

Temperature-Aware Design and Management for 3D Multi-Core Architectures

Mohamed M. Sabry and David Atienza
Embedded Systems Lab (ESL)
Ecole Polytechnique Fédérale de Lausanne (EPFL)
1015, Lausanne, Switzerland
mohamed.sabry@epfl.ch, david.atienza@epfl.ch

now

the essence of knowledge

Boston — Delft

Foundations and Trends[®] in Electronic Design Automation

Published, sold and distributed by:

now Publishers Inc.
PO Box 1024
Hanover, MA 02339
United States
Tel. +1-781-985-4510
www.nowpublishers.com
sales@nowpublishers.com

Outside North America:

now Publishers Inc.
PO Box 179
2600 AD Delft
The Netherlands
Tel. +31-6-51115274

The preferred citation for this publication is

M. M. Sabry and D. Atienza. *Temperature-Aware Design and Management for 3D Multi-Core Architectures*. Foundations and Trends[®] in Electronic Design Automation, vol. 8, no. 2, pp. 117–197, 2014.

This Foundations and Trends[®] issue was typeset in L^AT_EX using a class file designed by Neal Parikh. Printed on acid-free paper.

ISBN: 978-1-60198-775-4

© 2014 M. M. Sabry and D. Atienza

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, mechanical, photocopying, recording or otherwise, without prior written permission of the publishers.

Photocopying. In the USA: This journal is registered at the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923. Authorization to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by now Publishers Inc for users registered with the Copyright Clearance Center (CCC). The 'services' for users can be found on the internet at: www.copyright.com

For those organizations that have been granted a photocopy license, a separate system of payment has been arranged. Authorization does not extend to other kinds of copying, such as that for general distribution, for advertising or promotional purposes, for creating new collective works, or for resale. In the rest of the world: Permission to photocopy must be obtained from the copyright owner. Please apply to now Publishers Inc., PO Box 1024, Hanover, MA 02339, USA; Tel. +1 781 871 0245; www.nowpublishers.com; sales@nowpublishers.com

now Publishers Inc. has an exclusive license to publish this material worldwide. Permission to use this content must be obtained from the copyright license holder. Please apply to now Publishers, PO Box 179, 2600 AD Delft, The Netherlands, www.nowpublishers.com; e-mail: sales@nowpublishers.com

**Foundations and Trends[®] in
Electronic Design Automation**
Volume 8, Issue 2, 2014
Editorial Board

Editor-in-Chief

Radu Marculescu
Carnegie Mellon University
United States

Editors

Robert K. Brayton
UC Berkeley

Raul Camposano
Nimbic

K.T. Tim Cheng
UC Santa Barbara

Jason Cong
UCLA

Masahiro Fujita
University of Tokyo

Georges Gielen
KU Leuven

Tom Henzinger
*Institute of Science and Technology
Austria*

Andrew Kahng
UC San Diego

Andreas Kuehlmann
Coverity

Sharad Malik
Princeton University

Ralph Otten
TU Eindhoven

Joel Phillips
Cadence Berkeley Labs

Jonathan Rose
University of Toronto

Rob Rutenbar
*University of Illinois
at Urbana-Champaign*

Alberto Sangiovanni-Vincentelli
UC Berkeley

Leon Stok
IBM Research

Editorial Scope

Topics

Foundations and Trends[®] in Electronic Design Automation publishes survey and tutorial articles in the following topics:

- System level design
- Behavioral synthesis
- Logic design
- Verification
- Test
- Physical design
- Circuit level design
- Reconfigurable systems
- Analog design
- Embedded software and parallel programming
- Multicore, GPU, FPGA, and heterogeneous systems
- Distributed, networked embedded systems
- Real-time and cyberphysical systems

Information for Librarians

Foundations and Trends[®] in Electronic Design Automation, 2014, Volume 8, 4 issues. ISSN paper version 1551-3939. ISSN online version 1551-3947. Also available as a combined paper and online subscription.

Foundations and Trends[®] in Electronic Design
Automation
Vol. 8, No. 2 (2014) 117–197
© 2014 M. M. Sabry and D. Atienza
DOI: 10.1561/10000000032



Temperature-Aware Design and Management for 3D Multi-Core Architectures

Mohamed M. Sabry and David Atienza
Embedded Systems Lab (ESL)
Ecole Polytechnique Fédérale de Lausanne (EPFL)
1015, Lausanne, Switzerland
mohamed.sabry@epfl.ch, david.atienza@epfl.ch

Contents

1	Introduction	2
1.1	3D-ICs for Augmented Performance Per Unit Area	4
1.2	Thermal Issues in 3D MPSoCs	6
1.3	Thermal Impact on 3D MPSoC Reliability and Performance	8
1.4	Advanced Cooling Technologies for 3D MPSoCs	12
1.5	Survey Overview and Outline	15
2	Thermal Modeling for 3D MPSoCs	18
2.1	Heat Transfer Principles	18
2.2	Thermal Modeling Frameworks for 2D/3D ICs	21
2.3	Compact Thermal Modeling for 3D MPSoCs	22
2.4	Summary	26
3	Temperature-Aware Design Optimizations for 3D MPSoCs	28
3.1	On-Chip Design-Time Optimizations	30
3.2	Off-Chip Design-Time Optimizations	36
3.3	Summary	46
4	Temperature-Aware Runtime Management for 3D MPSoCs	47
4.1	Temperature-Affecting Control Knobs	49
4.2	Centralized Versus Decentralized Management Schemes . .	50
4.3	Reactive Versus Proactive Management Schemes	55

	iii
4.4 Heuristic Versus Algorithmic Management Schemes	60
4.5 The Global Perspective on Thermal Management	66
4.6 Summary	68
Conclusion	69
Acknowledgements	72
References	73

Abstract

Vertically-integrated 3D multiprocessors systems-on-chip (3D MP-SoCs) provide the means to continue integrating more functionality within a unit area while enhancing manufacturing yields and runtime performance. However, 3D MPSoCs incur amplified thermal challenges that undermine the corresponding reliability. To address these issues, several advanced cooling technologies, alongside temperature-aware design-time optimizations and run-time management schemes have been proposed. In this monograph, we provide an overall survey on the recent advances in temperature-aware 3D MPSoC considerations. We explore the recent advanced cooling strategies, thermal modeling frameworks, design-time optimizations and run-time thermal management schemes that are primarily targeted for 3D MPSoCs. Our aim of proposing this survey is to provide a global perspective, highlighting the advancements and drawbacks on the recent state-of-the-art.

