform distribution of power density across the chip. The DC component in the spatial temperature distribution is determined by Z_{Rs}, whereas the AC component is determined by Z_{Cs}. In addition, the total equivalent thermal spatial impedance is

$$Z_{s_{eq}} = Z_{Rs}||Z_{Cs}. \tag{6.19}$$

If we plot the Bode plot of $Z_{s_{eq}}$ with respect to the spatial frequency ω_s in Fig. 6.9, we can see that for low spatial frequencies (power sources with large dimensions), the thermal impedance is close to the DC component, that is the lumped $R_{th} = L/(kA)$ as we derived in earlier chapters (A is the corresponding vertical heat conduction area). But for high spatial

frequencies (power sources with small dimensions), the impedance attenuates to smaller values due to the presence of the "thermal spatial capacitance". This explains the spatial temperature filtering effect—for the same power density, structures with tiny dimensions have lower peak temperature comparing to their large counterparts applied with the same power density or heat flux.

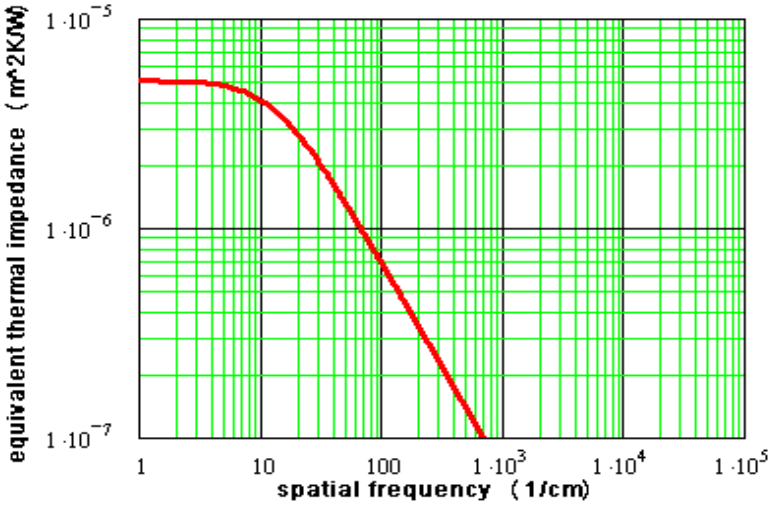

Fig. 6.9: The thermal spatial "RC" circuit is low-pass filter in the spatial frequency domain.

In the above derivations, we only considered the heat fluxes as sinusoidal waveforms with a specific spatial frequency. In real designs, the power density (or heat flux) is usually considered unform over the structure of interest, that is, the input heat fluxes should be modeled as pulses or square waves. Therefore, we need to find the spatial heat transfer response to a square wave heat flux input, with q_{low} and q_{high} representing the high and low levels of the heat flux pulses. Since we only care about the temperature rise here, it is convenient to set q_{low} to zero.

Again, we first resort to the first-order electrical R-C circuit (Fig. 6.6) for some insights. The transient response of this circuit to a square-wave input (with frequency ω) is a

sawtooth waveform as shown in Fig. 6.10. The expression for this sawtooth curve is

$$\begin{aligned} V_c(t) &= V_0 + (V_L - V_0)e^{-\frac{t}{RC}} \quad \text{for half period, and} \\ V_c(t) &= V_H e^{-\frac{t}{RC}} \quad \text{for the other half period.} \end{aligned} \tag{6.20}$$

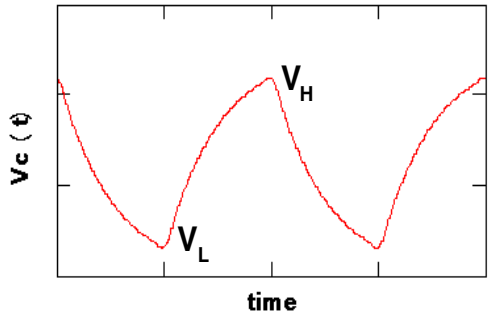

Fig. 6.10: The response of the first-order electrical RC circuit to square-wave input voltage is a sawtooth waveform over time.

V_H and V_L can be solved from Eq. (6.20) by setting t to half of a period ($T/2$), which results in

$$\begin{aligned} V_H &= V_0 \frac{1-e^{-\frac{T}{2RC}}}{1-e^{-\frac{T}{RC}}} \quad \text{and} \\ V_L &= V_0 \frac{e^{-\frac{T}{2RC}} - e^{-\frac{T}{RC}}}{1-e^{-\frac{T}{RC}}} \end{aligned} \tag{6.21}$$

If the thermal response in the spatial domain to the square-wave heat flux input is also a sawtooth-like temperature distribution, this is directly in conflict to what we derived in the first approach in this chapter. In addition, intuition tells us that the spatial temperature response to square-wave input should be symmetrical, i.e. no biased temperature waveform should be present along the x axis.

This paradox originates from a fundamental assumption that we take for granted during the analysis of the electrical R-C circuit—the circuit is solved in the time domain, and time has a direction! Time only travels in one direction ($+t$), nobody can travel back in time

(at least for now). Whereas in the case of solving the spatial thermal "R-C" circuit, this assumption is not true. Instead, the correct way is to model heat flux waveform in both directions, $+x$ and $-x$. So the input heat flux should be written as

$$q'(x) = \frac{1}{2}q(x) + \frac{1}{2}q(-x + \frac{\lambda}{2}) \tag{6.22}$$

where $q(x)$ is the heat flux written with components assuming only one direction, $q'(x)$ is the heat flux written in the symmetrical fashion. $q'(x)$ is the superposition of two waves in opposite directions and has the same shape as $q(x)$.

If we solve the spatial thermal "R-C" circuit using $q'(x)$ as the input, the resulting temperature distribution is the superposition of two sawtooth waveforms in opposite directions (Fig.6.11)

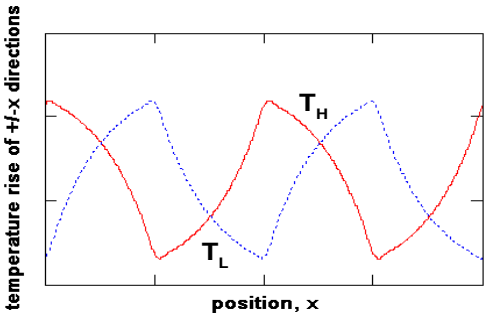

Fig. 6.11: In space, the thermal spatial low-pass filter's response can be modeled as the superposition of two sawtooth waveforms in opposite directions.

$$\begin{aligned} T_{+x}(x) &= \tfrac{T_0}{2} + (T_L - \tfrac{T_0}{2})e^{-\frac{x}{L}} \quad \text{for half period of +x,} \\ T_{+x}(x) &= T_H e^{-\frac{x}{L}} \quad \text{for the other half period of +x.} \end{aligned} \tag{6.23}$$

$$\begin{aligned} T_{-x}(x) &= \tfrac{T_0}{2} + (T_L - \tfrac{T_0}{2})e^{-\frac{-x+\lambda/2}{L}} \quad \text{for half period of -x, and} \\ T_{-x}(x) &= T_H e^{-\frac{-x+\lambda/2}{L}} \quad \text{for the other half period of -x.} \end{aligned} \tag{6.24}$$

where T_L and T_H can be solved in a similar way to Eq. 6.21. $T_0 = q_{high} \cdot Z_R$ has the same meaning as T_{max0} used in the first section of this chapter. The superposition of the two sawtooth curves is symmetrical as in Fig. 6.12. Notice the resemblance between Fig. 6.12 and Fig. 6.1(b).

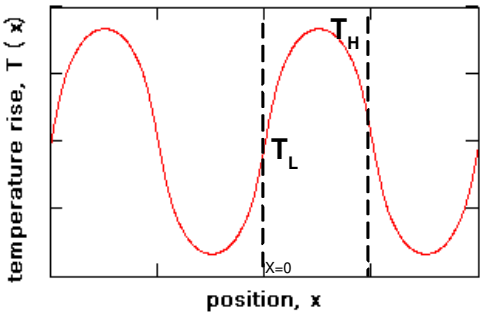

Fig. 6.12: The resulting temperature distribution of the superposition of two sawtooth waveforms in space.

When the size of the structure, i.e. the granularity, $w = \lambda/2 \rightarrow \infty$, which also leads to $T_L = 0$ and $T_H = T_0 = T_{max0}$. Therefore,

$$\begin{aligned} T(x) &= T_{max0} - \tfrac{T_{max0}}{2} e^{-\frac{x}{L}} \quad \text{for } x > 0 \\ T(x) &= \tfrac{T_{max0}}{2} e^{\frac{x}{L}} \quad \text{for } x <= 0 \end{aligned} \tag{6.25}$$

we notice that Eq. (6.25) is the same as the combination of Eq. (6.7) and (6.8). However, the second approach provides a more detailed and more accurate expression for spatial temperature distributions at different granularities. With the aid of Fig. 6.12, we can calculate the difference between hottest temperature (at the center of a structure, $x = w/2$) and the coolest temperature (at the edge, $x = 0$). Substituting T_L using Eq. (6.21), after some rearrangement of the equations, we get

$$\Delta T_{grid} = T_0 \left(\frac{1}{2} + \left(\left(\frac{e^{-\frac{w}{L}} - e^{-\frac{2w}{L}}}{1 - e^{-\frac{2w}{L}}} - 1 \right) e^{-\frac{w}{2L}} \right) \right) \tag{6.26}$$

Setting this temperature difference to some precision level $p\%$ that we desire, similar to what we did in the first approach, we can get a more accurate solution of w for different p. If we assume the same 4mm silicon-to-isotherm thickness as in the first section of this chapter, and let $p\% = 3\%$ and $T_{max0} = 30°C$, we get $w = 2.8mm$, which is much looser than the first approach. This is because ΔT_{grid} is not overestimated here. On the other hand, if we desire a $0.1°C$ accuracy with $T_{max0} = 30°C$, that is $p\% = 0.33\%$, the cell size derived from this analysis is around $1mm$.

The limitation of this analysis is that it takes into account the lateral spatial temperature gradient, but not the vertical gradient. A more accurate analysis would use multi-ladder, or ideally, distributed thermal spatial thermal R-C circuit, which results in about half spatial constant $\tau \approx L/2$. Thus, a more pessimistic and tightly-bounded granularity is about half of the granularity derived here. That is, for a $0.1°C$ accuracy with $T_{max0} = 30°C$, the proper cell size would be around $500\mu m$. For a rigorous analysis, we can use three-ladder RC circuit den write the temperature response to square-wave heat fluxes. This will be interesting future work, and is not treated in this dissertation.

Because the heat transfer in x and y directions are orthogonal, which is determined by the 2-D form of Fourier's Law, the above derivations can also be easily extended into two-dimensional space with the same results.

6.4 Implications

In addition to the more confidence added to the accuracy of temperature estimations of the thermal models, the above-mentioned granularity analysis also provides useful insights to temperature-aware architecture and circuit designs. An important question to ask during

thermal analysis of the design process is: following the technology scaling, what is the right power granularity that is necessary for accurate thermal analysis? Is it going to be modeled per transistor, standard cell, functional unit, or individual processor core? From the above analysis, it is more likely that we will be considering thermal impacts at the functional unit, or even at the core level in the near future. At that time, temperature distribution at lower granularities will more likely be uniform.

The spatial low-pass filtering phenomenon also helps explain why some of the tiny structures with very high local power density, such as clock buffers, do not impose serious thermal hazards in the chip.

