Abstract

HUA, HAO. Design and Verification Methodology for Complex Three-Dimensional Digital Integrated Circuit. (Under the guidance of Professor W. Rhett Davis).

Three-Dimensional Integrated Circuits (3DICs) have recently attracted great interest from researchers and IC designers as a possible solution to fill the gap between device and interconnect scaling. Various studies have demonstrated the potential performance improvement of 3DICs by eliminating long interconnects, repeaters, and clock buffers.

Though 3DICs are attractive, there are significant challenges associated with this topic. The most fundamental issue in 3DIC is heat dissipation. The thermal effect has impacted the conventional high-performance 2DICs in deep sub-micron technology nodes. Its effect will aggravate 3DICs due to two major reasons: higher power density, and lower thermal conductivity caused by more insulating dielectric layers. Furthermore, while 3D integration provides more design flexibility, this technology also introduces much higher design complexity. The existing 2D physical design methodology cannot be simply extended to a 3D case because of the huge obstacles in the z-direction and thermal constraints. Efficient design flows and algorithms must be developed to facilitate 3DIC design.

This dissertation proposes a design and verification methodology, along with analyses of delay, thermal, and reliability of a 3D system. The methodology uses commercial 2D CAD tools with Python and Tcl scripts to link them together. The scripts modify the output files

(or databases) of the commercial tools and add 3D features to them. The entire flow achieves RTL-to-GDSII physical design automation for 3DICs.

Design trade-offs and timing reliability of 3D systems are two other major issues of this dissertation. Non-idealities threaten to diminish the benefit and may cause reliability problems in 3D systems. These non-idealities must be monitored during the design procedure. With a fast yet accurate temperature dependency model, these non-idealities were successfully taken into consideration during both design and verification phases. The final performance analyses of two benchmark circuits show that the 3DIC achieves a maximum reduction of 27% on energy and 20% on delay compared to the conventional 2DIC approach. Finally, a performance trend study of 3DIC as the technology node shrinks is also performed to guide future 3DIC designs.

Design and Verification Methodology for Complex Three-Dimensional Digital Integrated Circuit

By

Hao Hua

A dissertation submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the Degree of
Doctor of Philosophy

ELECTRIAL ENGINEERING

Raleigh, 2006

APPROVED BY:

W. RHETT DAVIS Chair of Advisor Committee	PAUL D. FRANZON
XUN LIU	MARK A. L. JOHNSON

Biography

Hao Hua was born in Suzhou, China, which is considered the heaven on earth by Chinese people. He received the B.S. degree from electronics engineering department, Fudan University, Shanghai in 2000. From 2000-2001, he was with Shanghai Magima Electronics and later with Shanghai ZHSOC as an IC design engineer. From 2001-2006 he was pursuing his Ph.D. degree in department of electrical and computer engineering at North Carolina State University. His research interests include design and verification methodology for complex three-dimensional digital circuits, on-chip Signal Integrity, and System on a Chip (SoC) design methodology.

Acknowledgements

I would like to summarize my acknowledgements to the people who have been supporting and collaborating me on the work. Without them, this dissertation would never have been accomplished. I am grateful to be a part of an innovative project, a friendly group, and an enjoyable working environment. I would like to thank the people who gave me this opportunity.

First and foremost, I would like to thank my family: my parents Quanrong and Jingyu, my grandma Rong, my wife Yue, and my coming baby. Their love and support have been my source of inspiration, power, and ambition. Even the physical distance between my (grand)parents and I is great, I can feel a strong tie linking us together. My wife has been encouraging me when I felt depressed, and she has become the sunshine of my life. Without them, I could not imagine how I could ever be able to manage all the difficulties and challenges.

I would like to thank my advisor, Dr. W. Rhett Davis, who brought me into the design automation world and has guided me through my study. I would probably never realize the beauty of VLSI if I had not talked with him. He has also supported me in various aspects. Without him, my dream of being a Ph.D. would never come true.

I am grateful to my committee members, Dr. Paul D. Franzon, Dr. Xun Liu, and Dr. Mark A. L. Johnson. I personally think that exchanging idea is one of the most important things during any work. All my committee members were never at a loss of ideas. They have been

continuously giving me advice on how to improve my work. Dr. Franzon's sharp vision on the big picture impressed me every time I had a chance to talk with him.

Next, I would like to take this chance to thank my colleagues and friends. I would like to thank Ravi Jenkal, Samson Melamed, Chris Mineo, Kory Schoenfliess, and Ambarish Sule for their generous support. They have provided all the complementary materials to facilitate my research during my stay. It is always a great pleasure to work with smart people. I also would like to thank Jian Xu and Leon Zhang from whom I learned a lot about circuit design, which helped me greatly during the initial phase of my research.

Lastly, I would like to thank anyone that ever had a helpful discussion with me, including people inside and outside North Carolina State University.

Table of Contents

List of Figures

List of Tables

Chapter 1.

Introduction

1.1 Motivation

According to the ITRS roadmap 2005 [11], device delay is continuously scaled down while interconnect delay increases dramatically (Figure 1.1). Under ideal scaling, all dimensions of wires are shrunk 0.7x per generation. The wire capacitance per micron remains invariant while the resistance doubles, resulting in 1.4x increase in wire delay every generation [20]. Very large scale integrated (VLSI) circuits performance is restricted by long interconnects in deep sub-micron (DSM) technology nodes [41].

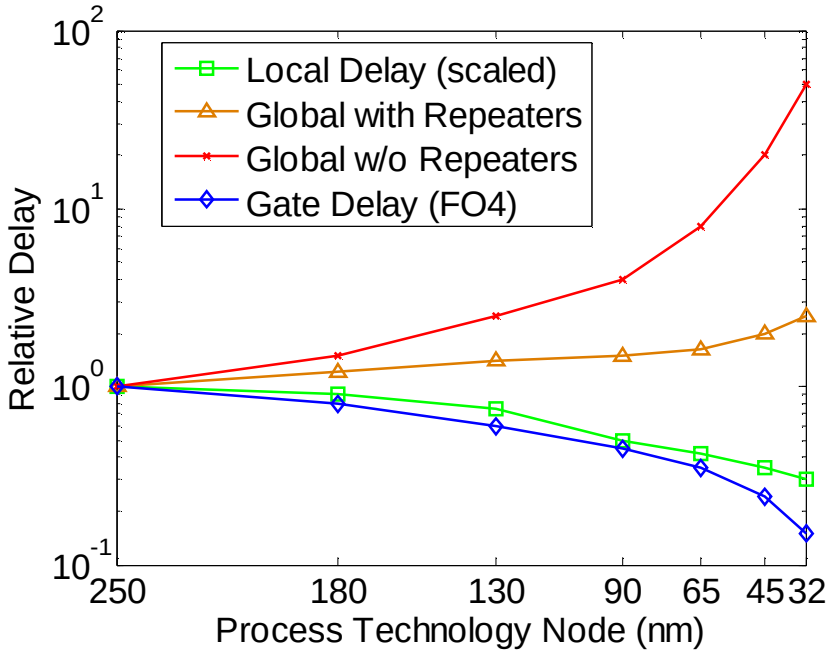

Figure 1.1: Interconnect roadmap.

The relationship known as Moore's Law [68] shows that the number of transistors on a chip doubles about every two years, and the chip complexity doubles every 18 months. This law drives the device to shrink significantly. Though device size becomes smaller, chip size is continually increasing due to the growing demand for functionality and higher performance [41]. Unlike local or intermediate interconnects, global interconnects do not scale in length since they communicate signals across a chip [43]. Longer interconnects require more repeaters to be inserted to achieve timing closure and signal integrity (SI) [15]. P. Saxena *et al* [20] expects that if the current trends continue, 35% of the total cells will be repeaters in 45nm technology node and this number will increase to 70% in 32nm technology node. The difficulty is that interconnect behavior and the number of repeaters are hard to predict before detailed physical design. Iterative back annotation, re-synthesis, placement and routing are required, and core utilization will also become another restriction that may force design iterations.

Copper (Cu) with low-k dielectric was introduced to alleviate interconnect problems [34]. Although copper has significantly reduced interconnect delay, some researchers predict that further reduction cannot be achieved by introducing new material after 130nm technology node [15]. Low-k materials are under investigation, but are problematic due to adhesion failure [11]. There is an urgent need for innovative design technologies that can reduce interconnect wirelength.

Three-dimensional integrated circuit (3DIC) technology aims to alleviate the above interconnect problems in today's VLSI design. A 3DIC is a circuit that has active devices on more than one silicon plane. One obvious advantage of 3DICs is the increase in device

density and the reduction of chip area. The reduced distance between blocks tends to reduce the interconnect wirelength. Studies showed that 3D integration can significantly reduce wirelength compared with its 2D counterpart [15] [33]. A shorter wirelength will reduce parasitic RC delays and lead to a smaller clock period as well as lower power dissipation.

In this work, we will concentrate on digital 3DICs. Performance trends of digital 3DICs such as timing and power were studied and proved to be attractive [32]-[33]. Tools regarding 3D floorplanning [13], 3D placement [9] and thermal optimization [5] were developed to facilitate 3D design. However, these investigations ignore non-idealities such as routing congestion and thermal impact on performance, which threaten to diminish the benefit of 3DIC. Therefore theoretical analysis of a 3DIC is not enough. We must be able to fabricate real 3D chips and perform quantitative comparisons of 3D and conventional (2D) ICs to show the real advantage of 3D integration.

1.2 Three Dimensional Integration Technologies

1.2.1 General 3D Integration Technology

There are many technologies that can be considered 3D integrations. When compared in terms of the method of assembly, 3D technologies can be defined as die scale integration or wafer scale integration. In the die scale integration, wafers are processed and cut into dies first. Dies are then aligned, stacked, and thinned. Though-silicon vias are rarely used in die scale integration due to fabrication cost and difficulty. Dies are connected through either pads on peripheral or solder bump. Because each die has its own IOs, it can be fully tested

before stacking to improve yield [76]. In the wafer scale integration, wafers are first stacked then cut into dies and the peripherals of the final chip are usually on top tier (in this dissertation, we give the name "tier" to each active layer and its associated metal layers). Wafers can be thinned to the scale of several microns in SOI technology or a little thicker than 10 μm in a bulk technology. Through-silicon vias are used extensively to communicate signals between different tiers. Wafer scale integration provides high-density inter-wafer interconnections and is more likely to improve performance. However, wafer scale integration usually results in higher cost because wafers are not testable before stacking.

Figure 1.2 [42] illustrates seven approaches. Figure 1.2 (a), (b), and (c) can be called die scale 3D integration. (d) and (e) belong to the AC coupling. They can be either die scale or wafer scale 3D integration. (f) and (g) are wafer scale back to face 3D integration. They represent bulk and SOI 3D integration respectively. More wafer level 3D integrations like face to face, back to back or a combination of these schemes are available from different foundries such as: RTI, Ziptronix, Terrzaron, IBM, and MIT Lincoln Lab. In this thesis, we will discuss the MIT wafer level SOI 3D integration only, which will be discussed in the following section.