Keywords: System-Level Design, Thermal Management, MPSoC Cooling, Temperature Optimization, Reliability, Vertical Integration.

1

Introduction

The last decades have seen a revolution in data gathering, processing, information storage and communication. This revolution has been caused by electronic computing systems, which nowadays are one of the key building blocks of the world's information technology (IT) infrastructure. In fact, computing systems and IT services are an essential pillar of the developed world, contributing up to 50% of its economy [1]. The IT and computing systems revolutions have been the result of the advancements in IC processing technology, where the number of components (transistors) on the same die area have been doubled every 18 months [2], which is also known as Moore's law. This has been the drive to generate more complicated computing systems with higher performance and computational functionality.

As feature sizes scale with advanced processing technologies, the performance of processing units has increased because of greater functionality and higher computational capabilities. This functionality augmentation was accompanied by an increase in the operating frequency of the processing unit. Micro-architects have conventionally used operating frequency as a measure for the processing unit performance, as higher frequency implies more instructions executed per unit time.




Figure 1.1: Number of processing units integrated in a single IC, as evolved with time. Blue-labeled ICs are single-core architectures, while red-labeled ICs are multi-core architectures.

However, at the sub-micron level, the circuit-level delays are getting dominated by wiring and interconnect delays, which led to frequency flattening. If the operating frequency would increase with this technology, significant additional power consumption is required by the processing unit, which results in an increased heat generation. For example, a 90 nm fabricated AMD processing unit would require 60% additional power consumption to increase the operating frequency by 400MHz [3].

Multi-core architectures have been proposed as an alternative design paradigm to frequency increase in single-core architectures, to continue performance improvement with technology scaling [4]. Multi-core architectures integrate two or more processing units, with shared or distributed memory modules, interconnected through an on-chip bus or a network-on-chip [5]. As feature sizes scale with advanced processing technologies, the number of processing units in multi-core architectures dramatically increases. Fig. 1.1¹ shows that the number of cores inte-

¹This figure is based on a similar figure found in the course slides given by Prof. S. Amarasinghe of MIT <http://groups.csail.mit.edu/cag/ps3/pdf/6.189-info-session.pdf>

grated in a single IC has started to ramp-up in the beginning of the 21st century. Recent multi-core architectures integrate a number of processing units, multi-level memory hierarchy, an interconnect module, and in the case of embedded domain, special peripherals such as analog-to-digital converters (ADC), co-processors, or wireless RF antennas. This architecture, which is known as *Multiprocessor Systems-on-Chip* (MPSoC), has been widely used in various domains. An example of recent MPSoC utilization at the high performance computing systems level is data-intensive computing systems, which is also known as the *fourth paradigm* [6]. Another example for MPSoC utilization is found at the embedded systems domain, with small- or tiny-size computing systems, such as on-body [7] or in-body [8] health monitoring systems.

1.1 3D-ICs for Augmented Performance Per Unit Area

While the performance of computing architectures has been enhanced by MPSoCs, MPSoCs' performance has been recently challenged by increased propagation delay, primarily due to longer interconnects [9]. This is mainly due to the increased wire-to-gate delay ratio. This delay would lead to degraded MPSoC performance or increased energy consumption.

This delay limitation, combined with the continued demands for increased integrated functionality while preserving the performance and area efficiency, have led to the development of vertically-integrated 3D ICs. 3D integration is viewed as an attractive solution to provide increased functionality with better yield, as well as a technique of combining several technologies in a single enclosure (package) [10]. From a design perspective, 3D integration can be split into the following categories:

- **Monolithic 3D integration** [11]. This integration technology fabricates the tiers serially at the transistor granularity, within a single fabrication process. From the bottom tier, the corresponding transistors are fabricated then a substrate layer is placed on top where another tier is fabricated. These layers are connected with vertical interlayer vias. Thus, this integration technology is

promising in terms of providing higher density and performance gains.

- **3D stacking** (also referred to as Parallel 3D integration) [10]. This integration technology stacks vertically 2D die layers to form a single 3D IC. To enable the communication and power delivery to these dies, there are several techniques adopted such as wire bonding, microbumps, and *through-silicon vias* (TSVs). TSVs are vertical wires that carry power and signals between different dies, which are etched in the silicon substrate between the 2D dies.

In this survey, we primarily target 3D stacked ICs with several digital logic dies and TSV-based interlayer communication. We refer to the targeted 3D ICs throughout this review by 3D multiprocessors systems-on-chip, or 3D MPSoCs.

3D MPSoCs are multi-layered stacked 3D ICs, where each die contains a number of processing units, memory modules, and other peripheral and interconnection units. Examples of typical 3D MPSoCs integrate a number of processing layers that contain all the processing units, and a number of memory layers. Another 3D MPSoC example is where the processing units and the memory modules are co-placed in each die of the 3D stacked layers. These examples have been shown in previous works [12, 13], and are shown in Fig. 1.2, which includes an UltraSPARC T1 [14] version of a 3D MPSoC. This vertical integration of logical modules brings several benefits to multi-core architectures, which are as follows:

- Vertical stacking shortens the wiring length between two modules. In this respect, the propagation delays, which are recently dominated by interconnect delays [15], are dramatically reduced leading to an increased performance of 3D MPSoCs. Thus, 3D MPSoCs would outperform 2D MPSoCs.
- 3D MPSoCs allow heterogeneous integration of different components, such as DRAM on multi-core architectures [13]. 3D MPSoCs enhance the memory access bandwidth and throughput, by bringing the memory modules (e.g. DRAM) to the top or bottom

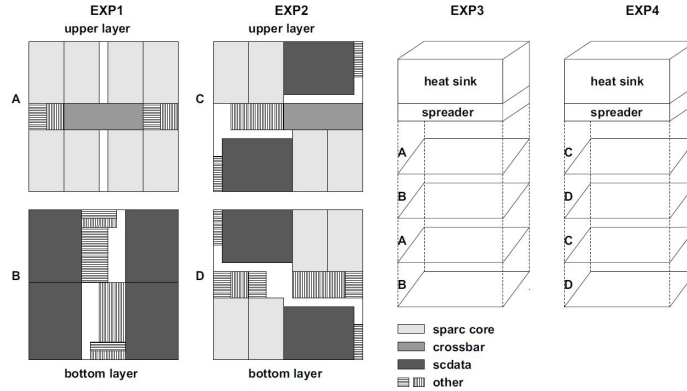


Figure 1.2: Schematic view of a 4-tier 3D MPSoC with different architectures [12].

of processing layers. This would enable high-speed, massively-parallel data access to these stacked DRAM layers.