Chapter 7

Validation

In this chapter, we present several works that we have done to validate the HotSpot compact thermal models. The validation for the primary heat transfer path has been done with a finite element simulator, thermal test chip measurement, and an FPGA system with embedded thermal sensors. The interconnect thermal model has been validated by comparing with finite element simulation result published in the literature. Another important verification of the HotSpot thermal model is its boundary condition independence (BCI).

7.1 Finite Element Simulations

We first build a finite-element model of the chip and the package components for the primary heat transfer path in a commercial finite-element software FloWorks, shown in Fig. 7.1. We use the same Alpha EV6 processor floorplan and and apply the same power distribution for the die in both the FloWorks model and the HotSpot thermal R-C model. We neglect the secondary heat transfer path in both models [74, 78]. Fig. 7.2 and Fig. 7.3 show the comparison of the results between the HotSpot compact model and the FloWorks FEM model. Also shown in Fig. 7.2 and Fig. 7.3 are the results from a "simplistic" thermal R-C model where the lateral thermal resistances are eliminated. HotSpot shows good agreement with FloWorks, with errors less than 5.8%. The simplistic model, however, has larger errors, as high as 16%. This indicates that the lateral heat transfer inside the silicon and the package cannot be omitted.

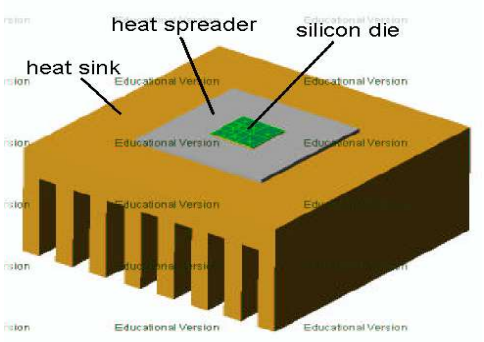

Fig. 7.1: Finite-element model in FloWorks [78].

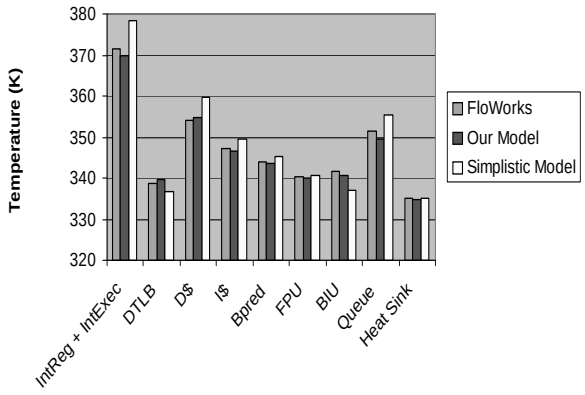

Fig. 7.2: Finite-element validation for steady-state temperatures with ALPHA EV6 floorplan and package [78].

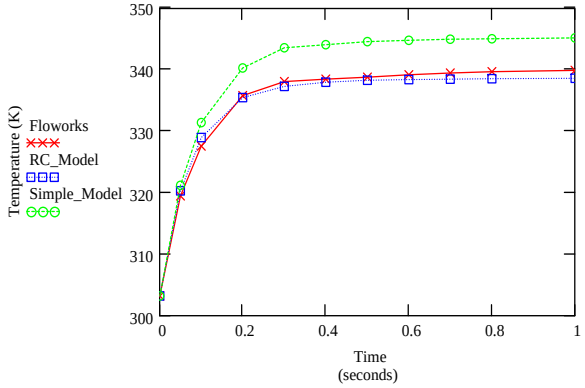

Fig. 7.3: Finite-element validation for transient temperatures with ALPHA EV6 floorplan and package [78].

7.2 Thermal Test Chip

We also validated our compact thermal modeling approach with a commercial thermal test chip [85]. The thermal test chip has a 9×9 grid of power dissipators, which can be turned on or off individually. Each grid cell also has an embedded thermal sensor. The test chip can be used to measure both steady-state and transient temperatures for each of the grid cells. We built the same 9×9 grid-like compact thermal model following the HotSpot approach. We then turned on particular sets of power dissipators in the test chip

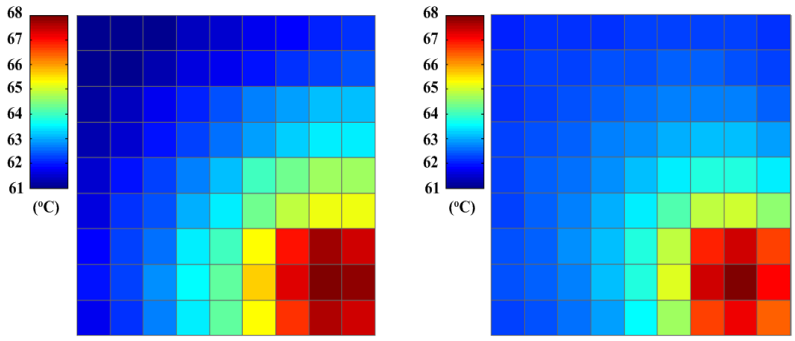

(a) (b)

Fig. 7.4: (a) Test chip steady-state measurements (b) Steady-state results from the model with errors less than 5% [36].

and assigned the same amount of power at the same locations in the thermal model. In the comparison, we neglected the secondary heat flow path, because the test chip is wire bonded and plugged in a plastic socket that has very low thermal conductivity.

Fig. 7.4(a) and (b) shows the steady-state thermal plots using measurements from the test chip (running for over 20 minutes) and results from our thermal model. Transient temperature data from the thermal model are also compared with the test chip transient measurements, as shown in Fig. 7.5. The percentage error between our model and the test chip measurements are calculated by $(T_{model} - T_{chip})/(T_{chip} - T_{ambient})$. The power density in this experiment is $50\text{W}/\text{cm}^2$ in the heat dissipating area (the 3×3 lower-right corner). As can be seen, the HotSpot model is quite accurate, with the worst case error values for steady-state temperatures and transient temperatures less than 5% and 7%, respectively. Note that the transient temperature response becomes flat and converges to the steady state after hundreds of seconds, which is beyond the time scale shown in Fig. 7.5. This long time interval to reach the steady-state temperature is caused by the huge thermal RC time constant of the heat sink.

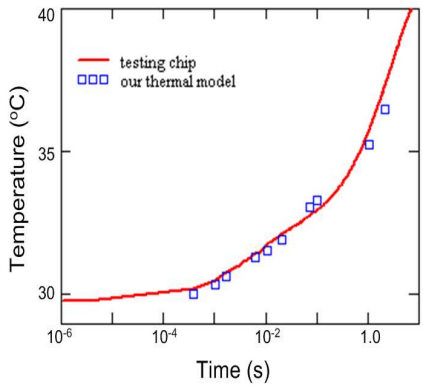

Fig. 7.5: (Transient validation of the compact thermal model. Percentage error is less than 7%. (Transient temperature response of one power dissipator is shown here.) [36]

Recently we have also built a full-chip and package compact thermal model using HotSpot modeling approach for an IBM microprocessor. The model has been validated both quantitatively with a detailed FEM model developed by IBM packaging engineers and qualitatively with on-chip temperature sensor measurements under a typical workload.

7.3 FPGA-Based System Prototype

We designed an FPGA-based system that monitors the temperature at various locations on a Xilinx Virtex-2 Pro FPGA[1] [89, 90]. The system is composed of a controller interfacing to an array of temperature sensors that are implemented on the FPGA fabric. We use ring oscillators as temperature sensors by exploiting the fact that the frequency of oscillation is approximately proportional to temperature [54]. Calibrations are done for 6 different sensors placed near the center of each unit on the die. Power consumption for different units is extracted through various methods. Using the floorplan shown in Fig. 7.6, we compare the sensor readings with values obtained from the corresponding HotSpot model.

[1] Xilinx Virtex-2 Pro user guide—http//direct.xilinx.com/bvdocs/publications/ds083.pdf

The results are in Table 7.1. We see that on average the temperatures predicted by the HotSpot thermal model and those obtained from the sensors differs by less than 0.2°C.

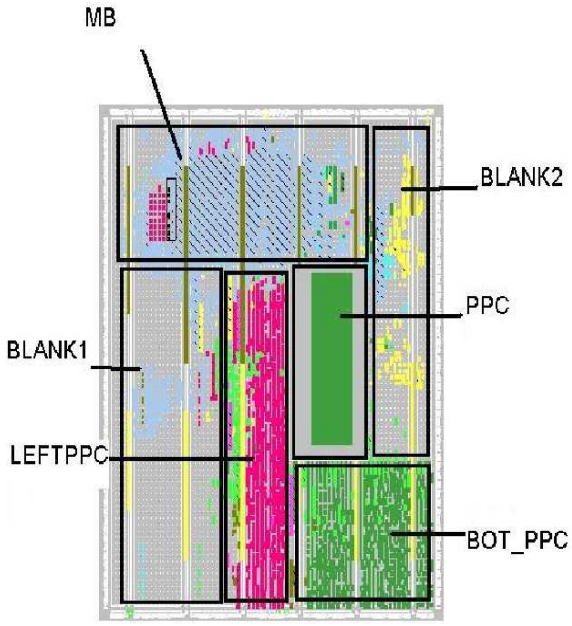

Fig. 7.6: Floorplan with 6 functional blocks implemented in FPGA, for HotSpot primary heat transfer path model validation [90].

The low temperature rise (4.1°C maximum) and small temperature difference across the FPGA chip (0.7°C maximum) is due to the fact that typical operating powers for the PPC and MB blocks on the FPGA are not significant enough to heat up the chip (because of this, the FPGA chip is not equipped with a heatsink). In order to achieve greater across-die temperature differences, we have intentionally left two "zero-power" blank blocks (blank1 and blank2). Regardless of the relatively cool die temperature, the errors between the HotSpot model and the thermal sensor measurements are within 10% of the measured temperatures, for example, for "MB" in Table 7.1, the percentage error is $(4.1 - 3.96)/4.1 = 3.4\%$. This confirms the validity of the HotSpot model, although the FPGA application itself doesn't show much interesting "hot" temperatures and temperature gradients.

Unit	Power(mW)	Sensor Temperature	HotSpot Temperature
blank1	0.1	3.4	3.37
left_ppc	75	3.5	3.69
bott_ppc	75	3.4	3.67
ppc	45	3.5	3.66
mb	313	4.1	3.96
blank2	0.1	3.4	3.38

Table 7.1: Comparisons of temperature readings from the FPGA and the HotSpot thermal model. Temperatures are with respect to ambient temperature. Errors are within 0.2°C ([31]).