The summary of these different approaches is shown in Table 1.1[1][42]. Wire bond is the most common approach. Connections between chips go through the board or chip carrier and back to other chips in the stack. This approach is limited by the resolution of wire bonders and becomes increasingly difficult as the number of IOs in the chip stack increases. In this approach, the signals can be only connected through periphery. The microbump uses only top or sometimes top two metals as bumps to communicate signals between dies. These

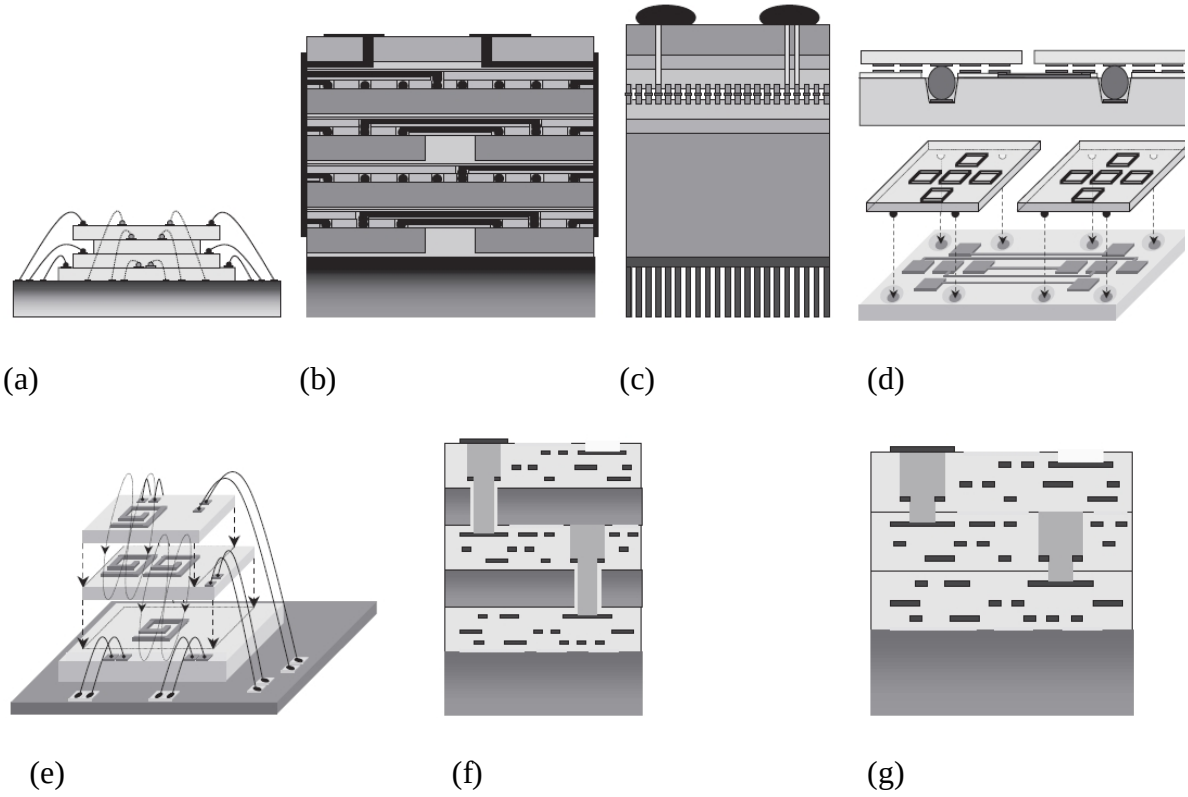

Figure 1.2: Illustration of 3D integration technologies. (a) wire bond; (b) microbump-3D package; (c) microbump face-to-face, 2 tier limited (d) contactless-capacitive; (e) contactless-inductive; (f) through-via bulk; and (f) through-via SOI (source [42]) [1]

bumps have a much smaller pitch and can be placed anyway over the chip, and hence offers a greater vertical interconnect density. However, microbump does not significantly reduce parasitic capacitances because signals still need to be routed to the periphery before sending them back to the destination. Hence, neither wire bond nor microbump technology can achieve very high vertical interconnect density. The Contact-less technology utilizes the

[1] Thanks to Dr. John Wilson for creating these figures and tables.

electro-magnetic field of a pair of coupled capacitors, inductors, or a combination of them to transmit signals between dies or wafers. This approach eliminates the step for creating inter-tier DC connectivity and eliminates the need to route signals to the periphery. The major limitations of AC coupling is the large size of the coupling capacitance (or inductance) and the complexity of designing an analog transceiver circuit. Through-via technology seems the most promising one compared to the other 3D integration technologies. This technology offers the greatest vertical interconnect density. Wafers are bonded and then ultra-thinned. The ultra-thinning process provides better thermal removal path than die scale integration technology. All communications between tiers are transmitted by through-vias. There is no limit on number of possible tiers unless heat inside does not allow more integration. The disadvantage of this technology is relatively high cost. Since the tiers are not known to be good before assembly, the yield drops quickly as tier number increases.

Table 1.1: Comparison of 3D interconnection technologies (Source [42]) [1]

Characteristic	Wire Bond	Micro-bump		Contact-less		Through via	
		3D Package	Face-to-face	Capacitive	Inductive	Bulk	SOI
Assembly level	Die	Die	Die	Die/Wafer	Die/Wafer	Wafer	Wafer
Tier Limit	Assembly process	Heat	Assembly process	Assembly process	Heat	Heat, yield	Heat, yield
Vertical Pitch (um)	35-100	25-50	10-100	50-200	50-150	50	5
Metal layers blocked by pad	All	Top 1 to 2	Top 1 to 2	Top	Top 1 to 2	All, top	All, top

Various processing technologies and bonding schemes for wafer level 3D integration are available. In a “bottom-up” scheme, devices are processed sequentially. When a bottom tier is fabricated, another device layer is formed and devices are fabricated on that layer. In a

"top-down" scheme, multiple 2D circuits are fabricated and then assembled to form a 3DIC [10]. The known processing candidates are: Beam Re-crystallization, Wafer Bonding, Silicon Epitaxial Growth and Solid Phase Crystallization. The advantages and disadvantages of each processing technology are listed in Table 1.2 [15][29][30]. All the technologies except wafer bonding need to do post process that will introduce process complication and hence degrade the device performance. This degradation is hard to quantify and thus makes them less promising for 3D integration. Wafer bonding technology is the best technology currently available for 3D integration because it is the easiest to implement and is the most reliable process. Since the wafers are process separately without any interaction, we can assume same device parameters at same operation conditions. Due to the above reasons, wafer bonding is selected as the 3D integration scheme for MIT Lincoln Lab 3D process. This technology provides high vertical connectivity, similar device properties and "System-on-a-Chip" integration [17] [21]. A process flowchart for wafer bonding technology is shown in Figure 1.3 [12][21].

Table 1.2: 3D fabrication technologies.

Beam Re-crystallization	Wafer Bonding	Silicon Epitaxial Growth	Solid Phase Crystallization
Deposit poly-sillicon and fabricate thin film transistors (TFTs) - high temp of melting poly-sillicon - high V_T - low carrier mobility - high-performance TFTs	Bond two or more fully processed wafers together - similar devices properties - independent of temp - good for chip assembly with different processing technologies - extra inter-tier vias and alignment	Epitaxially grow a single crystal Si - high vertical connectivity - significant degradation in quality of devices on lower layers - difficult to implement mixed technology	Low Temp alternative to Silicon Epitaxial - offers flexibility of creating multiple layers - compatible with current processing environments - good for Stacked SRAM and EEPROM cells

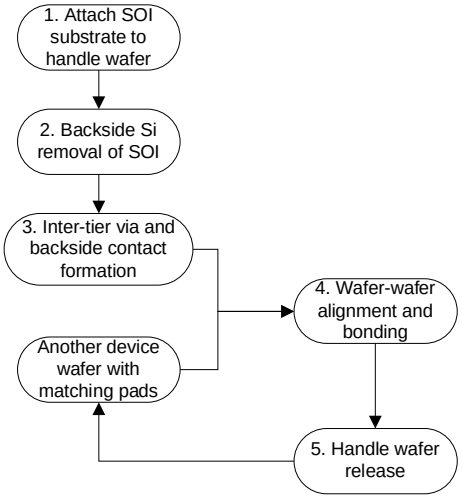

Figure 1.3: 3D wafer bonding process flowchart.

1.2.2 MIT Lincoln Lab 3D Integration Technology

The primary 3D processing technology used in this work is the MIT Lincoln Lab (MITLL) 0.18 μm 1.5 V low power fully depleted silicon on insulator (FDSOI) CMOS technology [12] [19], which has a single poly layer and three metal layers. When compared to conventional bump bond technology, this 3D technology offers better circuit-to-interconnect ratio, higher-density vertical interconnections, and reduced system power. The detailed process is briefly illustrated in Figure 1.4 [19]. In this technology, a total of three tiers, each fabricated individually, are vertically stacked. The special 3D via layers are known as 3DCUT and 3DLAND (shown in Figure 1.4 c). The 3DCUT defines the 3D via starting point on the higher numbered tier while 3DLAND indicates the 3D via stopping location on a metal layer of the lower numbered tier. When stacked, they electrically connect the separated silicon tiers.