1.2 Thermal Issues in 3D MPSoCs

Despite the performance and throughput enhancements that 3D MP-SoCs bring, 3D MPSoC designs face major challenges, particularly in the extreme elevated temperatures accompanied with high-performance designs.

While technology continues to scale-down the transistor features, MPSoCs voltage supply (V_{dd}) could not be scaled down accordingly [16]. Recent work [17] (Fig. 1.3(a)) has shown that the supply voltage scaling is saturating that, if combined with increased integration within a unit area due to reduced transistor features, leads to an increase in power consumption. In this respect, multi-core architectures design trends have taken the direction of increasing the power density by integrating more processing units on the chip (with a fixed chip area), as shown in Fig. 1.3(b). If the MPSoC power density keeps increasing, it will eventually reach the same magnitude of nuclear power plants [4, 18].

With such increased power densities, MPSoCs face a tremendous

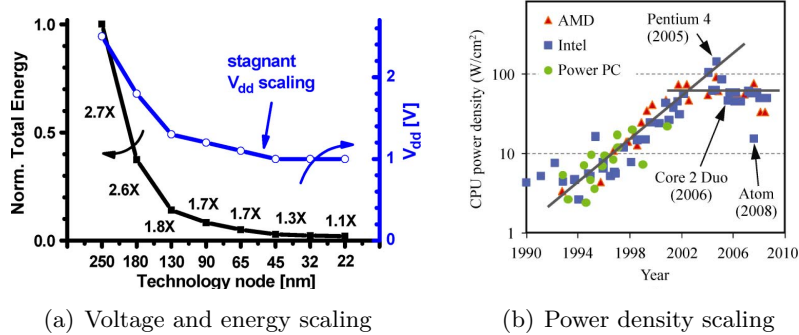


Figure 1.3: Voltage/Energy [17] and power density [18] scaling trends.

increase in heat generation that has a direct impact on the lifetime of MPSoCs. This increased heat generation leads to high temperature in the MPSoCs. While high-frequency single processing units have faced the similar case of high temperatures, the thermal profile of MPSoCs can have a more severe impact. This is mainly related to the localized heat generation of the several processing units of MPSoCs. Thus, the localized heat generation creates several localized high temperatures, which is known as thermal *hot spots*. The existence of several thermal *hot spots* would imply that there are other localized cold spots, which leads to the creation of the undesirable spatial *thermal gradients*. Moreover, the time-varying nature of workload processing requirements, or even when the processing goes to power-up and power-down cycles, leads to temporal *thermal cycles* [19] formation. To demonstrate the MPSoC thermal issues, an example of various hot spots location and thermal gradient is shown in Fig. 1.4(b). This figure shows a thermal response snapshot of the UltraSPARC T1 (Niagara) [20] MPSoC to a typical workload execution.

While high-density 2D MPSoCs face strong thermal challenges, these challenges are more prominent in the vertically-stacked 3D MPSoCs [21, 12]. Due to the vertical stacking of different dies, the thermal resistance of 3D MPSoCs is significantly increased to alarming values [12], compared to the increased temperature we demonstrate in the case of 2D MPSoCs. This is mainly due to the increased and non-

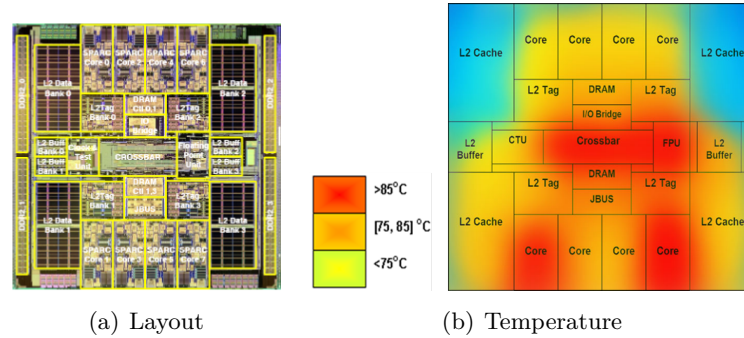


Figure 1.4: Floorplan layout and thermal response of the UltraSPARC T1 MP-SoC [20].

uniform thermal resistance at different stacked layers, based on their relative heat dissipation paths using conventional techniques such as placing heat sink on top of the top-most layer. For example, Fig. 1.5 shows the temperature of an emulated 3D MPSoC. This emulator is built by stacking 4 heat dissipation tiers on a substrate tier. In each of the heat dissipation tiers there is a number of controllable micro-heaters that are used to emulate the heat generation pattern of processing units similar to the actual pattern of each processing unit. In addition to these micro-heaters, there is a number of thermal sensors to capture the thermal profile of this emulator. This 3D thermal emulator has a heat sink placed at the top-most layer. The temperatures shown in this figure indeed confirm the expected high temperature and thermal gradient values of prospective high-performance 3D MPSoCs.

Thus, it is expected that high-density high-performance 3D MPSoCs are more prone to hot spots and thermal gradients. The existence of hot spots, thermal gradients, and thermal cycles heavily affect the MPSoC (2D and 3D) operation and lifetime, as shown in the following section.

1.3 Thermal Impact on 3D MPSoC Reliability and Performance

High temperature is undesirable in 3D MPSoCs operation due to the different device and interconnect reliability and degraded performance




Figure 1.5: Thermal state of a 5 tier (4 thermal dissipating tiers + 1 substrate base tier) high-density MPSoC thermal emulator. Temperature values are in Kelvin [22].

sources that are highly affected, directly and indirectly, by this rise in temperature. These sources would affect the reliability of 3D MP-SoCs by accelerating the processor aging or the *Mean-Time to Failure* (MTTF) [23], which is the statistical average time for the MP-SoC to breakdown permanently, as well as creating irreversible functional failures in the computation modules (e.g., storage) that limit the full utilization of these modules. In addition to the impacted reliability, high 3D MPSoC temperature would eventually lead to degraded performance by reducing the operating frequency due to increased propagation delays or reduced energy-efficiency resulted from the increased leakage power consumption. Thus, it is important to identify the temperature-induced reliability and degraded performance sources and elaborate more on their corresponding impact. The following paragraphs give an overview of these sources:

Bias Temperature Instability (BTI) [24]. This factor causes instabilities in the device behavior, due to the stress applied on the bias (e.g., a negative bias on the gate source voltage of a PMOS transistor). BTI can be split to two types, namely Negative BTI, which is related to the PMOS device stress, and Positive BTI, which is related to the NMOS device stress. The main degraded parameter due to BTI is the

threshold voltage, as shown in prior work [24]. The change in threshold voltage during stress (increase in threshold voltage) and release times (decrease in threshold voltage) has a dependency on a number of factors, which includes temperature [25]. Higher operating temperatures indeed have a direct impact on the threshold voltage based on BTI, which results in longer circuit delays and increase in dynamic power consumption.