7.4 Interconnect Thermal Model Validation

To validate the interconnect thermal model, we compare our model to the finite-element models in [69]. The authors of [69] use two interconnect test structures in FEM analysis software: one with individual metal wires on top of each other (corresponding to the case of Fig. 4.8(a)); and the other one with multiple metal wires within each layer (corresponding to the case of Fig. 4.8(b)). Both test structures have four metal layers at 0.6μm technology. We apply exactly the same settings to our interconnect thermal model as in [69], and perform the same two experiments—1) for the stacked single-wire test structure, apply different power for each wire and obtain the temperature rise with respect to ambient temperature; 2) for both test structures, apply different current density for each layer and obtain the temperature rise. Fig. 7.7(a) and (b) show the comparisons for both experiments. As can be seen, the results of our interconnect thermal model match FEM simulation results very well, which gives us confidence that the way we model the equivalent thermal resistance

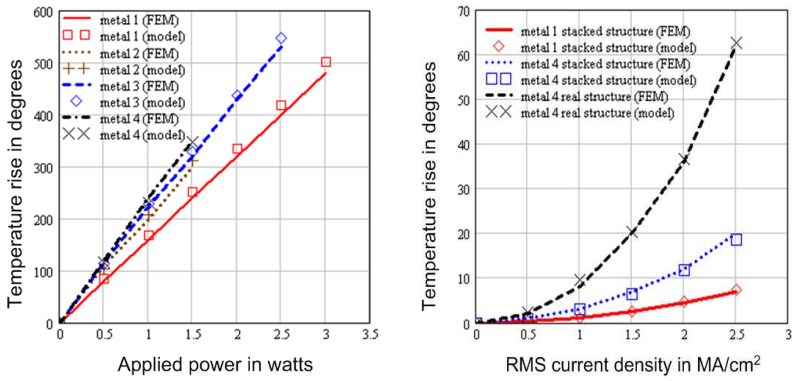

(a) (b)

Fig. 7.7: Interconnect thermal model validation. (Lines—FEM results from [69]. Markers—our thermal model results.) (a) stacked single wires—powers are applied to each wire (b) RMS current densities are applied to both test structures [36].

for wires and their surrounding dielectrics is appropriate for preliminary explorations of the design space.

7.5 Boundary Condition Independence (BCI)

Another important aspect of compact thermal modeling is boundary condition independence (BCI) [46, 47, 68]. Achieving BCI is essential to compact thermal models. If the model changes whenever the boundary conditions change, the model would be almost useless. Traditionally, researchers in the package compact thermal modeling community usually adopt the DELPHI approach to achieve BCI, that is, finding a thermal resistance network with minimum overall error when applied to different boundary conditions. The resistance values are extracted from detailed thermal simulations with the same package structure. Such simulations can be performed in numerical analysis tools.

When using our modeling approach, because there is no data extraction procedure and

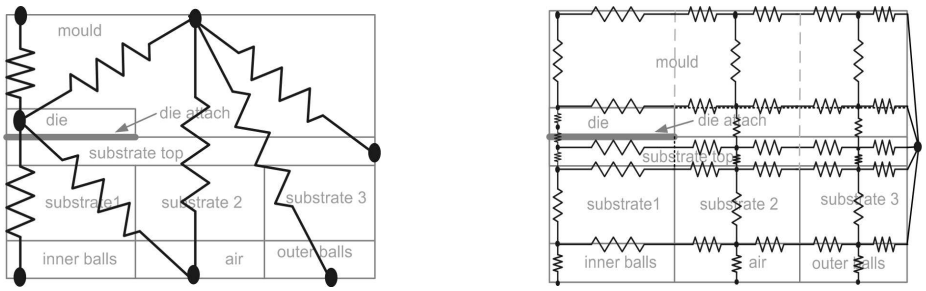

(a) (b)

Fig. 7.8: Thermal resistances network for (a) the BCI DELPHI model and (b) our thermal model of a DELPHI BGA benchmark chip, extracted from [43].

all the resistance values are calculated from the physical dimensions and properties of the materials, the model itself should be reasonably BCI if the major heat transfer paths are properly modeled. At the present state, in order to validate that our modeling method can indeed achieve reasonable BCI, we compare our compact thermal models with the DELPHI models. Since the DELPHI models have been extensively validated to be BCI for a large set of boundary conditions, we regard the DELPHI models as proper benchmarks for the BCI validations of our modeling approach. Comparisons with real designs will be part of our future work. One of our BCI validations is done by comparing with a DELPHI BGA benchmark chip in [43]. The dimensions of the BGA benchmark chip and package and the set of boundary conditions are both taken from the specifications in [43, 47]. The model structures of both the DELPHI model and our model are shown in Fig. 7.8. In this comparison, the notion of quarter symmetry can be applied because there is only one node needed for the die in both models. Therefore, only a quarter of the package is sketched for our model and the DELPHI models in Fig. 7.8. The first four standard boundary conditions in Table 7.2 correspond to different typical settings of heat transfer coefficients for the top, bottom, side and the lead of the package. The last boundary condition accounts for the

ideal case where the heat transfer coefficients of all the surfaces are infinite [43, 47].

#	b.c.	our model (HotSpot)	DELPHI	error
1	DCP-1	16.79	16.68	0.66%
2	DCP-2	19.94	20.00	-0.30%
3	DCP-3	66.42	62.78	5.80%
4	DCP-4	2960.00	3070.00	-3.58%
5	infinite	10.20	10.56	-3.41%

Table 7.2: Comparison of our compact thermal model and DELPHI model for the DELPHI BGA benchmark chip under the same set of boundary conditions. Temperatures are in Celsius and with respect to ambient temperature [35].

The temperature readings from both models are also listed in Table 7.2. The heat generated at the die surface is 2.5W, which is used as the input to both models. As can be seen from Table 7.2, our model achieves reasonable BCI. For the listed five standard boundary conditions, it yields almost the same temperature readings as the DELPHI model. The worst case percentage error is 5.8%. One possible reason of the error is that the top surface division ratio is fixed according to the area of the smaller neighbor layer, in this case, it is the die area. This division ratio might not be exactly the optimal ratio, but it is sufficiently near the optimal ratio. In a previous work [8], the author argues that the surface division ratio should be determined by the heat flux distribution on a particular surface. He also shows that the heat flux distribution function $f(-)$ of the top surface develops a peak just above the die area. Therefore, using the die area to divide the top surface as in our model is reasonable.

Chapter 8

HotSpot Applications

In this chapter, some successful applications of the HotSpot compact thermal modeling method are presented. These are works resulting directly from this dissertation research or collaborated works related to this dissertation research. In addition to these presented works, HotSpot has also been well adopted by both academia and industry. There are numerous independent research activities reported in the literature that take advantage of the infrastructures and facilities that HotSpot provides to make contributions to the general research area of thermal-related VLSI design, for examples, [12, 22, 38, 52, 81]..., etc.

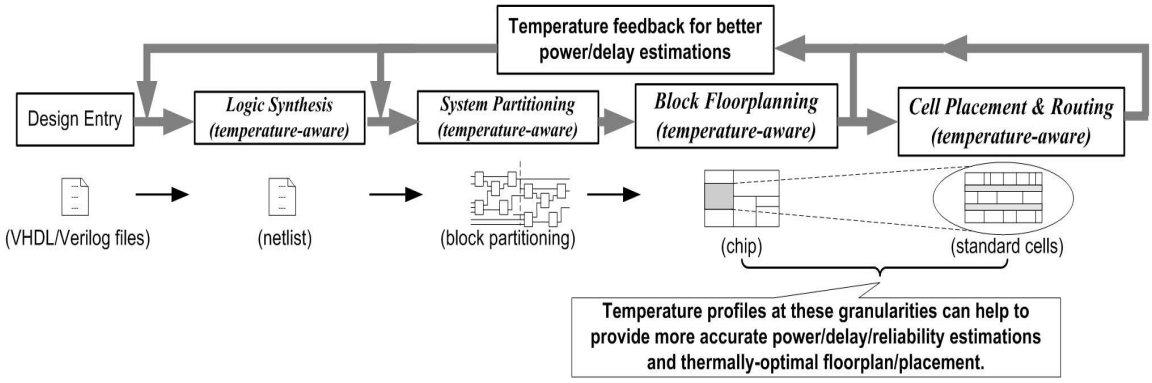

Fig. 8.1: An example of temperature-aware ASIC design flow [36].

8.1 Temperature-Aware ASIC and SoC Design

In sub-100nm and nano CMOS technologies, early accurate design estimation is key to high-level design convergence and should ensure careful consideration of deep submicron effects (including power, performance, reliability, etc.) [40]. Temperature plays an important role in early accurate estimations of power, performance and reliability. In addition, thermal effects are influenced by placement and routing; for example, putting two hot blocks adjacent to each other will exacerbate the hot spots, while surrounding a hot block by several colder blocks will actually help in cooling down the hot spot. Temperature should thus be included in the cost function in order to achieve optimal placement and routing in sub-100nm designs. Temperature can also affect manufacturability in terms of packaging and choices of process if the design is thermally limited. Fig. 8.1 shows a simplified ASIC design flow adapted to become temperature-aware. Temperature profiles are needed at both functional-block level and standard-cell level during the ASIC design flow. Similar arguments also apply to microprocessor and SoC design flows.

From above, we see that it is very important to be able to estimate temperature at differ-

ent granularities and at different design stages, especially early in the design flow. The estimated temperature can then be used to perform power, performance and reliability analyses, together with placement, packaging design, etc. As a result, all the decisions use temperature as a guideline and the design is intrinsically thermally optimized and free from thermal limitations. We call this type of design methodology *temperature-aware* design. The idea of temperature-aware design is unique because operating temperature is properly considered during the *entire* design flow instead of being determined only after the fact at the end of the design flow. There are a few examples of previous work about temperature-related design—for example, in [86], the authors present a design flow from digital simulations to a thermal map at the end of the design. This work is useful, but the design flow therein cannot be termed as a proper temperature-aware design since none of the intermediate design stages have closely considered temperature-related issues such as power or performance estimations, placement, thermal analysis, etc. Thus the design decisions of these stages are not optimized, and the design has to restart from the beginning if it turns out to be thermally limited.

The HotSpot thermal models can be utilized to achieve accurate preliminary design estimations and precise run-time thermal management techniques. As an example, die-level temperature estimations from HotSpot can be used as a guideline for temperature-aware design during the entire design flow [36].

First, we present some HotSpot simulation results for an imaginary microprocessor design at a future 45nm technology node as a case study. These results demonstrate the importance of using temperature as a guideline during design for high performance systems. Technology specifications used in this case study are shown in Table 8.1, with the second

physical parameters	across die	L1 D-cache
number of transistors	2200 million	70 million
Rent's parameters	$p_r = 0.6$, $k_r = 4.0$	$p_r = 0.6$, $k_r = 4.0$
feature size	45nm	45nm
wiring levels	12	12
area	3.10cm^2	9.56mm^2
power dissipation	218W	60.9W
power density	70.3W/cm^2	637W/cm^2

Table 8.1: A microprocessor example—across-die vs. L1 D-cache (based on ITRS 45nm technology node [2]).

column taken from ITRS data [2]. We use the on-die level-one data cache approximating that of the Alpha 21364 processor scaled to 45nm technology node as an example of localized heating. The scaling process is a linear scaling from known data at 130nm technology. Power consumption values of functional units are extracted from a technology-scaled version of an architecture-level power model [10].

We first show that at the die level, using estimated temperatures from HotSpot compact thermal model offers more accurate design estimations for power, delay and interconnect reliability than just using room temperature or worst-case temperature, as can be seen from Table 8.2. Simply using room temperature or worst-case temperature yields significantly more errors, therefore leading to possible incorrect design decisions and longer design convergence time.

From the same HotSpot thermal model, for the same 45nm microprocessor design, we also show the accuracy of temperature estimations with different grid densities. As we

Table 8.2: Temperature estimates using room temperature and worst-case temperature, normalized to the temperature estimates from the thermal model [36].

	model	room temp.	worst-case temp.
leakage power	1.0	0.61	2.85
intrinsic delay	1.0	0.83	1.25
wire lifetime	1.0	37.40	0.027

can see from Table 8.3, localized heating in a functional block such as L1 D-cache can have a significantly higher temperature than the average die temperature. Even within the L1 D-cache itself, there are also noticeable temperature differences. Therefore, during the design of specific functional blocks, using average die temperature yields inaccurate design estimates. From the last column of Table 8.3, we can also see the influence of the grid density on the precision of maximum L1 D-cache temperature predictions. At some point, further increasing the grid density no longer improve the temperature estimations in L1 D-cache, which confirms the accuracy analysis in Chapter 6.