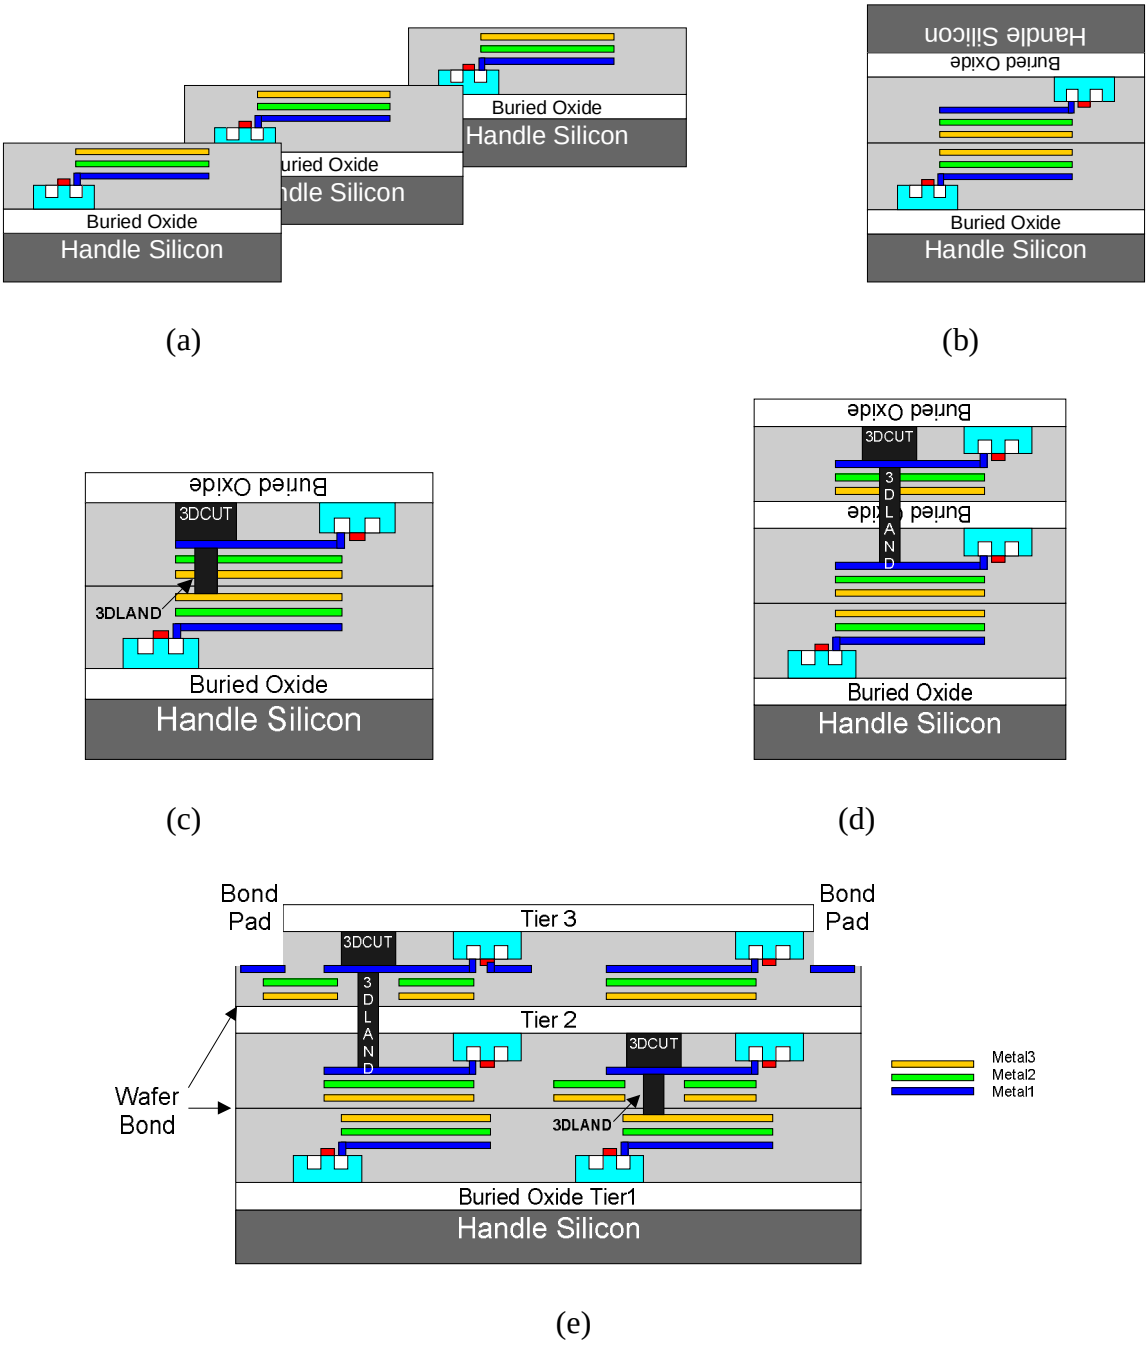

Figure 1.4: MITLL FDSOI wafer stacking example (source [19]). (a) Tier are fabricated with conventional IC processing; (b) Wafer2 is flipped and bond to wafer1; (c) Etching 3D vias, deposit tungsten (W), tier2 is electrically connected tier1; (d) Align and stack tier3, electrically connected to other tiers; (e) Three-tier assembly, bond pad etch

The basic sizing information is described here: design grid of this technology is 25 *nm*, the minimum-width poly features at 0.2 μm, the minimum metal width features at 0.25 μm with 0.35 μm spacing for all three metal layers, the routing pitch for all metal layers features at 1 μm, the contact and vias feature at $0.4 \times 0.4\ \mu m^2$, and the 3D signal vias feature at $3 \times 3\ \mu m^2$ [19].

1.2.3 Ziptronix and Tezzaron's 3D Integration Technology

Our 3DIC design flow is robust in a sense that it can adopt other wafer bonding 3D technologies besides MITLL FDSOI process. The wafer bonding technologies have similar processing steps and hence our design flow is convertible among technologies. This section will discuss two more 3D integration technologies

Unlike the SOI technology used by MIT Lincoln Lab, a bulk COMS technology would be preferred, due to the lower cost. Ziptronix and Tezzaron are the two major companies that provide 3D integration for bulk CMOS. Here we discuss the features of these technologies that differ significantly from the MITLL 3D FDSOI process.

Ziptronix's patented ZiROC and ZiCON bonding and interconnect technologies [69] provide solutions for 3D SoC design. ZiROC offers a new alternative to traditional bonding technologies, enabling room-temperature covalent bonding to occur in ambient fab conditions. This innovative process supports bonding at the wafer scale and die scale, providing a high degree of flexibility. The result of the bond is a single, uniform material whose strength typically exceeds the fracture strength of the bonded bulk materials. The

ZiROC process is specifically suited for die scale bonding. ZiCON is a room-temperature, non-adhesive bonding that enables die scale bonding and allows known-good-die selection prior to the bonding process, minimizing potential yield losses. Standard back grinding and CMP techniques are used to thin the bonded die, and a simple, five-mask-level process is used to expose the bonding points on both die, to interconnect the two chips, and to passivate the wafer.

The 3D integration technology from Tezzaron is called “FaStack” technology [70]. Tezzaron's “FaStack” technology creates fast, dense, highly integrated 3D chips. The heart of the process is wafer scale stacking. Device circuitry is divided into sections that are built onto separate wafers using standard processing. The wafers are then post-processed for thru-silicon interconnection, creating hundreds of thousands of vertical "Super-Via" (3D inter-tier via) connectors. The wafers are aligned with a precision of 0.5 micron, then bonded, thinned, and diced into individual devices. “FaStack” devices have many advantages over their single-layer counterparts. They are much more dense and their short vertical interconnects allow them to operate at higher speeds with a lower power budget. Tezzaron claims that “FaStack” devices match the tight integration of SoC devices while out-doing SiPs for high speed, low power budget, and tiny footprint.

The processing technologies for both SOI and bulk CMOS have similar steps. The first step is to fabricate the wafers separately. Then the through via pattern are formed and the wafers are bonded and thinned and finally the pad is formed. The biggest difference exists in the through via formation step. As shown in Figure 1.5, the metal can be directly deposited in the 3D via opening in SOI technology. However, a passive layer has to be deposited in the

bulk 3D via opening before the metal deposition to create isolation between metal via and silicon substrate.

A comparison of the three 3D integration technologies (MIT Lincoln Lab, Ziptronix, and Tezzaron) is listed in Table 1.3. Because of the lack of information, the Ziptronix data are estimated from their published figures. Symbol "~" means "approximately" because the data is based on manual measurement of a published figure. The tier thickness and 3D via size are of the same order of magnitude for all technologies. Generally speaking, the SOI process usually consumes smaller power but has poorer thermal conductivity. In this dissertation, we will focus on the MITLL process because of the detailed technology information that was available to us.

Table 1.3: 3D integration feature size comparison.

	MIT Lincoln	Ziptronix	Tezzaron
Tier thickness	6.64 – 7.59 μm	~ 9 μm	13 μm
Bonding material	SiO_2	Si	Si
Bonding material thickness	400 nm	~ 5 - 6 μm	5.5 μm
3D via size	3 ×3 μm^2	~ 4 ×4 μm^2	2 ×2 - 4 ×4 μm^2

Although our work focuses on MITLL process, it is compatible to other technologies. The results obtained from MIT process can be easily re-targeted for a different technology using the details shown in Table 1.3. To adopt other processing technologies, three modifications are required. The first is to modify the dimension of 3D through-via and its design constraints like spacing and enclosure requirements. The second is to develop a new standard cell library base on the select technology along with a Library Exchange Format

(LEF) file that contains the design rules and technology information. The last is to update the temperature dependency parameters that will be discussed in section 3.6.