Hot-Carrier Injection [26]. Hot-carrier injection occurs when a carrier gains enough energy to tunnel from the transistor source or drain to the dielectric material. This even accompanied with a rise in the device temperature. Hot-carrier injection occurs at normal temperature range, but the injection rate is increased as the operating (or stress) temperature is increased [27]. Based on the above observations, hot-carrier injection could lead to thermal positive feedback loop situation, i.e., the injection leads to increase in temperature that may trigger an increase in the injection rate. Consequently, hot-carrier injection would lead to thermal run-away.

Time-Dependent Dielectric Breakdown (TDDB) [28, 29] in high-k device dielectric and low-k interconnect dielectric. This is modeled as trap generated that leads to a leakage path through the oxide layer of the transistor. This is also referred in literature as *gate oxide breakdown*. TDDB has an exponential dependency on temperature [28] that accelerates the failure of a transistor by breaking down the dielectric, hence forming a constant conducting path. As a consequence, the faulty transistor would be permanently in a conducting state.

ElectroMigration (EM) in metallic interconnect [30]. Electromigration is a phenomenon that occurs in the IC interconnects (metal layers) due to high current densities. EM leads to a shift in the conducting ions location, hence causing a breakdown in the interconnects. EM has a strong dependency on the temperature resulted from the utilization of the IC and the joule self-heating of the interconnects due to the high current density. In this respect, higher operating temperatures would

1.3. Thermal Impact on 3D MPSoC Reliability and Performance 11

eventually lead to a breakdown of the metal layer, which may result in a complete failure of the IC.

Subthreshold Leakage Current [31]. Subthreshold leakage current is one of the sources of leakage current, and hence static power consumption in MPSoCs. It is the drain-source current of a transistor operating in the weak inversion region. Unlike the strong inversion region in which the drift current dominates, the subthreshold conduction is due to the diffusion current of the minority carriers in the channel. Subthreshold leakage is found to be the dominant component in the overall leakage current sources [32]. Subthreshold leakage current has a strong dependency on the operating temperatures with a sensitivity of $8 - 12x/100^{\circ}C$ [31].

The aforementioned sources can lead to system failure if no proper measures are taken. But these sources also degrade the MPSoCs (2D and 3D) operation from its original (also called time-zero) operating conditions. The following paragraphs highlight more on these affected parameters:

Mean-Time To Failure [23, 31]. The mean-time to failure would be heavily impacted due to the temperature impact on Time-Dependent Dielectric Breakdown (TDDB), Electromigration (EM), as well as stress migration and thermal cycling.

Temperature-Dependent Propagation Delays [33]. This change in the time required for a signal to travel between two modules is related to the thermally-induced delays in the logic gates (e.g. resulted from BTI), as well as the increase of the interconnect resistivity (e.g. resulted from Electro-migration). Another cause for the propagation delay is in the clock skew between different modules experiencing diverse thermal stress. In this respect, propagation delays have a strong dependency not just on the overall thermal state of the IC, but on the spatial thermal gradient as well. Indeed, previous work reports that a spatial gradient of $40^{\circ}C$ would create a 10% clock skew between different modules within a single IC.

Temperature-Dependent Leakage Power [31]. Leakage power is one of the sources of static power consumption in MPSoCs. The leakage power is a cause of various elements, such as the reverse-biased junction leakage current and the subthreshold leakage. These elements have a strong dependency on the operating temperatures. In fact, it has been shown by previous work that the leakage power has an exponential [34] dependency on temperature. Thus, it is crucial to prevent the operating MPSoCs from entering thermal runaway situations. Temperature-dependent leakage power may not be viewed as a failure mechanism, but high leakage power would cause a significant degradation in the power efficiency, as it would surpass the dynamic power consumption, where dynamic power is the effective power used in computations and is mainly workload dependent.

1.4 Advanced Cooling Technologies for 3D MPSoCs

To address the increasing thermal rise of 3D MPSoCs, several research initiatives have explored several advancing cooling strategies for the target architectures. For instance, there has been several research efforts to insert dummy thermal through-silicon vias (TTSVs) [35] to dissipate the heat generated from the layers further from the heat sink in a more efficient way.

Another advanced cooling technology uses injected fluids (single- or two-phase), between the different layers of the 3D MPSoCs. This cooling methodology, which is also known as interlayer liquid cooling [36], is achieved by manufacturing a cavity layer, for example rectangular microchannels or micro pin-fin structures [37], on the back-side of each silicon layer. A typical structure of 3D MPSoCs consists of two or more silicon tiers, which contain the processing and storage elements of the system. The communication between these tiers is realized with through-silicon vias (TSVs) that are etched in the residual silicon slab. To account for inter-tier liquid cooling, the porous cavity is realized by etching porous structures of different form and shapes (cf. Fig. 1.6). In the example shown in Fig. 1.6, the micro-channels are built, and distributed uniformly, in between the vertical layers for liquid flow. The




Figure 1.6: Layouts of the interlayer liquid-cooled 3D MPSoCs [38].

fluid flows through each channel at the same flow rate, but the liquid flow rate provided by the pump can be dynamically altered at runtime.

Manufacturing 3D MPSoCs with TSV interconnections and microchannels requires a series of microfabrication processes, namely (1) deep-reactive-ion-etching (DRIE) process for anisotropic silicon etching of both TSV openings and backside microchannels (cf. Fig. 1.7); (2) conformal thin film deposition for TSV sidewall insulation; (3) electroplating for conductive layer formation; (4) grinding for chip thinning, and finally (5) wafer- or die-level bonding for the stacking. A simplified illustration of a 3D stack with inter-tier liquid cooling is shown in Fig. 1.8.