8.2 Microarchitecture-Level Dynamic Thermal Management

Traditionally, dynamic thermal management (DTM) has been conducted at the chip level. While chip-level DTM techniques can significantly reduce cooling costs and still allowing peak performance for typical applications, they also can substantially reduce performance for applications that exceed the thermal design point. On the other hand,

Table 8.3: Temperatures with different grid densities (°C) [36].

# of grids (die)	die avg. T	D-cache avg. T	D-cache max T
25x25	72.8	115.4	120.5
30x30	72.8	115.4	123.7
35x35	72.8	115.4	126.7
40x40	72.8	115.4	128.1
45x45	72.8	115.4	128.8
50x50	72.8	115.4	129.1
55x55	72.8	115.4	129.2

Microarchitecture-level DTM techniques provide better performance-temperature trade-offs due to their unique ability to use runtime knowledge of different units of the chip, thus achieving finer grained control over the chip's thermal behavior [78]. A compact chip-level HotSpot thermal model at the microarchitecture level proves to be successful in such applications. At design time, the transient and steay-state temperature estimates of all functional units from HotSpot can be fed into a cycle-accurate microprocessor simulator such as SimpleScalar[1] and Turandot [59], where the DTM techniques are implemented and simulated, The cycle-accurate processor simulator in turn provides runtime information of the microprocessor to architecture-level power models such as Wattch [10]. The output power estimations are then input to HotSpot for simulated runtime temperature update, forming a

[1] SimpleScalar Architecture Research Tool Set (http://www.cs.wisc.edu/ mscalar/simplescalar.html).

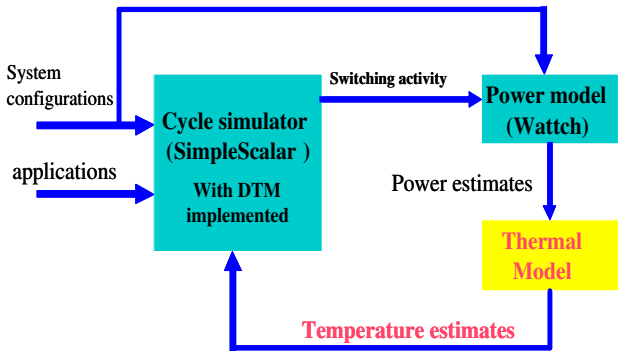

Fig. 8.2: Dynamic thermal management simulation environment.

control loop that makes the operation of the microprocessor temperature-aware. The loop is illustrated in Fig. 8.2.

HotSpot is used in our recent work [76, 77, 78, 74] to simulate and evaluate different DTM techniques. In [78], the concept of "warm-up" of HotSpot simulation is proposed. During the transient thermal simulation, the initial temperatures that are set at the beginning of simulation play a large role in thermal behavior. The most important temperature is that of the heat sink. Its thermal time constant is on the order of several minutes, so its temperature barely changes and certainly does not reach steady-state in short-time simulations. This means simulations must begin with the correct heat-sink temperature, otherwise dramatic errors occur. When we start simulations, we first run the simulations in full-detail cycle-accurate mode for 100 million cycles for microarchitectural "warm-up" to put functional block in reasonable states. We then warm up the temperatures by first setting the blocks' initial temperatures to the steady-state temperatures we calculate using the per-block average power dissipation for each benchmark. This accelerates thermal warm up, but a dynamic warmup phase is still needed because the sample we are at probably does not exhibit average behavior in all the units.

All these works from our group indicate that the performance penalty is improved significantly using micro-architecture-level DTM compared to chip-level techniques. For example, one microarchitecture DTM technique, migrating computation, outperforms global toggling and DVS and many other local DTM techniques for most of the benchmark programs with less than 10% slowdown. HotSpot compact thermal model plays a key role in achieving these temperature-aware microarchitecture design, and it has been also widely adopted in the computer architecture research community, such as [12, 22, 38, 52, 81] and many other on-going works.

8.3 Temperature-Aware Reliability Analysis

Although power density of VLSI circuits increases rapidly due to continued CMOS technology scaling, tolerable operating temperatures usually remain fixed from generation to generation. This is because many aging mechanisms in VLSI circuits proceed at a rate that is temperature dependent, and the requirements of product lifetime dictate the maximum aging rate. For most applications, maximum temperatures are therefore limited to around 100–120° C. Unfortunately, this means that rising power densities impose rising cooling costs that may eventually become so prohibitive that they limit the development of new and reliable products.

Historically, the temperature dependence of many aging processes, such as interconnect electromigration (EM), time-dependent dielectric breakdown (TDDB), and Negative Bias Temperature Instability(NBTI), can be empirically modeled by the Arrhenius Equation

$$\text{MTF} = \text{MTF}_0 \, exp\left(\frac{E_a}{kT}\right) \tag{8.1}$$

where MTF_0 is the mean time to failure at a specified reference temperature, E_a is the activation energy of the failure, and k is the Boltzmann constant.

Detailed steady-state temperature maps across the silicon die or the interconnect layers are crucial to achieve accurate reliability analysis. HotSpot proves to be competent in these applications [55, 56, 80].

Temperature gradients are also important for an accurate reliability analysis and reliability-aware design. In [33, 55, 56], we present a new approach to interconnect and device gate-oxide reliability analysis that accounts for temporal and spatial variations in temperature. We use HotSpot for simulating, at various levels of detail, the time-dependent evolution of on-chip temperatures across an IC. Reliability analysis using temporal and spatial gradient values obtained from a real application on a simulated processor show the importance of accounting for temperature gradients. Worst-case analysis can drastically underestimate expected lifetime, requiring either unnecessarily aggressive and costly cooling solutions or else reductions in power dissipation that incur unnecessary sacrifices in IC performance. Therefore, instead of designing for a maximum tolerated temperature based on a worst-case analysis, expected lifetime should be viewed as a resource that is consumed over time at a temperature-dependent rate. This dynamic, reliability-driven approach to manage operating temperature fits particularly well with the recent advent of dynamic thermal management techniques and shows that lifetime requirements are the proper objective function rather than fixed temperature thresholds.

For example, Fig. 8.3 indicates that the operating temperature can occasionally exceed the reliability equivalent temperature, meaning that the thermal threshold can be set higher than the reliability equivalent temperature, as long as the thermal profile for typical appli-

cations are available. By considering temperature-aware reliability, the slowdown of DTM techniques can be further reduced by up to 50% as reported in [56], hence reclaim the design margin.

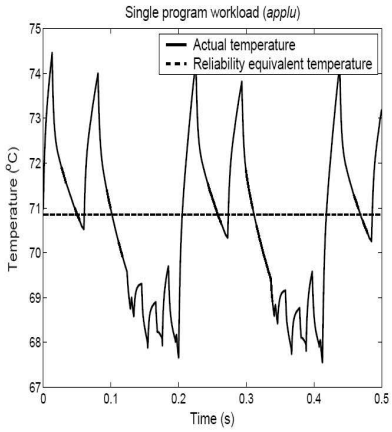

Fig. 8.3: The oscillation of transient temperature indicates that design margin can be reclaimed as long as the average temperature over time is no greater than the constant reliability temperature [56].

8.4 Exploring Thermal Packaging Options

In Chapter 3, we have mentioned that package components can greatly affect the temperature distribution across the silicon die. In this section, we show some example thermal analysis regarding one packaging component—thermal interface material (TIM) that glues the silicon die to the heat spreader. Other package components can be analyzed similarly.

With the flexibility of the compact chip and package thermal model in [36], we can easily investigate the thermal impacts of different TIM properties, such as its thickness, void size, and attaching surface roughness, in early design stages and provide important insights for circuit designers, computer architects and package designers.

We first show how the thickness of TIM affects silicon die temperature distribution. Fig. 8.4 plots the across-die temperature difference from the compact thermal model with different TIM thickness.

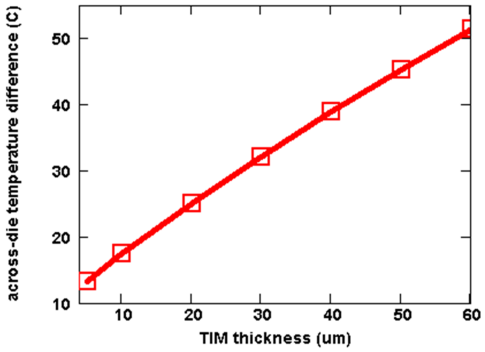

Fig. 8.4: The impact of thermal interface material (TIM) thickness to silicon die temperature difference [32].

As can be observed from Fig. 8.4, thicker TIM results in poor heat spreading which leads to large temperature differences across the die. We can see that thick TIM can lead to very large die temperature difference across the die ($>$50°C). Even with nominal TIM thickness, which is 20μm for this design, the temperature difference across the die is still 24°C. This means that the bottom surface of the die can not be modeled as an isothermal surface. If the TIM is thick enough, the resultant extremely large temperature differences across the die may be disastrous to circuit performance and die/package reliability. Using a better heat sink will only lower the average silicon temperature but will not help to reduce the temperature difference. This analysis suggests that using the thinnest possible TIM is one of the key issues for package designers to consider. On the other hand, with the known TIM thickness that can be best assembled in package with state-of-the-art packaging technology, it is the task of circuit designers and computer architects to design proper circuits and architectures to maintain the temperature difference across die within a manageable

level.

As another example, Fig. 8.5 shows the relationship between the size of TIM void and the hot spot temperature. During the packaging process, it is almost unavoidable to leave voids or air bubbles in the thermal interface material. In the compact thermal model, the void in TIM can be easily modeled by introducing higher vertical TIM thermal resistance to the grid cell where the void resides. Different sizes of the TIM void can be modeled by different sizes of the grid cell. For the simulations of Fig. 8.5, we put the TIM void right under the hottest grid cell, thus modeling the highest possible die temperature in the presence of a void with different sizes. As can be seen from Fig. 8.5, if the hot spot temperature of the design is 95°C, a void or air bubble in the TIM with a size of 0.25mm^2 can make the hot spot temperature drastically higher (290°C), which inevitably leads to thermal runaway of the chip. Therefore, it is desirable to improve the packaging techniques to make the size of the TIM void as small as possible. Package designers usually have the expertise to know typical TIM void sizes for different packaging processes. They can include this information in the thermal model. By doing this, the thermal model is now able to provide possible worst-case temperature regarding TIM void defects. The consequent architecture and circuit design decisions can thus avoid potential thermal hazards caused by the TIM void defects.