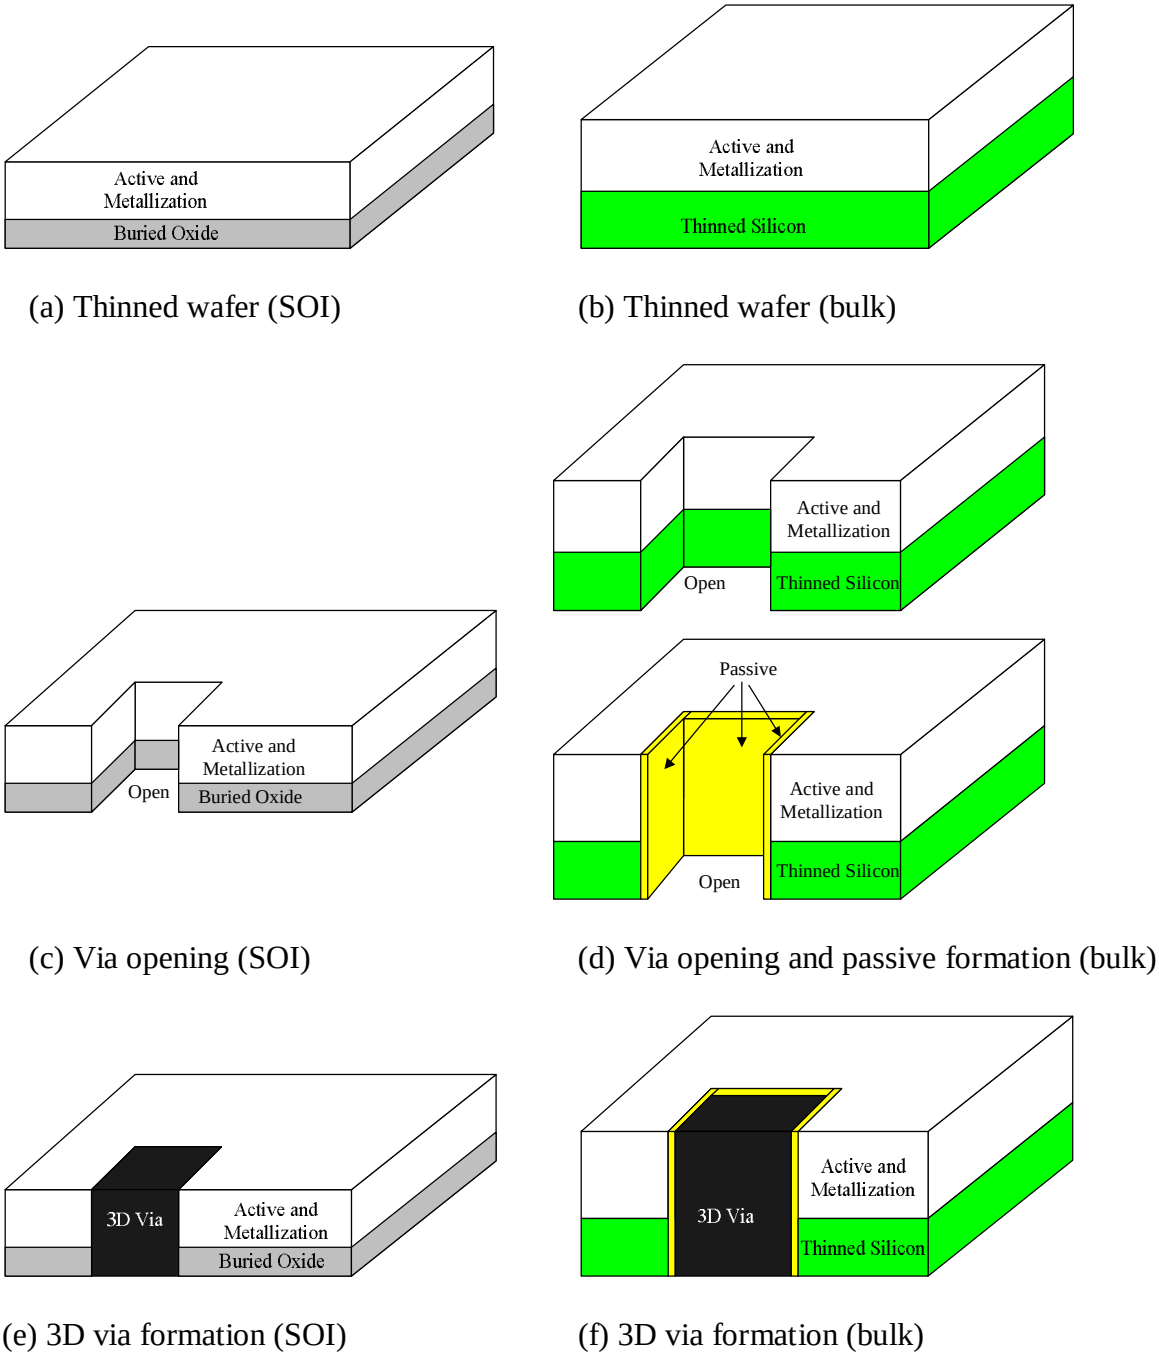

(a) Thinned wafer (SOI)

(b) Thinned wafer (bulk)

(c) Via opening (SOI)

(d) Via opening and passive formation (bulk)

(e) 3D via formation (SOI)

(f) 3D via formation (bulk)

Figure 1.5: 3D via formation in SOI and bulk 3D integration technology.

1.3 Goal of This Work

There are two major objectives of this work. The first objective aims to achieve 3DIC design automation using a combined CAD tools and algorithms. This 3DIC design flow seamlessly integrated various commercial tools with Python [38] and Tcl [39] scripts [45]. Although some potential performance is lost, it is relatively easier to explore the feasibility and advantages of 3DICs with this approach than to develop new 3D tools from scratch. This design flow provides a platform for designing 3DICs and ensures that the user can have a verified design to be fabricated. This work also demonstrates the efficiency of designing innovative 3DICs by reusing tools with a proper design flow.

The second objective of this work is to explore design trade-offs and circuit reliability among different applications. A direct solution to this problem is to create design guidelines for both low-power and high-performance 3D applications. The guidelines will give designers information regarding how many tiers and how many thermal vias are desirable once the specification is determined. However, the optimal tier and thermal via number for different designs would vary as power density and wiring complexity changes. Prototyping with a high level physical-design flow is an efficient way to obtain the guidelines. We change number of silicon tiers, number of metal layers and area of thermal vias to get different values of temperature and system performance for both applications. These values are presented as plots and tables to provide a set of trade-off selections. The thermal problem is fundamental in 3DIC as high temperature has a dramatic impact on device speed and

leakage power [35]. This impact is especially critical for clock trees, because clock branches experiencing a temperature gradient across tiers will have significant insertion delay and transition time variation in high-performance applications. To facilitate this trade-off exploration of 3D integration, we used the automated design flow mentioned above to complete 3DIC physical design and performance analysis.

An exhaustive comparison on all circuits is the best to achieve our goal but seems impossible. We select the benchmark circuits from two different kinds of applications: low-power and high-performance. An eight-point winograd Fast Fourier Transform (FFT) is implemented to explore the trade-offs in low-power applications. 2D implementation of this FFT runs at approximately 40MHz with 810mW. An Open Risk CPU core (OR1200) with 5 SRAMs is implemented to explore the trade-offs in high-performance applications. This design runs at 60MHz with 3.3W in a 2D implementation. We selected these two designs as a starting point because they have similar speed and chip area, but significantly different power consumption and placement restrictions. These two designs are good representatives of current digital designs.

1.4 Thesis Organization

Chapter 2 outlines the current state of the art in 3DIC research and applications.

Chapter 3 describes the 3DIC design flow and breaks it down to discuss details of every step. It begins by comparing the conventional ASIC design flow with our 3DIC flow to

show the compatibility. Next, it approaches how the flow is assembled by explaining sub-steps of the flow and the models employed.

Chapter 4 discusses the thermal issue in 3DICs. Heat flow is the one of the most fundamental problems in 3DIC and limits the integration level a 3D system can achieve. Thermal vias are used in many 3DICs to remove heat. This chapter gives the detail of our thermal model, via pattern generation, and thermal reliability analysis.

Chapter 5 studies the performance improvement of 3DICs in the original and extended MIT Lincoln Lab FDSOI 3DIC processing technology. Extensive comparison on wire-length, timing, power, and temperature are performed. The design trade-offs among timing, power, and temperature are also explored in this chapter.

Chapter 6 gives the conclusion of this work and discusses the future direction.

Appendix A - Appendix C gives the necessary supplemental materials to better understand the development and implementation of this 3DIC design flow.

Chapter 2.

Start of the Art 3DIC Tools and Progress

2.1 Summary

3D digital IC has been widely studied due to its potential to fill the gap between device scaling and interconnect scaling. The major advantages of 3DIC are the reduction of interconnect wirelength and chip area, which boost the performance of 3DIC. Figure 2.1 shows the complex interaction between the cause and effect, where the symbol "↓" means reduction. However, non-idealities will significantly cancel the advantages brought by 3DICs, and further study is required to investigate the 3D integration.

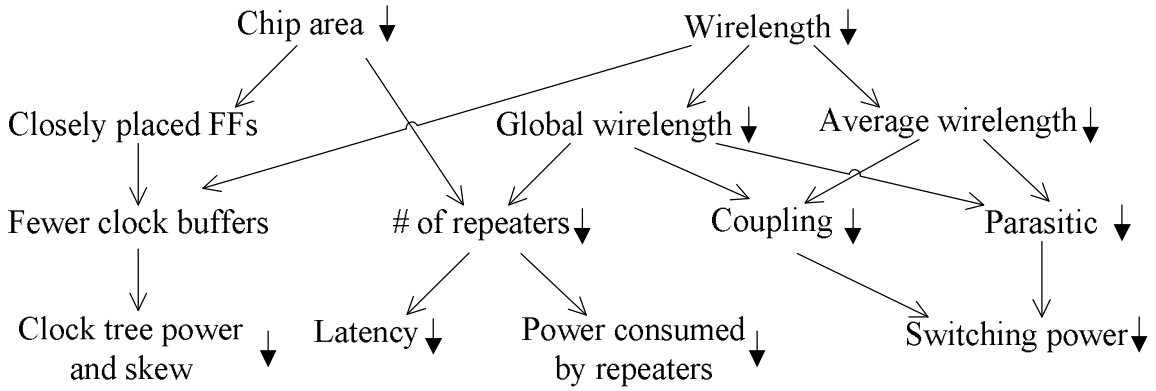

Figure 2.1: Performance boost in 3DIC.

Many researches have been conducted explorations of the advantage of 3D integration. In this chapter, previous theoretical and practical works about 3DIC are outlined to show the current picture of 3DIC research.

2.2 Empirical Wirelength Reduction Analysis of 3DICs

E. F. Rent of IBM wrote two memoranda in 1960 that later developed into well known Rent's rule. Rent's rule was applied by researchers and designers to derive the interconnect wirelength distribution. Rent's rule is written as (1) [15] [31], where T, N, α and p represent the number of signal inputs/outputs, the number of gates, the average number of fan-out per gate, and the degree of wiring complexity respectively. The wirelength distribution can be described by (2) [15]. $I(l)$ means the total number of interconnects that have length less than or equal to l. Since the Rent's rule is an empirical result obtained by observing existing designs, it is useful to apply the measured parameters of Rent's rule to a similar architecture but would be misleading to apply the parameters to a different architecture. The system architecture becomes the key parameter that affects the accuracy of Rent's rule. The Rent constants for various architectures are shown in Table 2.1 [73]. These constants act as guidance during our partition decision described in section 3.2.

$$T = aN^{p} \tag{1}$$

$$I(l) = \int_{1}^{l} I(x)dx \tag{2}$$

Table 2.1: Rent constants for various architectures.

Architecture type		p	α
Static Memory		0.12	6
Microprocessor		0.45	0.82
Gate Array		0.50	1.9
High-speed Computer	Chip level	0.63	1.4
	Board level	0.25	82

The performance of a conventional system can be roughly evaluated according to the wirelength estimation based on the Rent's rule, which is also true for 3DICs except that the estimation should be based on the 3D Rent's rule. Extending Rent's rule into a 3D structure with arbitrary number of active layers *m*, and supposing that *N* gates are evenly partitioned into *m* layers, the wirelength distribution is rewritten as (3) – (6) [32] [33], where α is related to the average fan-out which is defined by (6), ***V(d)*** is the number of connections with vertical distance of *d* device layers (d=1,2,……, m-1).