Despite the benefits that liquid cooling brings in terms of significant thermal reduction, liquid cooling adds additional challenges to obtain a balanced thermal state with low spatial thermal gradients. As the coolant flows in microchannels, it experiences sensible heat absorption along the path [41]. This results in the coolant temperature increase from inlet to outlet, which causes a thermal gradient formation on the MPSoC surface even when the heat dissipation is uniform, as shown in Fig. 1.9(a). More commonly, in 3D MPSoCs with non uniform heat




Figure 1.7: SEM photos the wafer Back-side with the inlet-outlet openings while showing the micro-channels [39].




Figure 1.8: Simplified illustration of 3D stack with inter-tier liquid cooling [39].

dissipations, the existing thermal gradients and hot spots are aggravated by this characteristic of interlayer liquid cooling, as shown in Fig. 1.9(b). As a result, thermal gradients proliferate in 3D MPSoCs with liquid cooling. These gradients cause uneven thermally-induced stresses on different parts of the MPSoC, significantly undermining overall system reliability [31] (cf. Section 1.3).

From these observations we deduce that these new advanced cooling technologies bring both additional benefits and challenges. Thus,

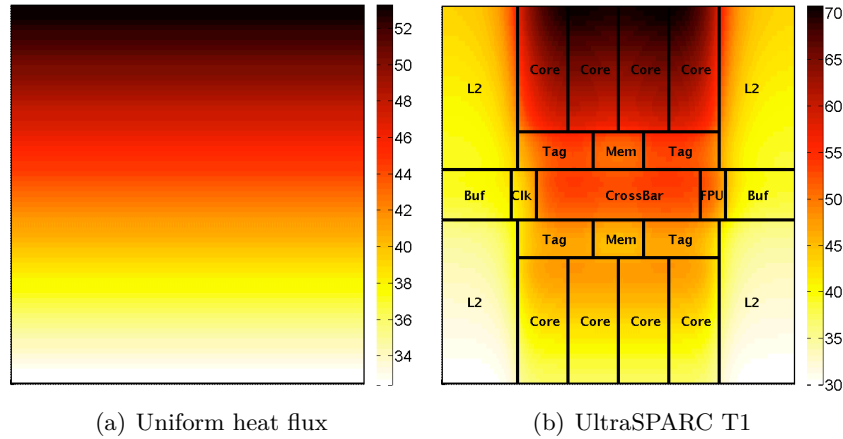


Figure 1.9: Steady-state temperature distribution of a 14mm x 15mm two-die 3D IC with (a) uniform (combined) heat flux density of $50\text{W}/\text{cm}^2$ and (b) the UltraSPARC T1 (Niagara-1) chip architecture [14]- the (combined) heat flux densities range from $8\text{W}/\text{cm}^2$ to about $64\text{W}/\text{cm}^2$. Direction of the coolant flow is from the bottom to the top of the figure [40].

it is crucial to find an optimized methodology to apply these technologies, alongside conventional 2D temperature balancing techniques in the context of resource-efficient temperature-aware 3D MPSoCs. Our interest in resource-efficiency varies from area to applied energy.

1.5 Survey Overview and Outline

In this survey, we provide an extensive survey that covers temperature-aware design optimizations and run-time management schemes for 3D MPSoCs, to avoid the rapid degradation of these architectures due to the thermal impact on reliability. The survey shows how the state-of-the-art tackles the thermal issues in the emerging 3D MPSoCs that include advanced cooling mechanisms, such as thermal-through-siliconvias (TTSVs) and interlayer liquid cooling. We perform this survey exploration in a top-down manner to cover all the temperature-aware aspects in the target architecture. In particular, we provide two main research directions that tackle the thermal issues, namely:

1. **Design-time technological solutions and temperature-aware optimizations.** In this category we explore various techniques that address, at *design-time*, the means of generating and dissipating heat in 3D MPSoCs. This includes the optimization of new advanced cooling and heat dissipation mechanisms, temperature-aware floorplanning, and design-time optimization of different parameters that would affect the thermal behavior of 3D MPSoCs.
2. **Run-time thermal management mechanisms.** We show in this category the various temperature-affecting knobs in the target 3D MPSoC, and how the state-of-the-art utilizes these knobs in developing several *run-time* thermal management mechanisms.

1.5.1 Related Survey Works

Our proposed survey overlaps with several surveys that exist in literature. An initial survey in thermally-aware design [42] explores the various design-optimization mechanisms for MPSoCs, both planar 2D and vertical 3D. In particular, this previous survey explains the various modeling framework approaches and explores thermal-aware floorplanning, and the means to recover from temperature-induced parameters degradation such as run-time shifts. However, there is no exploration for advanced cooling mechanisms, other design-time optimization mechanisms, or run-time thermal management techniques, which are explored in our proposed survey.

There is a recent survey work that explores various thermal management mechanisms for processing architectures [43]. In this previous survey, several mechanisms for balancing temperatures in 2D and 3D MPSoCs are explored, namely thermal sensor placement, run-time thermal management, floorplanning, operating system/compilation techniques, and a brief exploration on liquid cooling. However, this survey does not provide a systematic classification of the temperature optimization research field. In our proposed survey, however, we provide our classification in a more systematic methodology that follows a top-down reasoning to cover most of the research directions in 3D MPSoC

temperature optimization.

Finally, another survey explores thoroughly vertical integration in 3D ICs [44]. This previous work explores the various electronic-design automation tools needed for 3D architectures. Then it explores the various architecture flavors in 3D ICs, namely 3D field programmable gate arrays (FPGA) and 3D MPSoC designs. Thus, this previous survey work is complementary to our proposal.

1.5.2 Survey Organization

This survey starts with providing an overview of the recent thermal modeling approaches developed for the target 3D MPSoC architectures in Section 2. In Section 3, we explore the design-time optimization schemes to minimize the temperature-oriented issues in 3D MPSoCs. Section 4 shows our exploration in run-time thermal management schemes for 3D MPSoCs. We first start by stating the various temperature-affecting control knobs, then we show the classification of different management schemes that use these knobs to balance the dissipated heat in the target architecture. Finally, we summarize our work in this survey in the conclusion.