Another important thermal interface material property that affects the die temperature is the surface roughness, i.e. non-uniform TIM. In real-life chip packaging process, the bottom surface of the die and the TIM's attaching surface cannot be perfectly smooth. As shown in Fig 8.6, TIM is only attached to the die at the bumps of the TIM surface. This causes ineffective heat conduction and hence higher die temperature comparing to

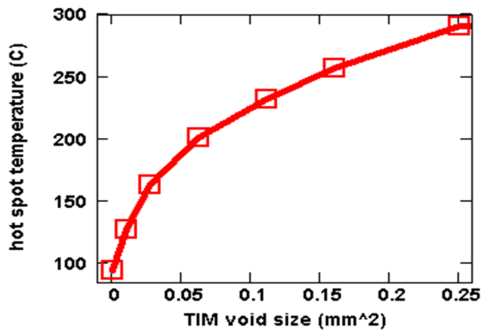

Fig. 8.5: The impact of the size of void defect in thermal interface material (TIM) to silicon die hottest temperature. Temperatures are normalized to the ideal case where there is no void defect in the TIM layer. TIM void sizes are with the unit of mm^2 [32].

the case where TIM and the die attach to each other perfectly. In order to investigate the impact of TIM non-uniformity to the die temperature, we change the thermal model of the TIM layer according to Fig 8.6, where we simply model the non-uniformity of the TIM surface as tiny bumps with spacing $2L$. The size of each grid cell is set to L. Therefore, heat can only be conducted through the grid cells representing the touching bumps. Grid cells representing the valleys are essentially tiny voids that do not touch the die and have extremely low thermal conductivity. The value of L thus can be used as an indicator of the non-uniformity of the TIM surface—the surface is rougher when L is larger and vice versa. Fig. 8.7 is the model results showing the relationship between L (non-uniformity) and die temperatures, where $L = 0$ means the TIM surface is perfectly uniform. As observed, even tiny non-uniform TIM surface (e.g. L=5μm) can significantly raise both the hottest and the average die temperature (by about 10 degrees). Package designers again usually have the specifications of the surface non-uniformities for different packaging processes. Without considering such package processing specifications, it is inevitable that a thermal model underestimates the die temperature and leads to designs that are not thermally optimized and designs with higher probability of premature failures.

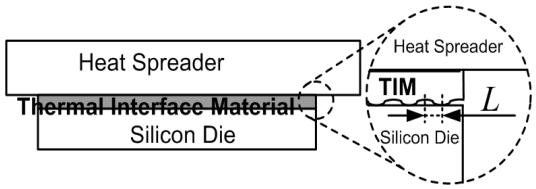

Fig. 8.6: Close-up view of the TIM/die attaching surface. Surface non-uniformity is indicated by L [32].

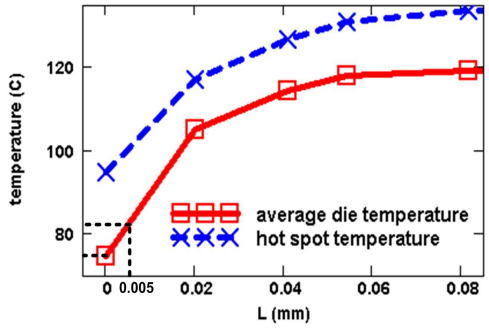

Fig. 8.7: Hottest die temperature and average die temperature vs. the non-uniformity of TIM attaching surface. The larger L is, the rougher the attaching surface. L is defined in Fig. 8.6 [32].

8.5 Thermally Self-Consistent Leakage Power Calculation

In this section, we show as an example that using HotSpot thermal model, one can achieve accurate thermally self-consistent leakage power calculations for a POWER4-like microprocessor design at 130nm technology node. The leakage power calculation flow-chart is shown in Fig. 8.8, which is replicated from Fig. 3.1(a) for convenience. Although similar leakage calculation methods have been described in [84] and [29], the literature still lacks for explicit data showing the impact of leakage power on die temperature distribution. In this section, we provide these data by showing the detailed die temperature maps with and without considering leakage power. We also show how the accuracy of temperature estimations impacts the accuracy of leakage power calculations by using temperature

readings from the thermal model versus using a constant heuristic temperature across the die.

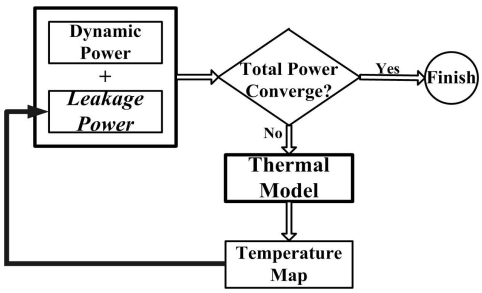

Fig. 8.8: A full-chip thermal model closes the loop for accurate leakage power calculations. The loop is iterated until either power/temperature convergence is achieved or thermal runaway is detected [31].

In this example, the floorplan of a POWER4-like microprocessor is generated by observing the real POWER4 floorplan. It is similar to the one shown in [52], except that the two cores are mirrored to each other, not simply duplicated. The full-chip and package compact thermal model is then constructed based on this floorplan and preliminary package data, which in a real design are from the package designers. For this example design, we use a packaging structure similar to the one in Fig. 4.1(a). The secondary heat transfer path from the die to C4 pads and to PCB is neglected since only a small amount of heat (less than 10%) is transferred through this path and including this path does not significantly change the results. Including the secondary path can be done by adding more layers of materials to the existing model, which will be part of our future work. In the thermal model, the silicon die, thermal interface material and the center part of the heat spreader that is covered by the thermal interface material are all divided into 40×40 grid cells to achieve detailed temperature distributions of these layers. In order to validate the above-mentioned thermal modeling approach, besides the the validation work shown in [36], we

have further quantitatively validated a similar thermal model for a real industrial design with a detailed ANSYS finite-element model simulations and qualitatively validated with on-chip temperature sensor measurements and infra-red temperature images for the same industrial design.

In order to get reasonably accurate initial power estimations of each functional unit for this POWER4-like microprocessor design, we combine IBM's cycle-accurate Turandot performance simulator [59] and PowerTimer power modeling tool [9] running benchmark program *bzip2*. The initial power inputs to the thermal model are the dynamic power values of running *bzip2* for each functional unit. The power numbers are further area-weighted into equivalent heat sources to each of the 40×40 grid cells. Leakage power of each functional unit is initially set to zero. After the thermal model is solved for the first time, leakage power of each unit is updated according to the updated temperature of that unit, which in turn updates the total power of the unit and changes the inputs to the thermal model and thus forms a loop as shown in Fig. 8.8.

The leakage power of grid cell i can be expressed as

$$P_{\text{leakage}_i} = A_i \cdot \alpha \cdot e^{\beta(T_i - T_{base})} \tag{8.2}$$

where A_i and T_i are the area and temperature of the grid cell. α and β are empirical factors that have different values for different technologies (e.g. $\alpha = 1\times10^5$W/m^2 and $\beta = 0.025$ for 130nm). Typical values of α can be found in [81], and typical values of β can be found in [30]. T_{base} is the reference temperature at which α and β are defined.

The loop in Fig. 8.8 is iterated until the operating temperature and the total power converge, or thermal runaway is detected. For this design, convergence is usually achieved within 5 to 7 iterations with zero initial leakage power. On the other hand, if the leakage

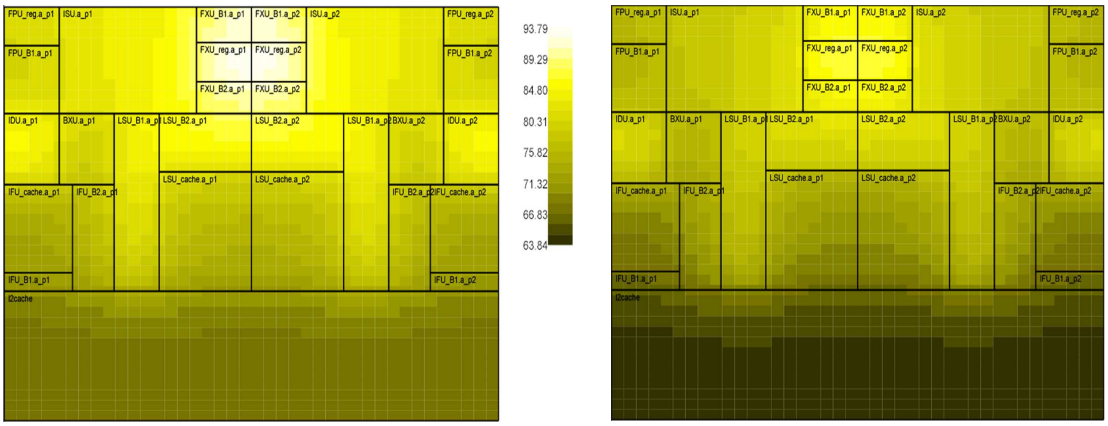

(a) (b)

Fig. 8.9: (a) Temperature map with considerations of the thermally self-consistent calculated leakage power for a benchmark workload on the POWER4-like microprocessor design at 130nm technology node. (b) Temperature map for the same design considering only dynamic power. Especially take note of the two FXU register files at the center of the top part, which are 7 degrees hotter in (a) compared to (b). (All temperatures are in Celsius.) [32]

power is initialized according to room temperature, about 4 iterations are enough to achieve convergence. The computation time for solving a thermal circuit consisting 40×40 nodes is less than one minute on an AMD MP 1.5GHz system during each iteration. The computation time is proportional to the number of grid cells. If needed, further node reduction techniques, such as algebraic multi-grid (AMG) method in [84], can be easily adapted to improve the computation time.

The major advantage of using the method in Fig. 8.8 is that it offers much more accurate leakage power calculations. Fig. 8.9(a) shows the converged temperature map which includes the effects from temperature-dependent subthreshold leakage power as well as dynamic power. For this particular design at 130nm technology node, the leakage power is 17.44% of the total chip power, and the temperature difference across the chip is 23.2°C for

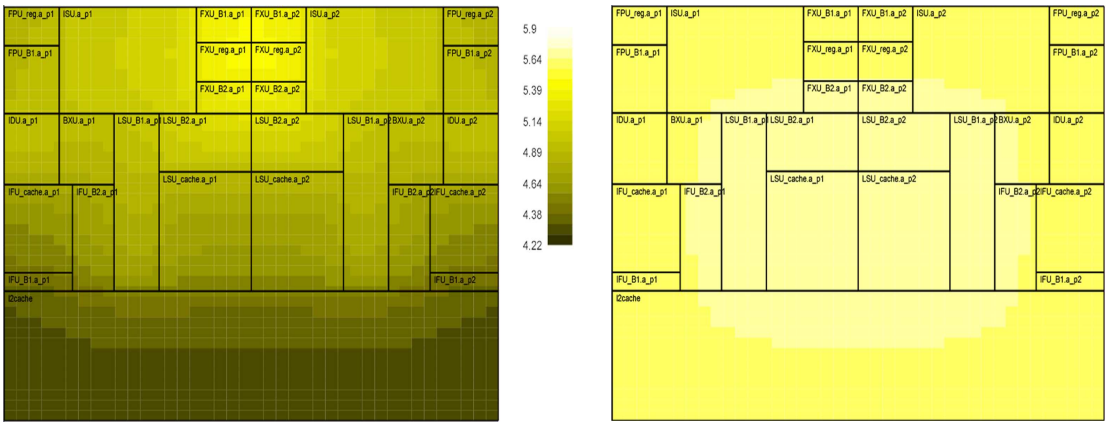

(a) (b)

Fig. 8.10: Imaginary temperature maps with only leakage power applied to the silicon, in order to isolate the thermal effect of leakage power. (a) Temperature map with the thermally self-consistent calculated leakage power applied for a benchmark workload on the POWER4-like microprocessor design. (b) Temperature map with leakage power calculated at constant 85°C across the silicon die. Notice the colored temperature scale in this figure is different from the one in Fig. 4. (Temperatures are in Celsius.) [32]

benchmark *bzip2*. Note that the locations of hot spots can change with different workloads. In comparison, Fig. 8.9(b) shows the temperature map with only dynamic power. Leakage power itself can raise the die temperature by 4 to 7 degrees for this 130nm design.