Horizontal wirelength distribution:

$$I(l)=\begin{cases}\Gamma(\frac{l^3}{3}-2l^2\sqrt{\frac{N}{m}}+2l\frac{N}{m})l^{2p-4}, & 1\le l<\sqrt{\frac{N}{m}};\\ \frac{\Gamma}{3}(2\sqrt{\frac{N}{m}}-l)^3 l^{2p-4}, & \sqrt{\frac{N}{m}}\le l\le 2\sqrt{\frac{N}{m}};\end{cases} \tag{3}$$

$$\Gamma=\frac{a\,km^{p}\,\frac{N}{m}(1-(\frac{N}{m})^{p-1})}{-(\frac{N}{m})^{p}\,\frac{1+2p-2^{2p-1}}{p(2p-1)(p-1)(2p-3)}-\frac{1}{6p}+\frac{2\sqrt{\frac{N}{m}}}{2p-1}-\frac{\frac{N}{m}}{p-1}} \tag{4}$$

Vertical wirelength distribution:

$$V(d)=\frac{2akN(1-N^{p-1}-m^{p-2}+m^{-1}N^{p-1})}{m(m-1)}(m-d) \tag{5}$$

$$a=\frac{f.o.}{1+f.o.} \tag{6}$$

Based on the 3D Rent's rule, R. Zhang *et al* [32] conducted a performance study on delay and power consumption of 3DICs using SGSOI (single gate SOI) and DGSOI (double gate SOI) device structure. To simplify the exploration, their work assumed each gate is a two-input NAND gate. The exponential Rent constant *p* is set to an arbitrary value of 0.45,

which implies a microprocessor. However, due to the arbitrary nature of value selection, the conclusion is not independent of the value of p [32]. This setup showed a maximum of 40% reduction for the longest wirelength and a 30% reduction for the average wirelength with 5 tiers at the 180 *nm* technology node, regardless of the technology used in their study (either SGSOI or DGSOI). They observed a dramatic improvement on circuit performance because 3D integration effectively reduced the number of long delay nets and repeaters. Compared to conventional IC circuits, the 3D integration is therefore predicted to achieve 2-3 technology generations of performance (timing) advantage. Furthermore, the DGSOI circuit can be clocked at 13%-20% higher speed while consuming 5% more power than the corresponding SGSOI circuit. However, the power delay product of DGSOI circuit is 8% lower than SGSOI circuit, which makes the DGSOI device more energy efficient.

R. Zhang *et al* [33] also studied the performance trend of 3DIC for the future technology generations based on a heuristic study of interconnect wirelength. Using the same simplification of the two-input NAND gate, this study showed that the total interconnect capacitance and power consumption went down first and then rose again due to a tremendous increase in vertical wirelength and the routing blockages caused by the vertical 3D vias as the number of tiers continued to increase. However, the worst case clock rate continued to speed up, even though the total capacitance increased.

In today's interconnect driven VLSIs, the system performance (timing) is usually determined by long nets. 3D interconnection reduces the number and wirelength of long nets, which directly benefits the global clock rate. J. W. Joyner *et al* [72] explored the global wirelength distribution rather than all the nets in a certain design, because they believe the

3D integration was expected to significantly reduce the longest interconnects while providing a lesser advantage for average interconnects. All the *m* modules are divided into *n* tiers and the global net distribution of each tier was calculated. With a certain netlist that provided the number of *N* nets, placement and routing models were created to obtain the layout information. The placement information described the average dimensions of the bounding area of a net connecting a group of modules, which predicted the edge length of a net. The routing information provided a distribution of the length of a net for a given number of terminals and net-bounding area. After that, a minimum rectilinear Steiner tree (MRST) model was used to calculate the minimum total wirelength of either a single-tier or a multi-tier net. With the models created in this study, the resulting distribution showed 3D integration reduced the wirelength as the square root of the number of tier. In addition, the maximum global clock frequency is shown to increase as T^2, where T is the number of tiers.

G. Chandra *et al* [8] evaluated the interconnect performance of 2D and 3D processors with memory due to the huge demand for high performance microprocessors. Their methodology started with an area evaluation of memory and logic followed by wirelength estimation. Since the memory and random logic have different wiring complexity, it is more desirable to study their wirelength distribution separately. With the significant amount of memory on chip, 3D integration leads to a number of possibilities for the partition. Figure 2.2 shows the possible two topologies. Figure 2.2 (a) put memory and logic on different tiers. Since the memory consumes less power, it was put on upper tiers. Figure 2.2 (b) split the memory and logic equally and then put the split parts to different tiers. With the reduction of chip area, the delay of the longest interconnect spanning the logic portion also reduced and

hence lead to significant performance advantages. Because of this reason, topology shown in Figure 2.2 (a) does not lead to a large reduction in interconnect length and hence less delay reduction.

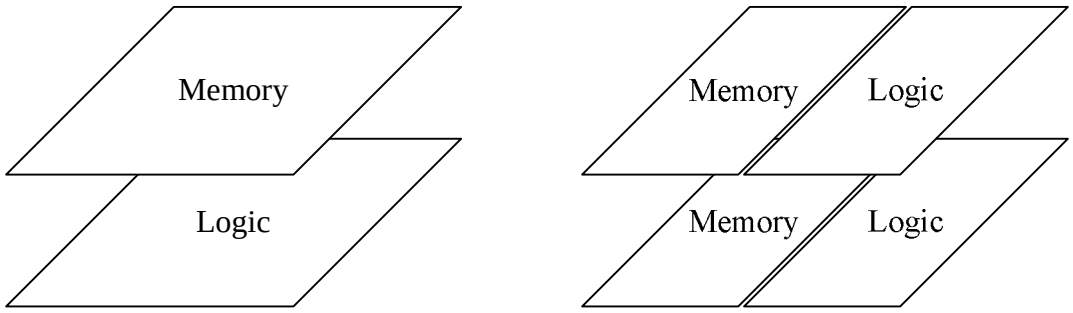

Figure 2.2: Possible 3D partition topology: (a). Memory and logic are put on separate tiers; (b). Memory and logic are split equally to put on different tiers.

The above studies were limited in a sense that temperature was not considered one of the major limiting factors in 3DICs. Another restriction is that they did not consider the routability issue in a real design. However, the wirelength distribution investigation of their studies was the base of any theoretical analysis of 3DICs. In our case, equations (3) – (6) act as partition criterions in our partitioner (discussed in section 3.2).

2.3 Computer Aided Design (CAD) Tools

Much research has focused on 3D CAD tools to facilitate the 3DIC design. These studies have shown that the performance improvements of 3DIC are expected. The point-tools have already been widely researched and will be discussed in this section. However, lack of a complete flow to design 3D VLSI makes the result of point-tool less persuasive.

Floorplanning is vital for high level temperature reduction because hot blocks can be placed away from each other. The problem caused by this approach is that wirelength may be increased so that both timing and power degrade. Heat dissipation is the most critical challenge in 3DICs. J. Cong *et al* [13] developed an efficient thermal-drive floorplanning algorithm for 3DICs. Their algorithm aims to minimize chip area, wirelength, and number of inter-tier vias, while taking temperature into consideration in both vertical and lateral directions. They defined seven operations to move a block to a proper location: (1). Rotation, which rotates a block; (2). Swap, which swaps two blocks in one tier; (3). Reverse, which exchanges the relative position of two blocks in one tier; (4). Move, which moves a block from one side of a block to another side; (5). Inter-tier swap, which swaps two blocks at different tiers; (6). Z-neighbor swap, which swaps two blocks at different tiers and are close to each other; and (7). Z-neighbor move, which moves a block to a position at another tier and are close to the current position. With these seven operations, a simulated annealing technique is employed to find the best trade-offs. Every time a floorplan was generated, a weight cost of the optimization objectives and the constraints would be evaluated. The cost function was written as (7), where *n_wl*, *n_area*, and *n_vc* were normalized wirelength, chip area, and inter-tier via count; c_T was the cost of a certain temperature; α, β, γ, η were the predetermined weight. The simulated annealing technique was applied to multiple thermal models using compact resistive network (CBA-T), closed-form equations (CBA-T-fast), and a hybrid of them (CBA-T-hybrid). These three models traded off runtime and quality. Compared to other published results, their floorplan result achieved a 5.5% improvement in wirelength with GSRC benchmarks. Their most accurate CBA-T model reduced temperature by 56% in average with a 21% increase in chip area, a comparable wirelength and inter-tier

via number. They also achieved a 40% improvement on average temperature with CBA-T-fast model and a 50% improvement with CBA-T-hybrid model, which trades off between runtime and quality.

$$cost = a \cdot n\_wl + b \cdot n\_area + g \cdot n\_vc + h \cdot c_T \quad (7)$$

Cong's work demonstrated that the maximum on-chip temperature can be effectively controlled through their thermal-driven floorplan algorithm. The white space or thermal via insertion can guarantee the routability issue and satisfaction of temperature constraints. However, a lack of performance analysis makes the floorplan algorithm somewhat insufficient.

3D place and route tools were widely studied to investigate the impact of 3D integration technology. S. Das *et al* [22] discussed a 3D standard-cell-based placement tool, global routing tool, and layout editor. The placer recursively partitions the blocks into halves with a pre-set tolerance, searching for a min-cut partition (cut in z-dimension means an inter-tier via) between these halves. Their placer trades off increased total wirelength for fewer inter-tier vias. A concurrent global router following the placement performs a global routing. It determines the location and quantity of inter-tier vias in addition to route the wires on each tier. With the tools described above, the authors test the placer and router in two configurations: (1). optimize for total wirelength; and (2). optimize for least number of inter-tier vias. Their results showed significant reduction in terms of wirelength in 3D designs. If a total of 3 tiers are available, the maximum wirelength reduction was about 40% with configuration 1 and 10% with configuration 2 respectively. The final layout was edited in an extension of Magic to design 3DICs. The circuits were first designed separately in Magic

and then bonded together to form a single entity. The 3D Magic editor facilitated low-level design of multi-wafer integrated circuits as well as management of hierarchical 3D designs. Their further study [23] gave an extensive comparison on 3DICs using timing, energy, and thermally driven algorithms respectively. Their results with a FFT datapath and Data Encryption Standard (DES) chip showed the energy dissipation can be reduced by 24% to 40% using two to five tiers when compared to a 2D counterpart. And an energy driven place and route algorithm resulted in smaller energy-delay product. Their thermally driven algorithm achieved much lower temperature with same speed, which agrees with the finding from [6]. In both cases, there was a trade-off between energy performance and thermal issue.

Other placement tools have been developed to address thermal problem besides the concern of wirelength [9] [18] [44]. Balakrishnan *et al* [44] developed a wire congestion and thermally aware 3D global placer. They used a multilevel min-cut algorithm to minimize the congestion and power dissipation within confined areas. After an initial solution was obtained, a simulated annealing based technique was performed to minimize congestion from global wire as well as to improve temperature distribution. In a congestion-driven case, their results showed that temperature became a problem, while in a power-driven case, routability became a problem. A trade-off among wirelength, wire congestion, power and temperature must be performed to get a good result. B. Goplen *et al* [9] developed a placer for standard cells using a force directed approach. In their approach, an extra thermal force directs cells away from hot areas while wirelength is minimally affected. Temperature of certain area is determined by the surrounding points. If temperature at some position is high, the thermal repelling force is also considerably high that no cells can be placed there. With their method,

a 12% reduction in maximum temperature, 1.3% reduction in average temperature and 17% reduction of thermal gradient were achieved, but there was a 5.5% increase in total wirelength compared to the same placer that does not consider thermal effects.