References

- [1] M. Pedram. Energy-efficient datacenters. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 31(10):1465–1484, 2012.
- [2] G. E. Moore. Cramming more components onto integrated circuits. *Electronics Magazine*, 38(8):114–117, 1965.
- [3] W. Knight. Two heads are better than one [dual-core processors]. *IEE Review*, 51(9):32–35, 2005.
- [4] P. P. Gelsinger. Microprocessors for the new millennium: Challenges, opportunities, and new frontiers. In *ISSCC*, pages 22–25, 2001.
- [5] L. Benini and G. De Micheli. Networks on chips: a new soc paradigm. *Computer*, 35(1), 2002.
- [6] T. Hey et al. *The Fourth Paradigm: Data-Intensive Scientific Discovery*. Microsoft Research, 2009.
- [7] A. Burns et al. Shimmer a wireless sensor platform for noninvasive biomedical research. *IEEE Sensors Journal*, 10(9), 2010.
- [8] X. Xie et al. A low-power digital ic design inside the wireless endoscopic capsule. *IEEE Journal of Solid-State Circuits*, 41(11), 2006.
- [9] K. Nomura et al. Performance analysis of 3d-ic for multi-core processors in sub-65nm cmos technologies. In *ISACS*, pages 2876–2879, 2010.
- [10] W. R. Davis et al. Demystifying 3d ics: the pros and cons of going vertical. *IEEE Design and Test of Computers*, 22(6):498–510, 2005.

- [11] H. S. Lee and K. Chakrabarty. Test challenges for 3d integrated circuits. *IEEE Design and Test of Computers*, 26(5):26–35, 2009.
- [12] A. K. Coskun et al. Dynamic thermal management in 3D multicore architectures. In *DATE*, pages 1410–1415, 2009.
- [13] J. Meng and A. K. Coskun. Analysis and runtime management of 3d systems with stacked dram for boosting energy efficiency. In *DATE*, 2012.
- [14] A. Leon et al. A power-efficient high-throughput 32-thread SPARC processor. *ISSCC*, 42(1):7 – 16, 2007.
- [15] Yole Development. 3d ic and tsv report. yole.fr/pagesAn/products/Report_sample/3DIC.pdf, 2007.
- [16] M. Horowitz et al. Scaling, power, and the future of cmos. In *IEDM*, pages 7–15, 2005.
- [17] R. G. Dreslinski et al. Near-threshold computing: Reclaiming moore’s law through energy efficient integrated circuits. *Proc. of the IEEE*, 98(2), 2010.
- [18] E. Pop. Energy dissipation and transport in nanoscale devices. *Nano Research*, 3(3):147–169, 2010.
- [19] J. Donald and M. Martonosi. Techniques for multicore thermal management: Classification and new exploration. In *ISCA*, pages 78–88, 2006.
- [20] A. K. Coskun, T. Simunic Rosing, and K. Whisnant. Temperature aware task scheduling in MPSoCs. In *DATE*, pages 1659–1664, 2007.
- [21] C. Zhu et al. Three-dimensional chip-multiprocessor run-time thermal management. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 27(8):1479–1492, August 2008.
- [22] P. G. Del Valle and D. Atienza. Emulation-based transient thermal modeling of 2d/3d systems-on-chip with active cooling. *Microelectronics Journal*, 42:564–571, 2011.
- [23] D. Atienza, G. De Micheli, L. Benini, J. L. Ayala, P. G. Del Valle, M. DeBole, and V. Narayanan. Reliability-aware design for nanometer-scale devices. In *ASPDAC*, pages 549–554, 2008.
- [24] S. M. Lim et al. Effects of BTI during AHTOL on SRAM VMIN. In *IRPS*, 2011.
- [25] T. Grasser et al. Simultaneous extraction of recoverable and permanent components contributing to bias-temperature instability. In *IEDM*, pages 801–804, 2007.

- [26] E. Takeda and N. Suzuki. An empirical model for device degradation due to hot-carrier injection. *IEEE Electron Device Letters*, 4(4):111–113, 1983.
- [27] P. Heremans et al. Temperature dependence of the channel hot-carrier degradation of n-channel mosfet. *IEEE Transactions on Electron Devices*, 37(4):980–993, 1990.
- [28] R. Moazzami et al. Temperature acceleration of time-dependent dielectric breakdown. *IEEE Transactions on Electron Devices*, 36(11):2462–2465, 1989.
- [29] D. Qian and D. J. Dumin. The electric field, oxide thickness, time and fluence dependences of trap generation in silicon oxides and their support of the e-model of oxide breakdown. In *IPFA*, 1999.
- [30] J. R. Black. Electromigration—a brief survey and some recent results. *IEEE Transactions on Electron Devices*, 16(4), 1969.
- [31] M. Pedram and S. Nazarian. Thermal modeling, analysis, and management in vlsi circuits: principles and methods. *Proceedings of the IEEE*, 94(8), 2006.
- [32] Y. Taur et al. Cmos scaling into nanometer regime. *Proceedings of the IEEE*, 85(4):486–504, 1997.
- [33] T. Sato et al. On-chip thermal gradient analysis and temperature flattening for soc design. In *ASP-DAC*, pages 1074–1077, 2005.
- [34] Jayanth Srinivasan et al. The case for lifetime reliability-aware microprocessors. In *ISCA'04*, pages 276–287, 2004.
- [35] J. Cong and Y. Zhang. Thermal via planning for 3-d ics. In *ICCAD*, 2005.
- [36] T. Brunschwiler et al. Interlayer cooling potential in vertically integrated packages. *Microsyst. Technol.*, 15(1):57 – 74, 2009.
- [37] T. Brunschwiler et al. Hotspot-optimized interlayer cooling in vertically integrated packages. In *MRS Fall Meeting*, 2009.
- [38] M. M. Sabry et al. Energy-Efficient Multi-Objective Thermal Control for Liquid-Cooled 3D Stacked Architectures. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 30(12):1883–1896, 2011.
- [39] M. M. Sabry et al. Towards thermally-aware design of 3d mpsoCs with inter-tier cooling. In *DATE*, pages 1466–1471, 2011.
- [40] M. M. Sabry, A. Sridhar, and D. Atienza. Thermal balancing of liquid-cooled 3d-mpsoCs using channel modulation. In *DATE*, 2012.