In addition, in order to decouple the temperature rise caused by the leakage power to that caused by the dynamic power, Fig. 8.10(a) shows the temperature map where only the thermally self-consistent leakage power is applied to the design. For comparison, Fig. 8.10(b) shows the temperature map where only non-temperature-aware leakage power is applied to the design, i.e., in Fig. 8.10(b), leakage power is calculated by assuming constant 85°C across the die. This is still the most common method for leakage power estimation in industry, but it is not accurate. As listed in Table 8.4, using heuristic constant temperature for leakage calculation for this design underestimates leakage power at

the hottest spot by about 15.3%, and overestimates leakage power at the coolest spot by about 51.6%. In total, using heuristic constant temperature overestimates the overall chip leakage power by about 15.7%. (Notice that the colored temperature scale in Fig. 8.10 is not the same as the one in Fig. 8.9.) From these results, it is obvious that considering chip temperature variation with the actual temperature map from a thermal model is required for accurate leakage power calculation for this particular 130nm design, not to mention designs at future sub-100nm technologies. The inaccuracy of using constant heuristic temperature for the across-die leakage power calculations can cause unnecessary packaging cost, hence adding packaging cost (in the case of leakage overestimation); or put the chip in the danger of thermal hazards, hence lowering the final yield (in case of leakage underestimation).

scenarios	leakage and overall powers	
	hottest spot	coolest spot
actual thermal map	(FXU Regfile) leakage:0.144W overall:1.654W	(L2 Cache) leakage:2.406W overall:2.661W
85°C constant temp.	leakage:0.122W overall:1.632W	leakage:3.648W overall:4.263W

Table 8.4: Comparison of the leakage power values calculated using the actual temperature map from the thermal model to that calculated with a heuristic constant 85°C across the die, for both the hottest spot (FXU Regfile) and coolest spot (L2 Cache) on the die. (FXU Regfile is hotter because it has higher *power density* than L2 Cache, although its overall power is less than that of the L2 Cache.) [32]

Using the thermally self-consistent leakage calculation method in Fig. 8.8 can also detect whether the design could possibly run into thermal runaway. One example is that if

the preliminary package design doesn't have enough capability to dissipate all the generated heat, the loop in Fig. 8.8 can turn into a positive feedback loop. In this case, the leakage power and the temperature don't converge, hence thermal runaway occurs. The criterion for the occurrence of thermal runaway is indicated in [29, 53] as

$$\frac{\partial^2 T}{\partial t^2} > 0$$

where T is temperature and t is time.

Table 8.5 shows the results of our investigation of potential thermal runaway for the 130nm process technology used in the example design. As can be seen, thermal runaway can be caused by elevated power dissipation of the design (from 55.74W to 139.35W) with the same thermal package as the example design. It can also be caused by defects in the package, e.g. voids or air bubbles in the thermal interface material, or imperfect attaching surface of the thermal interface material. Defects in the package can equivalently increase the thermal resistance from the die to the ambient. As shown in Table 8.5, an increase in the equivalent package thermal resistance from 0.25°C/W to 0.8°C/W could result in thermal runaway for the design.

	No thermal runaway	thermal runaway
total power	55.74W	139.35W
package R_{th}	0.25°C/W	0.8°C/W

Table 8.5: For the 130nm process of the example design, an increases in total power from 55.74W to 139.35W, or an increase in equivalent junction-to-ambient thermal resistance from 0.25°C/W to 0.8°C/W can make thermal runaway happen [32].

As CMOS processes continue to scale into the sub-100nm regime, both operating temperature and leakage power increase significantly. To make things worse, industry usually

tries aggressive techniques, such as controlling the gate length of a transistor to the lower-end of the gate-length variation (e.g. -1 to -3σ) in order to gain more transistor performance. This in turn exacerbates the leakage power consumption and makes thermal runaway of a design much easier to happen. Therefore, the accurate thermally self-consistent leakage power calculation method shown in Fig. 8.8 becomes more and more important for future technologies.

8.6 A Thermally Optimized Design Flow

From the above thermal analysis examples in Section 8.5 and 8.4, it is obvious that for optimal designs at future technologies, operating temperature needs to be modeled as accurate as possible in early design stages. In order to model temperature more accurately, important aspects of package information should also be included in the model. Ultimately, the full-chip and package thermal model should include all the needed package information (e.g. heat dissipation capability, geometries, materials, potential packaging defects such as the ones in Section 8.4, etc.) for different available package designs that circuit designers and architects can choose from and evaluate. Essentially, this requires more collaborations among circuit designers, computer architects and package designers. A compact thermal model that models detailed temperature distributions for both the silicon and the package can act as a convenient medium for such purpose. Fig. 8.11 illustrates a pre-layout design flow reflecting the collaborations among designers at different design levels. This design flow can detect potential thermal hazards early in the design process and lead to thermally optimized design.

As shown in Fig. 8.11, circuit designers first design basic circuit blocks called macros,

and each macro has a simulated dynamic power for certain workload. It also has an estimated layout bounding box. The macros are then assembled into a preliminary microarchitecture and a floorplan according to work of computer architects. At this stage, initial total power, including rough estimation of leakage power, can be used for a package designer to propose a preliminary package design. All the information of power, floorplan and package are used to construct a compact thermal model which can perform thermally self-consistent leakage power calculations as shown in Section 8.5. The resulting temperature map can then be utilized to perform temperature-critical reliability analysis (e.g. interconnect electromigration, gate-oxide breakdown and package deformation) and temperature-related performance analysis (e.g. interconnect/device delay, power grid IR drop).

The results of all these analysis, together with the total powers, are then compared to the design goals. If the goals are not satisfied, different tradeoffs can be made—circuit designers may need to invent novel circuits with lower power dissipation, computer architects may think more about new architectures and different floorplans to better manage power and temperature, or package designers may need to propose more advanced, usually more expensive, packages. On the other hand, if the design goals are fully satisfied, we still need to check whether the design is too conservative and the design margin is too large for the application. We can then improve the conservative design by either introducing more aggressive circuit and/or architecture solutions to enhance performance, or using simpler and cheaper packages to reduce the cost of final product. These decisions and tradeoffs can then be evaluated by the thermal analysis again following the same flow until an optimal design point is reached. Then one can proceed to the physical design stage.

With the above design flow, the potential thermal hazards can be discovered and dealt

with early and efficiently, thus the design is optimized from a thermal point of view.

8.7 Other Applications

The HotSpot thermal models also have been used to assist the analysis of interactions between temperature and other sources of variations, such as the threshold voltage variation. Temperature exacerbates the power and performance variations caused by the systematic and random variations of V_{th} in the following ways: (1) Higher temperature results in lower V_{th}; (2) The thermally related subthreshold leakage distribution is worse with the presence of local hot spots; (3) Circuit performance variation due to variations V_{th} and degraded carrier mobility cannot be accurately modeled without considering temperature distribution. In [37], HotSpot model is used achieve better estimations for the impacts of V_{th} and other process parameters on power and performance.

The estimated temperature map of a chip under different workloads can also be used as a guideline for on-chip temperature sensor placement. In [49], an analysis of thermal sensor placement and the avoidance of undetected hot spots are addressed with the aid of HotSpot thermal models.

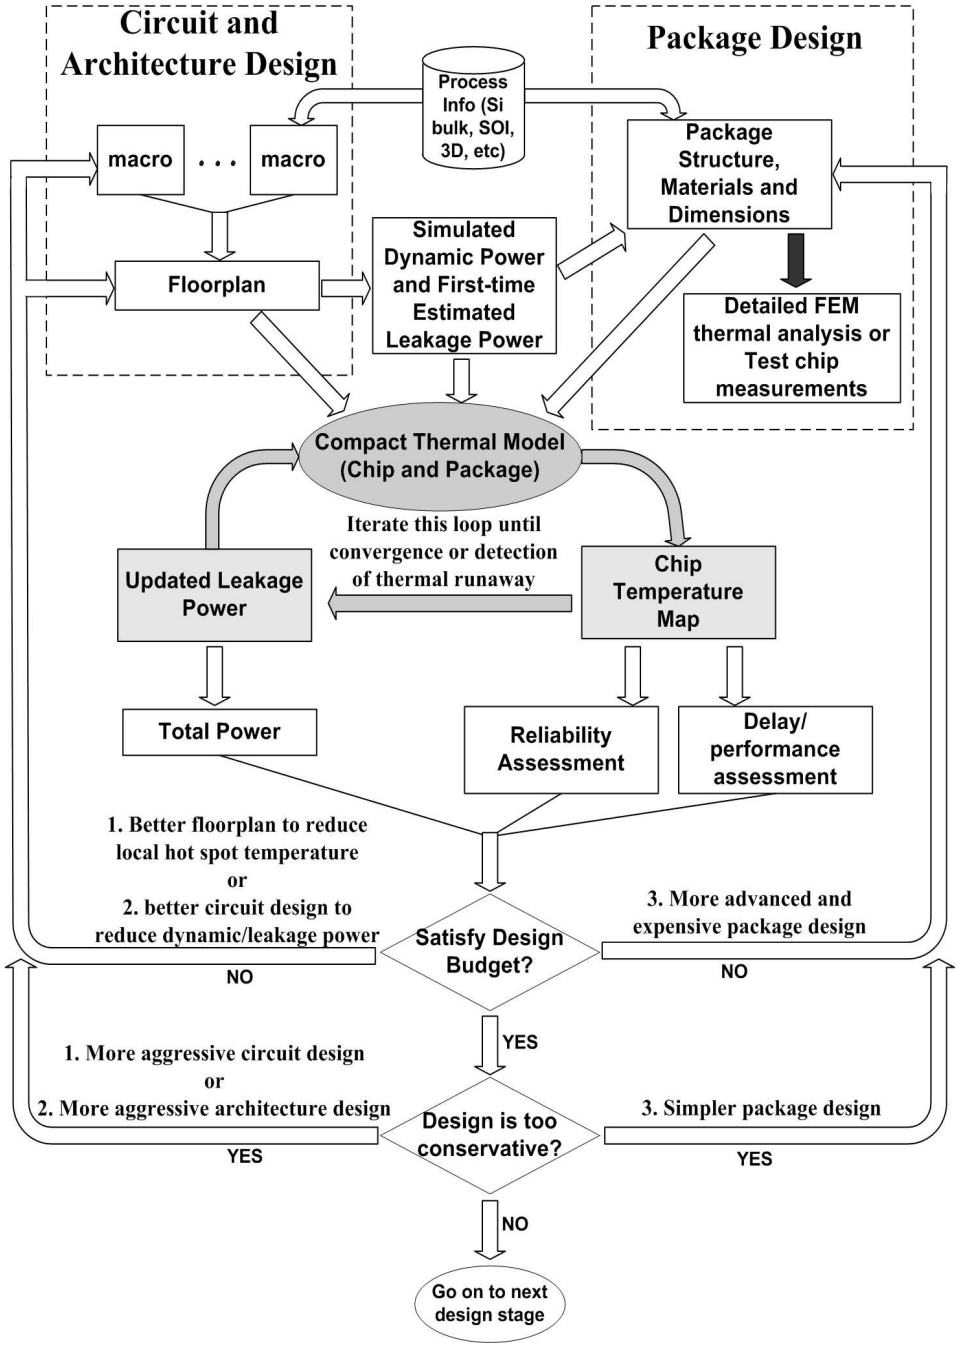

Fig. 8.11: A design flow showing the compact thermal model acts as a convenient medium for productive collaborations for designers at the circuit, architecture and package levels [32].