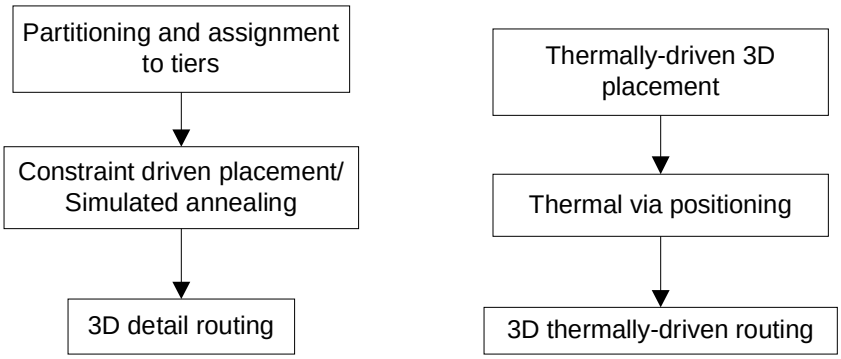

(a) FPGA style design flow (b) ASIC style design flow

Figure 2.3: Design style for FPGA and ASIC in 3D analysis and design optimization.

A combination of placement and routing algorithm for 3D FPGA and ASIC was developed by C. Ababei *et al* [71], which addressed wirelength, delay, area minimization, as well as thermal optimization during place and route phases. The FPGA and ASIC style design flows are shown in Figure 2.3. The FPGA style design started by partitioning the circuits among tiers using *hMetis*, then the tier specific circuits were mapped to the 3D FPGA accordingly. Each FPGA tier was placed with a simulated annealing technique and finally detail routed. Results showed a 20% improvement on delay and wirelength compared to its 2D counterpart. The ASIC style design flow started with a fast thermal analysis of 3DICs, and followed by a force directed placement with thermal constraints. Thermal vias are interested to the pre-reserved thermal via region after the placement. With this thermal via insertion approach, an improvement in the average temperature of 30% was observed. When

thermal optimization is done, the entire circuit is detail routed as a whole. After routing, the maximum observed timing improvement approached 30%.

An automated design method was also proposed for standard cell based 3DICs [46]. In that work, the researchers developed 3D placement and global routing tools, and used MAGIC as the layout editor. The wire-length and performance characteristics of 3DIC were analyzed and shown to be attractive. However, their research showed that if heat cannot be removed from 3DIC efficiently, the performance will degrade before more tiers can be stacked.

2.4 Heat Removal

Heat is considered the most fundamental problem in 3D integration. Thermal vias are widely used to remove heat in 3DICs. A. Rahman ***et al*** [3] developed simple power consumption model and boundary conditions for a thermal analysis. Though their simulation results showed that there was a maximum of 35%-45% power reduction associated with normal signal interconnects and 30%-35% power reduction associated with clock network, thermal is still a problem due to higher power density. A possible solution to remove heat is to insert thermal vias in the 3D system. The thermal resistance in the vertical direction is strongly dependent on each silicon layer's thickness and underlying bonding material. Thermal vias provide extra heat removal paths and hence significantly reduce the equivalent thermal resistance. Their numerical simulations showed that inserting thermal vias for about

4% of the total die area could achieve two to three times reduction in effective thermal resistance.

B. Goplen *et al* [5] used Finite Element Analysis (FEA) to calculate temperature. To simplify their computation, they used pre-reserved regions that were uniformly placed within standard cell rows for thermal via insertion. These regions were blockages during the placement phase. After placement, thermal vias are inserted to the pre-reserved regions to reduce temperature. The entire chip was divided into sub-blocks and power was randomly generated for each block. The thermal via density of each thermal region is calculated by their FEA method. With a pre-set maximum temperature target, the thermal via density in each pre-reserved region was updated iteratively until convergence. Their experimental results showed that both temperature and thermal gradients are reduced significantly after thermal via insertion. Besides this, their results also showed that lateral thermal conductivity is much lower than vertical thermal conductivity.

Current studies showed that the use of thermal vias is a simple yet efficient method of heat removal. Thermal vias reduce both vertical and lateral temperature gradient, and hence we will use this technique in our work to remove heat as well.

2.5 3D Integration Examples

With the extensive effort spent on the development of 3D integration, 3DIC has recently become feasible. Researchers from MIT Lincoln Lab fabricated a 64×64 active pixel sensor array circuit [12] with their SOI 3D processing technology. Each pixel was composed of a

photodiode on one wafer and an analog to digital converter (ADC) on the other wafer joined by a through via. The construction of these circuits consisted of bonding and interconnecting a SOI wafer with imaging circuits and inverters to a SOI wafer with ADC circuits (illustrated in Figure 2.4). This work showed the feasibility of stacking SOI circuits to build 3DICs with dense vertical interconnects (through via). However, this study did not show a performance comparison between 3D and 2D integration.

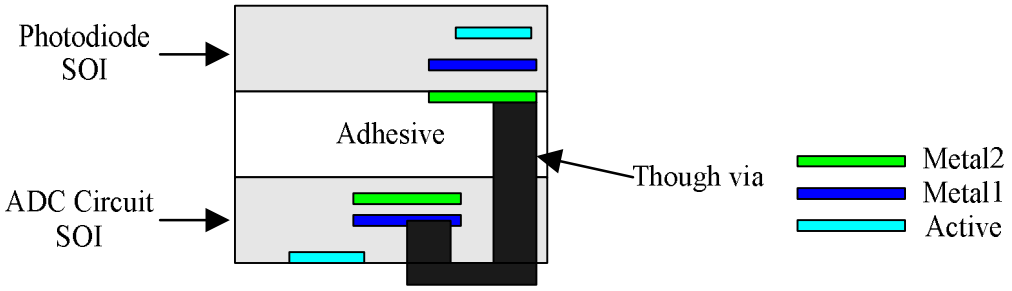

Figure 2.4: Illustration of the 3D pixel sensor array.

Prior to this work, the best comparison of 2D vs. 3D performance in a working example comes from B. Black *et al* [4], in which the authors fabricated a microprocessor chip to evaluate the impact of 3D processing technology. The chip was fabricated on 2 tiers and then bonded together face to face. They repartitioned the iA32 processor into a 3D topology such that cache memory and functional units are split. In other words, one tier has the memory while the other has the functional units of the original design. They have pointed out that folding (a strategy similar to Figure 2.2 a) large blocks is preferable to stacking (a strategy similar to Figure 2.2 b) because internal wire and latencies can be eliminated. Besides, folding large blocks also reduces the distance traveled by external signals that are crossing the block. In their example, the worst case path with 2D floorplan was when the load must travel from the far edge of the data cache and across the data cache to the farthest

functional unit. In the 3D floorplan, cache memory and function units overlap so that the load data only traveled to the center of the data cache. The same worst case path contained half as much routing distance. Since data was only traversing half of the data cache and half of the functional units, 1 clock cycle of delay was eliminated in the load execution delay. Performance improvement of the 3D iA32 microprocessor was obvious. The clock delay of store retirement was reduced by 30%, Floating Point load latency was reduced by 35%, register file read was reduced by 25%, retirement and de-allocation were reduced by 20% and other important latencies were also improved. A total of 25% pipe stages were eliminated by the 3D floorplan. In general, a 15% performance improvement was achieved by eliminating piped wire stages, reducing delay between blocks, and eliminating wire within blocks. In terms of power, a 15% total power reduction was achieved by eliminating 50% of repeaters and 50% reduction in the clock wire. The thermal problem was studied in their work. A naïve floorplan can increase heat by 10-15%. However, the thermal problem was found to be addressable at block level. In their design, the "hot" areas of the die are due to several blocks that utilize dynamic circuits to meet timing requirements. Splitting the hot blocks across two tiers could reduce internal wire delay sufficiently to allow a relaxation in the implementation of the power inefficient circuits. The power consumption of those fast hot blocks was expected to have a 50% reduction if they can be properly folded.

However, B. Black's work only worked on integrating two tiers and was insufficient to explore the full advantages brought by 3D integration. Better performance is expected by theoretical analysis of wirelength reduction. To explore this area more thoroughly, our design flow supports up to 10 tiers.

2.6 Unsolved Problems

Though previous researches regarding 3DICs have been proven attractive, there are still a set of unresolved problems. These problems are described below.

First, a design methodology that can assemble the tools to fabricate real 3D chips is currently not available. We developed verification models and an automated design flow to address these problems. The models and design flow are discussed in Chapter 3.

Second, 3D clock tree synthesis (CTS) is not proposed by previous research. A modern IC design will not function properly without a carefully designed clock tree. This problem becomes more critical in 3DICs because clock buffers may be operating at significant different temperatures. The synthesis and verification schemes of 3D CTS need to be developed to successfully implement the clock tree in any 3DIC.

Third, though 3DICs can significantly improve system performance, how to achieve most out of 3D integration is unclear due to the inevitable electro-thermal delay-power coupling. The increased temperature causes longer delay. If the system is running at its highest frequency ($f = 1/delay$), the higher temperature reduces dynamic power but increases leakage power. Both dynamic and leakage power will in turn change the system temperature profile. Besides this coupling, the effect of thermal vias is unclear because they may reduce temperature but may also increase wire-length. How the 3D integration and thermal via will benefit our system remains unknown. As more tiers are integrated, heat

dissipation becomes prominent. The trade-off between system performance and thermal issue must be identified.

The major features of this work are listed in Table 2.2. It tells what are provided by this work but not by previous researches. These features aims to solve the unresolved problems mentioned above.

Table 2.2: Features provided by this work.

Previous Research	**This Work**
Conventional methodology handles 2DICs	This methodology handles 3DICs
Temperature range of for conventional 2DICs thermal analysis is [60°C, 110°C]	Temperature range of 3DIC thermal analysis is [0°C, 250°C]
3D tools are developed but are not applied to tape out 3DICs	This methodology is capable of taping out 3DICs
No 3D Clock tree synthesis tool or methodology has ever been reported	This methodology can handle 3D clock tree synthesis and verification
No parasitic extraction tool has been developed for 3DICs	Our methodology is able to create parasitic deck for 3DICs

Chapter 3.

Design and Verification Flow for 3DICs

3.1 Overview

Because 3DIC design methodologies are relatively new, there is not any standard design flow in this area. Researchers have investigated many aspects of 3D integration such as floorplanning [13], placement [9] and routing [22]. These academic tools are not verified and would require significant effort if they were used to design real 3D system. Instead of spending huge amount of time to develop and verify new tools ourselves, existing CAD tools are utilized to assemble an efficient yet reliable flow for 3DIC design. With this flow, designing 3DICs becomes feasible and their performance can be evaluated and compared to their 2D counterparts. This flow also minimizes format exchanges by using standard input/output file formats. Therefore, it will be easy to incorporate new tools into this flow as long as these new tools use standard file formats.