- [41] A. Sridhar et al. 3D-ICE: Fast compact transient thermal modeling for 3D-ICs with inter-tier liquid cooling. In *ICCAD*, pages 463–470, 2010. <http://esl.epfl.ch/3d-ice.html>.
- [42] Y. Zhan et al. Thermally aware design. *Foundatin and Trends in Electronic Design Automation*, 2(3):255–370, 2008.
- [43] J. Kong et al. Recent thermal management techniques for microprocessors. *ACM Computing Surveys*, 44(3):13:1–13:42, 2012.
- [44] G. Sun et al. Three-dimensional integrated circuits: Design, eda, and architecture. 5(1-2):1–151.
- [45] John Lienhard-IV and John Lienhard-V. *A heat transfer textbook*. Phlogiston Press, Cambridge, Massachusetts, 2006.
- [46] Wei 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.
- [47] F. Incropera, D. Dewitt, T. Bergman, and A. Lavine. *Fundamentals of heat and mass transfer*. John Wiley and Sons, 2007.
- [48] Z. Chen, X. Luo, and S. Liu. Thermal analysis of 3D packaging with a simplified thermal resistance network model and finite element simulation. In *ICEPT-HDP*, pages 737–741, aug. 2010.
- [49] T Wang and C Chen. 3-D thermal-ADI: a linear-time chip level transient thermal simulator. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 21(12):1434–1445, December 2002.
- [50] www.ansys.com/products/fluid-dynamics/cfx.
- [51] J.-M. Koo, S. Im, L. Jiang, and K. E. Goodson. Integrated microchannel cooling for three-dimensional electronic circuit architectures. *Journal of Heat Transfer*, 127(1):49–58, 2005.
- [52] H. Mizunuma et al. Thermal modeling and analysis for 3D-ICs with integrated microchannel cooling. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 30(9):1293–1306, 2011.
- [53] Y. J. Kim et al. Thermal characterization of interlayer microfluidic cooling of three-dimensional integrated circuits with nonuniform heat flux. *Journal of Heat Transfer*, 132(4):1–9, 2010.
- [54] A. Sridhar et al. Compact transient thermal model for 3D-ICs with liquid cooling via enhanced heat transfer cavity geometries. In *THERMINIC*, pages 1–6, 2010.

- [55] A. Sridhar et al. 3D-ICE: A compact thermal model for early-stage design of liquid-cooled ics. (*to appear in*) *IEEE Transactions on Computers*, 2013.
- [56] H. Qian et al. Thermal simulator of 3d-ic with modeling of anisotropic tsv conductance and microchannel entrance effects. In *ASPDAC*, 2013.
- [57] A. H. Ajami et al. Analysis of non-uniform temperature-dependent interconnect performance in high performance ics. In *DAC*, 2001.
- [58] A. H. Ajami, K. Banerjee, and M. Pedram. Modeling and analysis of nonuniform substrate temperature effects on global ULSI interconnects. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 24(6):849–861, June 2005.
- [59] H. Su et al. Full-chip leakage estimation considering power supply and temperature variations. In *ISLPED*, pages 78–83, 2003.
- [60] M. Ni et al. An analytical study on the role of thermal tsvs in a 3dic chip stack. In *DATE*, 2010.
- [61] M. Healy et al. Multiobjective microarchitectural floorplanning for 2-D and 3-D ICs. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 26(1), Jan 2007.
- [62] M. S. Bakir et al. 3d heterogeneous integrated systems: Liquid cooling, power delivery, and implementation. In *CICC*, pages 663 – 670, 2008.
- [63] K. Kota, P. Hidalgo, Y. Joshi, and A. Glezer. Thermal management of a 3d chip stack using a liquid interface to a synthetic jet cooled spreader. In *THERMINIC*, 2009.
- [64] A. Furmanczyk et al. Multiphysics modeling of integrated microfluidic-thermoelectric cooling for stacked 3d ics. In *SEMI-THERM*, pages 35–41, 2003.
- [65] H. N. Phan and D. Agonafer. Experimental analysis model of an active cooling method for 3d-ics utilizing a multidimensional configured thermoelectric. In *SEMI-THERM*, pages 55–58, 2010.
- [66] P. Greenalgh. Big.little processing with arm cortex-a15 and cortex-a7. www.arm.com/files/downloads/big.LITTLE_Final.pdf.
- [67] N. Xu et al. Thermal-aware post layout voltage-island generation for 3d ics. *Journal of Computer Science and Technology*, 28(4):671–681, 2013.
- [68] K. Puttaswamy and G. H. Loh. Thermal herding: Microarchitecture techniques for controlling hotspots in high-performance 3d-integrated processors. In *HPCA*, pages 193–204, 2007.

- [69] M. Bao et al. Temperature-aware idle time distribution for energy optimization with dynamic voltage scaling. In *DATE*, 2010.
- [70] B. C. Schafer et al. Temperature-aware compilation for vliwprocessors. In *RTCSA*, 2007.
- [71] Y. Han et al. Temperature aware floorplanning. In *Workshop on Temperature Aware Computing Systems*, 2005.
- [72] K. Sankaranarayanan, S. Velusamy, M. Stan, and K. Skadron. A case for thermal-aware floorplanning at the microarchitectural level. *Journal of Instruction-Level Parallelism*, 8:1–16, 2005.
- [73] W.-L. Hung et al. Thermal-aware floorplanning using genetic algorithms. In *ISQED*, 2005.
- [74] J. Cong, J. Wei, and Y. Zhang. A thermal-driven floorplanning algorithm for 3D-ICs. In *ICCAD*, pages 306–313, 2004.
- [75] W.-L. Hung et al. Interconnect and thermal-aware floorplanning for 3D microprocessors. In *ISQED*, pages 98–104, 2006.
- [76] M. Ekpanyapong et al. Thermal-aware 3d microarchitectural floorplanning. Georgia Institute of Technology, 2004.
- [77] X. Zhou et al. Temperature-aware register reallocation for register file power-density minimization. *ACM TODAES*, 14(2), 2009.
- [78] M. M. Sabry et al. Thermal-aware compilation for system-on-chip processing architectures. In *GLSVLSI*, 2010.
- [79] A. Jorio et al. *Carbon nanotubes: advanced topics in the synthesis, structure, properties and applications*. Springer, 2008.
- [80] M. Freitag et al. Energy dissipation in graphene field-effect transistors. *Nano Letters*, 9(5):1883–1888, 2009.
- [81] A. A. Balandin et al. Superior thermal conductivity of single-layer graphene. *Nano Letters*, 8(3):902–907, 2008.
- [82] B. Goplen and S. Sapatnekar. Thermal via placement in 3d ics. In *ISPD*, pages 167–174, 2005.
- [83] X. Li et al. Lp based white space redistribution for thermal via planning and performance optimization in 3d ics. In *ASPDAC*, pages 209–212, 2008.
- [84] T. Zhang et al. Temperature-aware routing in 3d ics. In *DAC*, 2006.
- [85] Hai Wei et al. Cooling three-dimensional integrated circuits using power delivery networks. In *IEDM*, pages 14.2.1–14.2.4, 2012.