Chapter 9

Conclusion

In this dissertation, we have presented a novel full-chip and package compact thermal modeling modeling methodology—HotSpot—to deal with the thermal challenges faced by circuit designer, computer architects and package designers in the deep submicron and nanometer era.

HotSpot simplifies the heat diffusion equation with a compact network of thermal resistances and thermal capacitances representing the heat transfer paths through both the silicon die and the package components. HotSpot is an accurate and efficient by-construction thermal modeling approach which is fully parameterized. HotSpot takes into account the heat spreading effect by modeling lateral heat transfer using lateral thermal resistances. It also takes care of the time-domain temperature filtering effect by using thermal capacitances to model the transient responses. The spatial temperature filtering effect is taken care of in HotSpot by rigorous spatial granularity analysis to accurately model the spatial temperature gradient.

9.1 Future Research Directions

The HotSpot thermal modeling method presented in this dissertation has been widely adopted in academia. However, it is far from perfect. There are several additional research directions that extend the existing HotSpot models to make more robust and versatile.

1. HotSpot needs to be able to easily accommodate more types of thermal packages, such as multi-chip module (MCM), advanced liquid cooling structures such as micro-channel liquid cooling, therefore making HotSpot support any type of chip besides high-performance microprocessors.

2. A true system-level thermal modeling method, which includes not only one integrated circuit chip, but also multiple packages and other components such as hard disks, fans, chasses, racks and their environment.

3. Make HotSpot more compatible with existing VLSI design methodologies at different abstract levels. One example is to make HotSpot interact with transaction-level SystemC modeling and other pre-RTL/pre-sysnthesis stages during the design for ASIC and SoC designs.

4. Using hybrid grids for full-chip modeling in HotSpot by taking advantage of the spatial granularity analysis will be an important improvement over the existing regular-grid approach. This will make the HotSpot models more accurate and more compact.

5. The solver for the linear R-C network needs to be improved by advanced numerical methods such as the multi-grid algorithm to further speed up the computation.

6. In addition to the examples shown in this dissertation, there should be many other potential applications to be identified for HotSpot thermal models. One example is to utilize HotSpot results to guide the efficient and accurate placement of on-chip temperature sensors.

Bibliography

[1] http://www.intel.com/research/silicon/mobilepackaging.htm.

[2] The international technology roadmap for semiconductors (ITRS), 2003.

[3] K. Banerjee, S. C. Lin, A. Keshavarzi, and V. De. A self-consistent junction temperature estimation methodology for nanometer scale ICs with implications for performance and thermal management. *Proc. Intl. Elec. Dev. Meeting (IEDM)*, pages 36.7.1–36.7.4, 2003.

[4] K. Banerjee and A. Mehrotra. Global (interconnect) warming. *IEEE Circuits and Devices Magazine*, 17(5):16–32, September 2001.

[5] K. Banerjee, S. J. Souri, P. Kapur, and K. C. Saraswat. 3-D ICs: A novel chip design for improving deep-submicrometer interconnect performance and systems-on-chip integration. *Proceedings of the IEEE*, 89(5):602–633, May 2001.

[6] A. Bar-Cohen, T. Elperin, and R. Eliasi. θ_{jc} charaterization of chip packages – justification, limitations and future. *Components, Hybrids, Manufacturing Technology, IEEE Transactions on*, 12:724–731, December 1989.

[7] S. Borkar. Design challenges of technology scaling. *IEEE Micro*, 19(4):23–29, July/August 1999.

[8] E. G. T. Bosch. Thermal compact models: An alternative approach. *Components and Packaging Technologies, IEEE Transactions on*, 26(1):173–178, March 2003.

[9] D. Brooks, P. Bose, V. Srinivasan, M. Gschwind, P. G. Emma, and M. G. Rosenfield. New methodology for early-stage microarchitecture-level power-performance analysis of microprocessors. *IBM Journal of Research and Development*, 47(5/6):653–670, 2003.

[10] D. Brooks, V. Tiwari, and M. Martonosi. Wattch: A framework for architectural-level power analysis and optimizations. *Proc. Intl. Symp. of Computer Architecture (ISCA)*, pages 83–94, June 2000.

[11] Y. Cao, C. Hu, X. Huang, A. B. Kahng, I. L. Markov, M. Oliver, D. Stroobandt, and D. Sylvester. Improved a priori interconnect predictions and technology extrapolation in the GTX system. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 11(1):3–14, 2003.

[12] P. Chaparro, Jose Gonzalez, and Antonio Gonzalez. Thermal-effective clustered microarchitectures. *Proc. of Proceedings of the First Workshop. of Temperature-Aware Computer Systems (TACS) at ISCA 04*, 2004.

[13] D. Chen, E. Li, E. Rosenbaum, and S. Kang. Interconnect thermal modeling for accurate simulation of circuit timing and reliability. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 19(2):197–205, 2000.

[14] Y. Cheng, P. Raha, C. Teng, E. Rosenbaum, and S. Kang. ILLIADS-T: An electrothermal timing simulator for temperature-sensitive reliability diagnosis of CMOS VLSI chips. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 17(8):668–681, 1998.

[15] T. Y. Chiang, K. Banerjee, and K. Saraswat. Analytical thermal model for multilevel vlsi interconnects incorporating via effect. *IEEE Electron Device Letters*, 23(1):31–33, 2002.

[16] P. Christie and D. Stroobandt. The interpretation and application of rent's rule. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 8(6):639–648, 2000.

[17] J. J. Clement. Reliability analysis for encapsulated interconnect lines under DC and pulsed AC current using a contimuum electromigration transport model. *Journal of Applied Physics*, 83(9):2399–2401, May 2003.

[18] L. Codecasa, D. D'Amore, and P. Maffezzoni. An Arnoli based thermal network reduction method for electro-thermal analysis. *Components and Packaging Technologies, IEEE Transactions on*, 26(1):186–192, March 2003.

[19] P. Dadvar and K. Skadron. Potential thermal security risks. *Proc. of IEEE Semiconductor Thermal Measurement, Modeling, and Management Symposium (SemiTherm)*, pages 229–234, 2005.

[20] J. Dambre, P. Verplaetse, D. Stroobandt, and J. V. Campenhout. A comparison of various terminal-gate relationships for interconnect prediction in vlsi circuits. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 11(1):24–34, 2003.

[21] J. A. Davis, V. K. De, and J. D. Meindl. A stochastic wire-length distribution for gigascale integration (GSI)—part I: Derivation and validation. *Electron Devices, IEEE Transactions on*, 45(3):580–589, March 1998.

[22] J. Donald and M. Martonosi. Temperature-aware design issues for SMT and CMP architectures. *Proc. of Workshop on Complexity-Effective Design*, 2004.

[23] W. E. Donath. Wire-length distribution for placements of computer logic. *IBM Journal of Research and Development*, 2(3):152–155, May 1981.

[24] K. Etessam-Yazdani, H. F. Hamann, and M. Asheghi. Impact of power granularity on chip thermal modeling. *Proc. of 10th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronics Systems (ITHERM)*, pages 666–670, June 2006.

[25] M. Furmanczyk, A. Napieralski, K. Szaniawski, W. Tylman, and A. Lara. Reduced electro-thermal models for integrated circuits. *Proc. 1998 Intl. Conf. on Modeling and Simulation of Microsystems, Semiconductors, Sensors and Actuators*, pages 139–144, April 1998.

[26] S. Gunther, F. Binns, D. M. Carmean, and J. C. Hall. Managing the impact of increasing microprocessor power consumption. *Intel Tchnology Journal*, Q1 2001.

[27] K. D. Hagen. *Heat Transfer with Applications*. Prentice-Hall, Inc., 1999.

[28] L. He, W. Liao, and M. R. Stan. System level leakage reduction considering the interdependence of temperature and leakage. *Proc. of 41st Design Automation Conference (DAC)*, pages 12–17, June 2004.

[29] L. He, W. Liao, and M. R. Stan. System level leakage reduction considering the interdependence of temperature and leakage. *Proc. 41st Design Automation Conference (DAC)*, pages 12–17, June 2004.

[30] S. Heo, K. Barr, and K. Asanovic. Reducing power density through activity migration. *Proc. of Intl. Symp. on Low Power Electronics and Design (ISLPED'03)*, pages 217–222, August 2003.

[31] W. Huang, S. Ghosh, S. Velusamy, K. Sankaranarayanan, K. Skadron, and M. R. Stan. Hotspot: A compact thermal modeling methodology for early-stage vlsi design. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 14(5):501–513, May 2006.

[32] W. Huang, E. Humenay, K. Skadron, and M. Stan. The need for a full-chip and package thermal model for thermally optimized IC designs. *Proc. Intl. Symp. on Low Power Electronic Design (ISLPED)*, pages 245–250, August 2005.

[33] W. Huang, Z. Lu, S. Ghosh, J. Lach, M. Stan, and S. Skadron. The importance of temporal and spatial temperature gradients in ic reliability analysis. Technical Report CS-2004-07, University of Virginia, Computer Science Department, 2004.

[34] W. Huang, M. R. Stan, and K. Skadron. Physically-based compact thermal modeling—achieving parametrization and boundary condition independence. *Proc. of 10th Int'l Workshop on THERMal INvestigations of ICs and Systems*, October 2004.

[35] W. Huang, M. R. Stan, and K. Skadron. Parameterized physical compact ther-

mal modeling. *Components and Packaging Technologies, IEEE Transactions on*, 28(4):615–622, December 2005.

[36] W. Huang, M. R. Stan, K. Skadron, K. Sankaranarayanan, S. Ghosh, and S.Velusamy. Compact thermal modeling for temperature-aware design. *Proc. 41st Design Automation Conference (DAC)*, pages 878–883, June 2004.

[37] E. Humenay, W. Huang, M. Stan, and S. Skadron. Toward an architectural treatment of parameter variations. Technical Report CS-2005-16, University of Virginia, Computer Science Department, 2005.

[38] W. Hung, C. Addo-Quaye, T. Theocharides, Y. Xie, N. Vijakrishnan, and M. J. Irwin. Thermal-aware IP vvrtualization and placement for networks-on-chip architecture. *Proc. of IEEE International Conference on Computer Design (ICCD)*, pages 430–437, October 2004.

[39] S. Im, Srivastava N, K. Banerjee, and K. E. Goodson. Scaling analysis of multilevel interconnect temperatures for high-performance ICs. *Electron Devices, IEEE Transactions on*, 52(12):2710–2719, December 2005.

[40] T. Kam, S. Rawat, D. Kirkpatrick, R. Roy, G. S. Spirakis, N. Sherwani, and C. Peterson. EDA challenges facing future microprocessor design. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 19(12):1498–1506, December 2000.

[41] A. Kaveh. The history of power dissipation. *Electronics Cooling Magazine*, January 2000.

[42] C. J. M. Lasance. Thermal characterization of electronic parts with compact models: Interpretation, application, and the need for a paradigm shift. *Proc. 13th IEEE SEMI-THERM Symposium*, pages 40–48, 1997.

[43] C. J. M. Lasance. Two benchmarks to facilitate the study of compact thermal modeling phenomena. *Components and Packaging Technologies, IEEE Transactions on*, 24(4):559–565, December 2001.

[44] C. J. M. Lasance. Recent progress in compact thermal models. *Proc. 19th IEEE SEMI-THERM Symp.*, pages 290–299, March 2003.