The conventional IC design flow (shown in Figure 3.1 a) has been a great success in facilitating the IC design and manufacture. At the Register Transistor-Level (RTL) design step, the system is implemented in RTL code (usually Verilog or VHDL) and verified to match the functionality. In the logic-synthesis step, RTL code is mapped to real circuits. In this step, designers will compare different mapping methods and select the one that results in optimal performance (usually in terms of critical path delay or power consumption). The output of the

synthesis is a gate level netlist. The physical designers take the gate level netlist and use place and route (P&R) tools to generate layout mask patterns for fabrication while verifying the performance constraints and design rules are met [7]. The flow was developed to help designers to avoid excessive interaction between different groups by clearly defining group tasks.

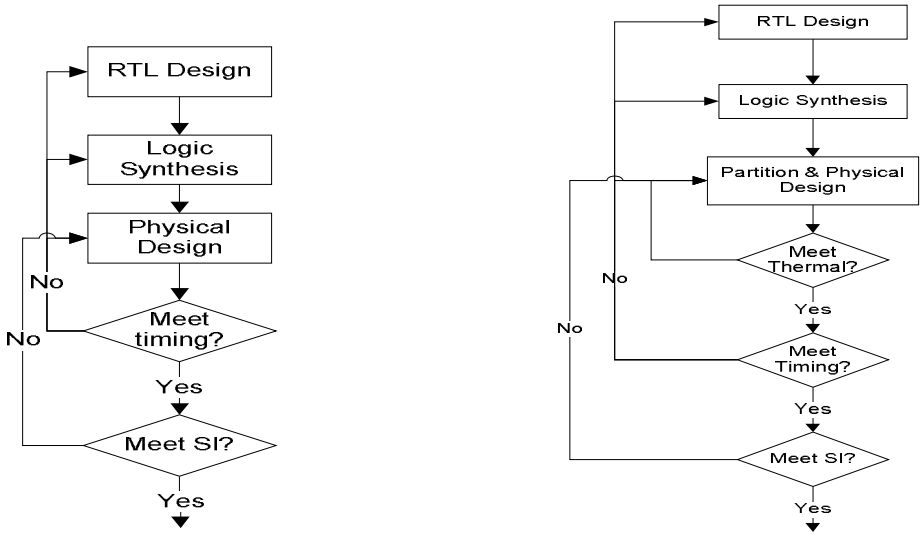

(a) 2DIC design flow (b) NCSU 3DIC design flow

Figure 3.1: Comparison between conventional 2DIC design flow and our 3DIC design flow.

Since the conventional 2DIC design flow has been successfully used for many years and is widely accepted, our 3DIC design flow is designed to be as close to the conventional flow as possible. The block diagram of our 3DIC design flow is shown in Figure 3.1 (b). The only modification was the physical design flow. To facilitate physical design, partitionning is executed prior to detailed physical design. The physical design flow is similar to the

conventional flow, too, except that there are more design constraints. Another difference is that a thermal check is implemented to verify the performance and reliability of 3DICs.

An expanded top level design flow is shown in Figure 3.2 (modified from [40]). The basic job of this flow is to invoke commercial and self-developed tools/scripts in the correct order to complete 3DIC design and verification. The initial flow is based on the three-tier, three-metal-per-tier, 180*nm* FDSOI technology from MIT Lincoln Lab [19] discussed in section 1.2.2. As shown in the figure, the flow begins with the netlist result from standard-cell synthesis, from which the block level area, connection, and power information is obtained. Based on this information, our partitioner partitions the design for minimum inter-tier cuts using *k-METIS* [2]. The design is then floorplanned, using automated or manual module placement and power-planning in the commercial 2DIC tool *Encounter* from *Cadence* [56]. A manual adjustment will be enforced if two hot blocks are stacked on top of each other. Next, thermal-design is performed, which involves even insertion of a pre-determined number of thermal-vias into the chip. To facilitate the design flow, the thermal via is implemented as a standard cell. The standard cells are then placed in each tier independently using *Encounter*, followed by an inter-tier via alignment step to ensure consistent via positions on each tier. The clock-tree is then inserted, after which the thermal via cells are relocated to eliminate any possible hotspots in the design. Then the design is routed, and RC parasitics are extracted for each tier. These parasitics are then merged into a single Standard Parasitic Exchange Format (SPEF) file, inserting a value for each inter-tier signal via determined from 3D field-solver simulations. However, the thermal via are not involved in the parasitic extraction. The delays and power are then analyzed with nominal temperature using the cell-based analysis tools *PrimeTime* and

PowerCompiler from *Synopsys* and switching-activity annotations from Verilog simulation. Lastly, an electro-thermal coupling analysis was executed to determine the actual performance and determine if any design iteration will be invoked based on weather or not there exists any performance violation. The final flow supports place-and-route for up to 10 tiers and 7 layers of metal per tier.

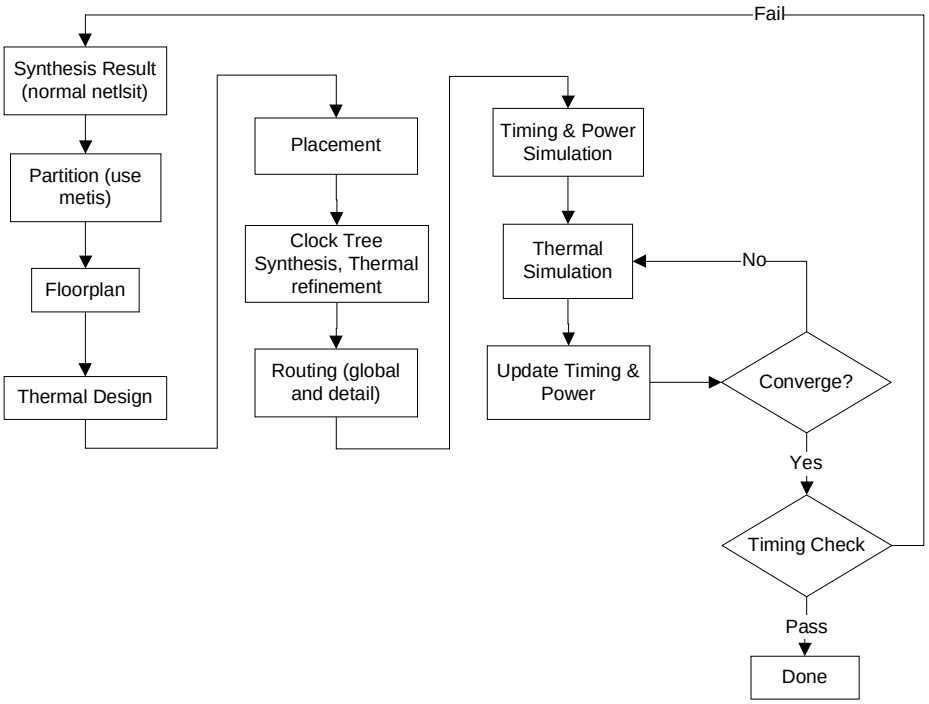

Figure 3.2: 3DIC physical design flow.

An electro-thermal delay-power coupling analysis is definitely necessary in 3DIC because the power, timing and temperature cannot be determined independently especially when temperature is high. Instead, power, timing and temperature are co-variables and will be iteratively simulated and calculated until they converge on a final solution. When the temperature and performance converge and design passes final constraints check, DRC and LVS checks will be invoked to assure no logical or fabrication problems will occur. However,

DRC and LVS check are split from the automatic design flow because excessive manual fixes are required whenever an error occurs, which makes it hard to be automated. In the following sections of this chapter, each design step will be discussed in detail.

3.2 Circuit Partitioning

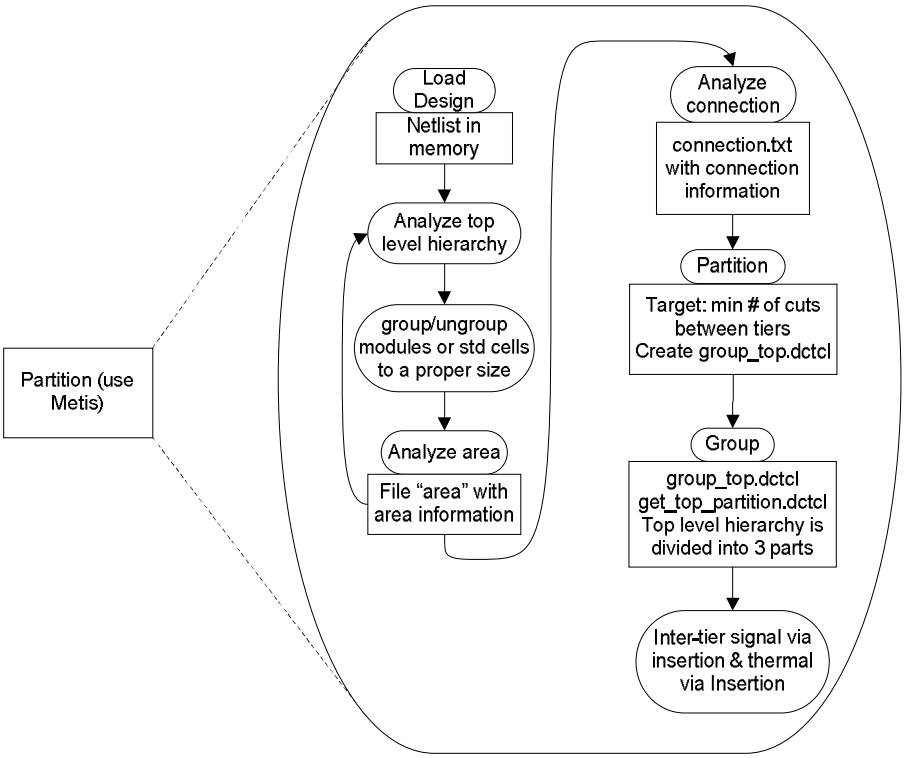

Figure 3.3: Partition flow.

The purpose of the circuit partitioning step is to partition the netlist into multiple tiers for physical design. The first step of our 3DIC design flow is the module based circuit partitioning. The circuit partition flow is shown in Figure 3.3 and is implemented in Tcl within *Design Compiler* [55]. There are total eight separate steps in this flow. Step1: load design; step2:

analyze top level hierarchy; step3: group/ungroup modules or standard cells; step4: analyze module area; step5: analyze connection; step6: circuit partition; step7: grouping different modules into partitions; and step8: insert inter-tier signal via and thermal via.