- [86] M. Saini and R. L. Webb. Heat rejection limits of air cooled plane fin heat sinks for computer cooling. *IEEE Transactions on Components and Packaging Technologies*, 26(1):71–79, 2003.
- [87] S. Zimmermann et al. Aquasar: A hot water cooled data center with direct energy reuse. *Energy*, 43(1):237–245, 2012.
- [88] S. Szczukiewicz et al. Two-phase flow boiling in a single layer of future high-performance 3d stacked computer chips. In *ITHERM*, 2012.
- [89] T. Brunschwiler et al. Angle-of-attack investigation of pin-fin arrays in nonuniform heat-removal cavities for interlayer cooled chip stacks. In *SEMI-THERM*, pages 116–124, 2011.
- [90] F. Alfieri et al. 3d integrated water cooling of a composite multilayer stack of chips. *Journal of Heat Transfer*, 132(12), 2010.
- [91] H. Qian et al. Cyber-physical thermal management of 3D multi-core cache-processor system with microfluidic cooling. *ASP Journal of Low Power Electronics*, 7(1):1–12, 2011.
- [92] H. Qian, Ch. Chang, and H. Yu. An efficient channel clustering and flow rate allocation algorithm for non-uniform microfluidic cooling of 3d integrated circuits. *Integration, the VLSI Journal*, 46(1):57–68, 2013.
- [93] B. Shi et al. Non-uniform micro-channel design for stacked 3D-ICs. In *DAC*, pages 658–663, 2011.
- [94] B. Shi and A. Srivastava. Cooling of 3d-ic using non-uniform micro-channels and sensor based dynamic thermal management. In *Annual Allerton Conf. Communication, Control, and Computing*, 2011.
- [95] H. Qian and C. H. Chang. Microchannel splitting and scaling for thermal balancing of liquid-cooled 3dic. In *ISCAS*, 2013.
- [96] M. M. Sabry et al. Greencool: An energy-efficient liquid cooling design technique for 3d mpsoCs via channel width modulation. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 32(4):524–537, 2013.
- [97] R. Shah and A. London. *Laminar flow forced convection in ducts*. New York: Academic Press, 1978.
- [98] Y. Tan et al. Modeling and simulation of the lag effect in a deep reactive ion etching process. *Journal of Micromechanics and Microengineering*, 16, 2006.
- [99] A. K. Coskun et al. Energy-efficient variable-flow liquid cooling in 3D stacked architectures. In *DATE*, pages 111–116, 2010.

- [100] A. K. Coskun, T. Simunic Rosing, K. A. Whisnant, and K. C. Gross. Static and dynamic temperature-aware scheduling for multiprocessor socs. *IEEE Transactions on VLSI*, 16(9):1127–1140, Sept. 2008.
- [101] A. Bartolini et al. A distributed and self-calibrating model-predictive controller for energy and thermal management of high-performance multicores. In *DATE*, 2011.
- [102] J. Kim et al. Correlation-aware virtual machine allocation for energy-efficient datacenters. In *DATE*, 2013.
- [103] Y. U. Ogras, R. Marculescu, D. Marculescu, and E. G. Jung. Design and management of voltage-frequency island partitioned networks-on-chip. *IEEE Transactions on VLSI*, 17(3):330–341, 2009.
- [104] P. Bogdan, S. Jian, R. Tornero, and R. Marculescu. An optimal control approach to power management for multi-voltage and frequency islands multiprocessor platforms under highly variable workloads. In *ISNoC*, pages 35–42, 2012.
- [105] J. Choi et al. Thermal-aware task scheduling at the system software level. In *ISLPED*, 2007.
- [106] A. K. Coskun et al. Temperature management in multiprocessor socs using online learning. In *DAC*, pages 890–893, 2008.
- [107] Festo electric automation technology. <http://www.festo-didactic.com/ov3/media/customers/1100/00966360001075223683.pdf>.
- [108] T. Emi et al. Tape: Thermal-aware agent-based power economy for multi/many-core architectures. In *ICCAD*, pages 302–309, 2009.
- [109] X. Zhou et al. Thermal management for 3D processors via task scheduling. In *ICPP*, pages 115–122, 2008.
- [110] T. Ebi et al. Agent-based thermal management using real-time i/o communication relocation for 3d many-cores. In *PATMOS*, pages 112–121, 2011.
- [111] F. Allgöwer and A. Zheng. *Nonlinear Model Predictive Control*. Birkhäuser, 2000.
- [112] H. Wang et al. Thermal management via task scheduling for 3d noc based multi-processor. In *ISOC*, pages 440–444, 2010.
- [113] S. Liu et al. Thermal-aware job allocation and scheduling for three dimensional chip multiprocessor. In *ISQED*, pages 390–398, 2010.
- [114] F. Zanini, M. M. Sabry, D. Atienza, and G. De Micheli. Hierarchical thermal management policy for high-performance 3d systems with liquid cooling. *IEEE JETCAS*, 1(2):88–101, 2011.

- [115] F. Mulas et al. Thermal balancing policy for multiprocessor stream computing platforms. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 28(12):1870–1882, 2009.
- [116] Y. Diao et al. Using mimo linear control for load balancing in computing systems. In *ACC*, pages 2045–2050, 2004.
- [117] M. Su and H. Chang. Application of neural networks incorporated with real-valued genetic algorithms in knowledge acquisition. *Fuzzy Sets and Systems*, 112(1):85 – 97, 2000.
- [118] H. T. Nguyen and N. R. Prasad. *Fuzzy modeling and control, selected works of M. Sugeno*. CRC press, 1999.
- [119] F. Zanini, D. Atienza, and G. De Micheli. A combined sensor placement and convex optimization approach for thermal management in 3d-mpsoc with liquid cooling. *Integration, the VLSI journal*, 46:33–43, 2013.
- [120] A. Bemporad et al. The explicit linear quadratic regulator for constrained systems. *Automatica*, 38(1):3 –20, 2002.
- [121] G. Zhao et al. Processor frequency assignment in three-dimensional mpocs under thermal constraints by polynomial programming. In *APC-CAS*, pages 1668–1671, 2008.
- [122] A. Sridhar et al. Steam: a fast compact thermal model for two-phase cooling of integrated circuits. In *ICCAD*, 2013.
- [123] P. Gupta et al. Underdesigned and opportunistic computing in presence of hardware variability. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 32(1), 2013.
- [124] J. Henkel et al. Design and architectures for dependable embedded systems. In *CODES+ISSS*, 2011.