[45] C. J. M. Lasance. The urgent need for widely-accepted test methods for thermal interface materials. *Proc. 19th IEEE SEMI-THERM Symposium*, pages 123–128, 2003.

[46] C. J. M. Lasance. The influence of various common assumptions on the boundary-condition-independence of compact thermal models. *Components, Packaging, and Manufacturing Technology–Part A, IEEE Transactions on*, 27(3):523–529, September 2004.

[47] C. J. M. Lasance, H. I. Rosten, and J. D. Parry. The world of thermal charaterization according to DELPHI—part II: Experimental and numerical methods. *Components and Packaging Technologies, IEEE Transactions on*, 20(4):392–398, December 1997.

[48] C. J. M. Lasance, H. Vinke, H. Rosten, and K. L. Weiner. A novel approach for the thermal charaterization of electronic parts. *Proc. 11th IEEE SEMI-THERM Symposium*, page 1, 1995.

[49] K. J. Lee, K. Skadron, and W. Huang. Analytical model for sensor placement on microprocessors. *Proc. of IEEE International Conference on Computer Design (ICCD)*, pages 24–27, October 2005.

[50] S. Lee, S. Song, V. Au, and K. Moran. Constricting/spreading resistance model for electronics packaging. *Proc. American-Japan Thermal Engineering Conf. (AJTEC)*, pages 199–206, March 1995.

[51] P. Li, L. Pileggi, M. Asheghi, and R. Chandra. Efficient full-chip thermal modeling and analysis. *Proc. Intl. Conf. on Computer-Aided Design (ICCAD)*, 2004.

[52] Y. Li, D. Brooks, Z. Hu, and K. Skadron. Performance, energy, and thermal considerations for SMT and CMP architectures. *Proc. High Performance Computer Architecture (HPCA)*, pages 71–82, February 2005.

[53] W. Liao and L. He. Coupled power and thermal simulation and its application. *3rd Workshop on Power-Aware Computer Systems*, December 2003.

[54] S. Lopez-Buedo and E. Boemo. Making visible the thermal behavior of embedded microprocessors on FPGAs. a progress report. *Proc. FPGA*, pages 79–86, February 2004.

[55] Z. Lu, W. Huang, J. C. Lach, M. R. Stan, and K. Skadron. Interconnect lifetime prediction under dynamic stress for reliability-aware design. *Proc. Intl. Conf. on Computer-Aided Design (ICCAD)*, pages 327–334, November 2004.

[56] Z. Lu, W. Huang, M. Stan, K. Skadron, and J. Lach. Interconnect lifetime prediction with temporal and spatial temperature gradients for reliability-aware design and run-

time management: Modeling and applications. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, in press.

[57] R. Mahajan, C-P. Chiu, and G. Chrysler. Cooling a microprocessor chip. *Proceedings of the IEEE*, 94(8):1476–1486, August 2006.

[58] G. E. Moore. Progress in digital integrated electronics. *Proc. of International Electron Devices Meeting*, 1975.

[59] M. Moudgill, J. D. Wellman, and J. H. Moreno. Environment for PowerPC microarchitecture exploration. *IEEE Micro*, 19(3):15–25, 1999.

[60] S. Naffziger, B. Stackhouse, T. Grutkowski, D. Josephson, J. Desai, E. Alon, and M. Horowitz. The implementation of a 2-core, multi-threaded itanium family processor. *Solid-State Circuits, IEEE Journal of*, 41(1):197–209, January 2006.

[61] R. H. J. M. Otten and R. K. Brayton. Planning for performance. *Proc. Design Automation Conference (DAC)*, pages 122–127, June 1998.

[62] H. Pappe, D. Schweitzer, J. H. Janssen, A. Morelli, and C. M. Villa. Thermal transient modeling and experimental validation in the european project PROFIT. *Components and Packaging Technologies, IEEE Transactions on*, 27(3):530–538, September 2004.

[63] J. Parry, H. Rosten, and G. B. Kromann. The development of component-level thermal compact models of a C4/CBGA interconnect technology: The motorola PowerPC 603 and PowerPC 604 RISC microproceesors. *Components, Packaging, and Manufacturing Technology–Part A, IEEE Transactions on*, 21(1):104–112, March 1998.

[64] C. Poirier, R. McGowen, C. Bostak, and S. Naffziger. Power and temperature control on a 90nm itanium-family processor. *Proc. of IEEE International Symposium on Solid-State Circuits (ISSCC)*, pages 304–305, February 2005.

[65] E. Pop, K. Banerjee, P. Sverdrup, R. Dutton, and K. Goodson. Localized heating effects and scaling of sub-0.18 micron CMOS devices. *Proc. Intl. Elec. Devices Meeting (IEDM)*, pages 677–680, 2001.

[66] M. Rencz, G. Farkas, V. Székely, A. Poppe, and B. Courtois. Boundary condition independent dynamic compact models of packages and heat sinks from thermal transient measurements. *Proc. Electronics Packaging Technology Conference*, pages 479–484, 2003.

[67] H. Rosten and C. Lasance. Delphi: The development of libraries of physical models of electronic components for an integrated design environment. *Proc. Conf. Int. Elec. Pack. Soc.*, 1994.

[68] H. I. Rosten, C. J. M. Lasance, and J. D. Parry. The world of thermal charaterization according to DELPHI—part I: Background to DELPHI. *Components and Packaging Technologies, IEEE Transactions on*, 20(4):384–391, December 1997.

[69] S. Rzepka, K. Banerjee, E. Meusel, and C. Hu. Characterization of self-heating in advanced VLSI interconnect lines based on thermal finite element simulation. *Components, Packaging, and Manufacturing Technology–Part A, IEEE Transactions on*, 21(3):406–411, September 1998.

[70] M-N. Sabry. Compact thermal models for electronic systems. *Components and Packaging Technologies, IEEE Transactions on*, 26(1):179–185, March 2003.

[71] B. Shapiro. Creating compact models of complex electronic systems: An overview and suggested use of existing model reduction and experimental system identification tools. *Components and Packaging Technologies, IEEE Transactions on*, 26(1):165–172, March 2003.

[72] W. Shockley, J. Bardeen, and W. Brattain. The first transistor, December 1947.

[73] K. Skadron, T. Abdelzaher, and M. R. Stan. Control-theoretic techniques and thermal-rc modeling for accurate and localized dynamic thermal management. *Proc. High-Performance Computer Architecture (HPCA)*, pages 17–28, 2002.

[74] K. Skadron, K. Sankaranarayanan, S. Velusamy, D. Tarjan, M. R. Stan, and W. Huang. Temperature-aware microarchitecture: Modeling and implementation. *ACM Transactions on Architecture and Code Optimization*, 1(1):94–125, March 2004.

[75] K. Skadron, M. R. Stan, and *et al*. Hotspot: Techniques for modeling thermal effects at the processor-architecture level. *Proc. 8th THERMINIC*, pages 169–172, October 2002.

[76] K. Skadron, M. R. Stan, W. Huang, K. Sankaranarayanan, Z. Lu, and J. Lach. Monitoring temperature in fpga based socs. *Proc. of International Conference on Thermal and Mechanical Simulation and Experiments in Microelectronics and Microsystems, EuroSimE.*, pages 415–422, 2004.

[77] K Skadron, M. R. Stan, W. Huang, S. Velusamy, K. Sankaranarayanan, and D. Tarjan.

Temperature-aware computer systems: Opportunities and challenges. *IEEE Micro*, 23(6):52–61, 2003.

[78] K. Skadron, M. R. Stan, W. Huang, S. Velusamy, K. Sankaranarayanan, and D. Tarjan. Temperature-aware microarchitecture. *Proc. Intl. Symp. on Computer Architecture (ISCA)*, pages 2–13, June 2003.

[79] K. R. Spangenberg. *Vacuum Tubes*. McGraw-Hill, 1948.

[80] J. Srinivasan, S. V. Adve, P. Bose, and J. A. Rivers. The case for lifetime reliability-aware microprocessors. *Proc. of The Intl. Symp. on Computer Architecture (ISCA'04)*, June 2004.

[81] J. Srinivasan, S. V. Adve, P. Bose, and J. A. Rivers. The impact of technology scaling on lifetime reliability. *Proc. of The Intl. Conf. on Dependable Systems and Networks (DSN-04)*, June 2004.

[82] M. R. Stan and K. Skadron. Power-aware computing. *IEEE Computer*, 36(12):35–38, December 2003.

[83] M. R. Stan, K. Skadron, M. Barcella, W. Huang, K. Sankaranarayanan, and S. Velusamy. Hotspot: a dynamic compact thermal model at the processor-architecture level. *Microelectronics Journal*, 34:1153–1165, 2003.

[84] H. Su, F. Liu, A. Devgan, E. Acar, and S. Nassif. Full chip estimation considering power supply and temperature variations. *Proc. Intl. Symp. of Low Power Elec. Design (ISLPED)*, pages 78–83, August 2003.

[85] V. Székely, C. Márta, M. Rencz, G. Végh, Z. Benedek, and S. Török. A thermal benchmark chip: Design and applications. *Components, Packaging, and Manufacturing Technology–Part A, IEEE Transactions on*, 21(3):399–405, September 1998.

[86] K. Torki and F. Ciontu. IC thermal map from digital and thermal simulations. *Proceedings of the 2002 International Workshop on THERMal Investigations of ICs and Systems (THERMINIC)*, pages 303–08, Oct. 2002.

[87] S. R. Turns. *Thermodynamics: Concepts and Applications*. Cambridge University Press, 2006.

[88] N. Vasseghi, K. Yeager, and et al. 200-mhz superscalar risc microprocessor. *Solid State Circuits, IEEE Journal of*, 31(11):1675–1686, Novemebr 1996.

[89] S. Velusamy, W. Huang, J. Lach, and K. Skadron. Monitoring temperature in FPGA-based SoCs. Technical Report CS-2004-39, University of Virginia, Computer Science Department, 2004.

[90] S. Velusamy, W. Huang, J. Lach, M. Stan, and K. Skadron. Monitoring temperature in FPGA based SoCs. *Proc. of IEEE International Conference on Computer Design (ICCD)*, pages 634–637, October 2005.

[91] H. Vinke and C. J. M. Lasance. Compact models for accurate thermal charaterization of electronic parts. *Components, Packaging, and Manufacturing Technology–Part A, IEEE Transactions on*, 20(4):411–419, December 1997.

[92] T-Y. Wang and C. C-P. Chen. 3-D thermal-ADI: A linear-time chip level transient

thermal simulator. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 21(12):1434–1445, December 2002.

[93] T.-Y. Wang, J.-L. Tsai, and C. C.-P. Chen. Thermal and power integrity based power/ground networks optimization. *Proc. Design Automation and Testing of Europe (DATE)*, pages 830–835, February 2004.

[94] P. Zarkesh-Ha, J. Davis, and J. Meindl. Prediction of net-length distribution for global interconnects in a heterogeneous system-on-a-chip. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 8(6):649–659, 2000.

[95] P. Zarkesh-Ha, J. A. Davis, W. Loh, and J. D. Meindl. On a pin versus gate relationship for heterogeneous systems: Heterogeneous rent's rule. *Proc. IEEE Custom Integrated Circuit Conf. (CICC)*, pages 93–96, May 1998.

[96] P. Zarkesh-Ha and J. D. Meindl. Stochastic net-length distribution for global interconnects in a heterogeneous system-on-a-chip. *Proc. IEEE VLSI Symposium on Technology.*, pages 44–45, June 1998.