Once the design is loaded into memory, this partitioner will analyze the top level hierarchy. The purpose of step2, step3, and step4 is to create groups of cells in the netlist that have similar area and can there be floorplanned conveniently. Step2 analyzes the top level module size and sparse standard cell count. Huge module should be avoided because it will destroy the balance of cell area on different tiers, while large number of sparse standard cells significantly increases runtime. Hence, a grouping and/or ungrouping is employed in step3. Grouping is a manual process of combining some standard cells and ungrouping other units of hierarchy until all modules in the top-level of hierarchy have similar area. As a rule of thumb, the difference in area between the largest and smallest modules should be around 5. Exception occurs only if the original top-level hierarchy contains some very small modules. Iterations between step2 and step4 could occur until we got a proper top level hierarchy. After the proper hierarchy is obtained, area is annotated to each module and this information is saved as a text file with name "area". Then step5 is invoked to analyze the connections among the top-level modules and standard cells. This information is saved to another text file "connection.txt". Then the partition step is executed based on the area and connection information. The goal of the partition step is to minimize the total number of cuts (in this case, one cut stands for one inter-tier signal via) and balance the standard cell area of each tier so that we do not waste silicon resources. This partitioner analyzes file "area" and "connection.txt" obtained in step4 and step5, and then translates them into a graph that the graph partitioning tool ***k-METIS*** understands. The

area information is used to balance the cell area in different tiers and the connection information is used to minimize the cuts between tiers. The partitioner calls *k-METIS* to create an N-way partition with minimized cuts and analyzes *k-METIS*' output. Then a file named "group_top.dctcl" is created (step6), which contains *Design Compiler* commands to create the necessary groups. The tier that consumes the most power should be assigned closer to the heat sink and subsequent tiers should be ordered from hottest to coolest, going away from the heat sink. File "group_top.dctcl" is executed in step7. Also in this step, the Tcl file "get_top_parititon.dctcl" is executed to traverse all the sparse standard cells and analyzes connections of their inputs and outputs. If a cell is connected to a module or a standard cell on certain tier, this standard cell is assigned to that tier too. If a cell is connected to multiple tiers, then the cell will be assigned to the tier where its output pin is connected to.

When the partition is done, an inter-tier signal via insertion procedure is invoked. The partitioner analyzes all nets in the top level hierarchy. Any net that spans more than one tier (group) is broken down and inter-tier signal vias are inserted to connect its branches on different tiers. Each branch is assigned a suffix with "_tier#", where # represents the tier number or tier name. The original nets are removed from the netlist. If the net connects multiple tiers or goes through certain intermediate tiers without any connection, through vias will be inserted. Thermal via insertion follows the signal via insertion. Currently, the number of thermal vias on each tier is predefined and identical for all tiers. After all vias are inserted, tier specific netlists are generated.

A partition criterion is studied based on equations (3) – (6). This criterion determines if a module should be partitioned or not based on wirelength reduction analysis. The benchmark

circuits used in this study are a microprocessor and a logic circuit. According to Table 2.1 in section 2.2, we set the exponential rent's constant p to 0.45 and 0.5 respectively, while α is set to 0.82 and 1.9 respectively. The total wirelength is analyzed with different gate count and tier count and the results are plotted in Figure 3.5. The left part of Figure 3.4 shows how the total wirelength in a microprocessor changes as gate count varies from 100 to 1000000 with different number of tiers, while the right part of Figure 3.4 enlarges the region where gate count is between 1000 and 5000. Figure 3.5 shows the results for a random logic. As can be seen from the figure, for both design types with the provided parameters, gate count equals 5000 is the threshold. Any module with less than 5000 gates does not worth partition because the wirelength improvement is very small. The cell area of a 2-input NAND gate is 48 μm^2 in our library, so we consider any module smaller than $5000 \times 48 \mu m^2 = 240000 \mu m^2$ does not deserve partition too. Finally, we conclude that 5000 cells or $240000 \mu m^2$ cell area, whichever comes first, is the partition threshold for a module.

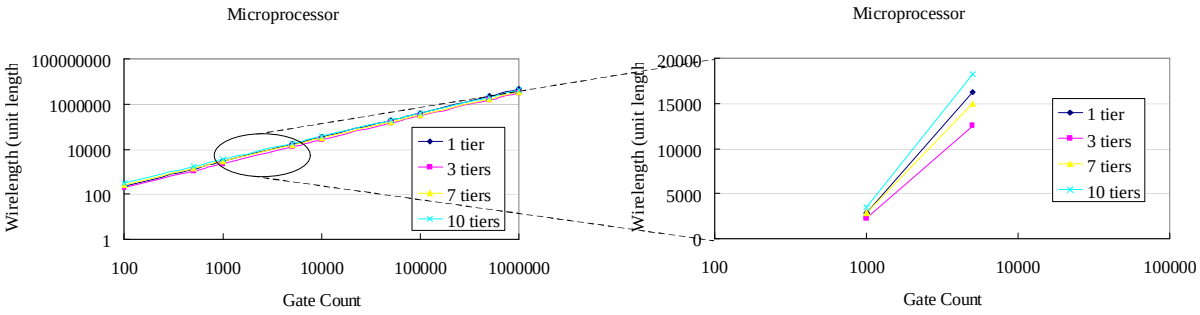

Figure 3.4: Total wirelength vs. gate count at different tier count for a microprocessor.

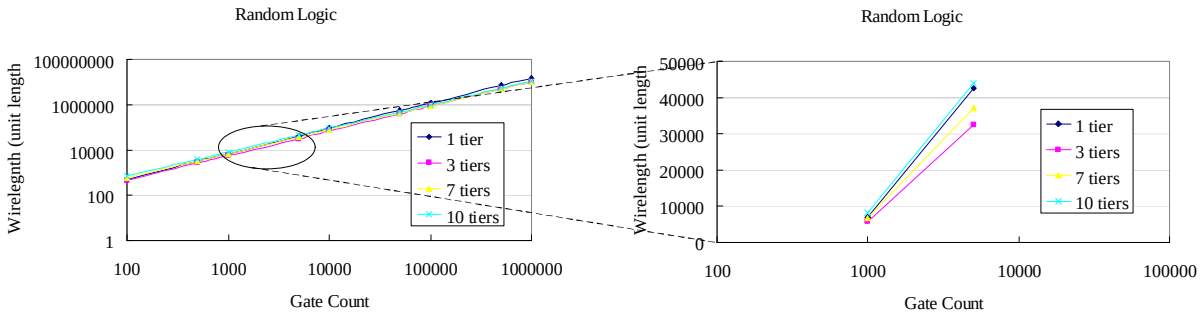

Figure 3.5: Total wirelength vs. gate count at different tier count for a random logic.

3.3 Floorplanning and Thermal Design

The floorplanning flow is shown in Figure 3.6. Floorplanning is a very important step in conventional IC design as well as in 3DIC design. In 3DIC design, floorplanning also needs to incorporate thermal design, and iterations between floorplanning and thermal design may occur to find an optimal floorplan. In this step, we place modules, thermal vias, signal vias, IO pads, and power rings/stripes, which can be done either automatically or manually. The starting point is the tier specific netlists generated by the partitioner. Since the cell area of each tier will not be identical, the bounding-box dimension for all tiers must be determined prior to floorplanning. The purpose of doing this is to facilitate 3D via alignment and thermal via insertion. If the designer does not choose the bonding-box dimensions, then an automatic method sets the same placement density for each tier and picks the largest area. However, this approach is not advised, because it usually does not give an optimal dimension, so the manual decision is required to help achieve the smallest chip area. When the dimensions are decided, we can start floorplanning.

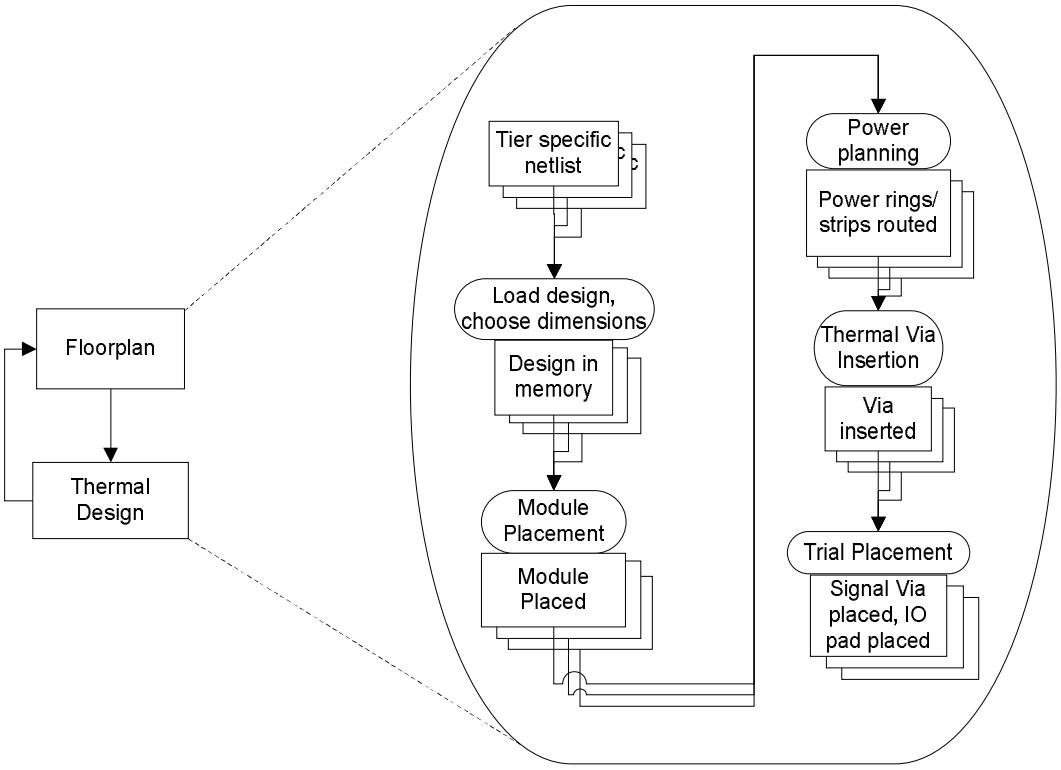

Figure 3.6: Floorplan and thermal design.

When the dimensions are set for all tiers, tier specific netlists are loaded individually. Then a *First Encounter* "module placement" is executed to place modules and set their aspect ratios such that wire lengths on each tier are minimized. Designers can manually change the location of some modules to optimize overall wirelength, timing, and maximum temperature, if desired. To minimize the temperature rise, we should avoid stacking hot blocks on top of each other. Then, power rings and stripes are added to the floorplan. The power rings are added to surround the die area and the vertical stripes are planned across the die. Special rings can be added to IP blocks with some manual manipulations. Next is the thermal via insertion. Via cells can be placed anywhere along with other stand cells without affecting the previous floorplan. Generally, thermal vias are evenly distributed across the chip, but in some cases,

thermal vias can be inserted according to the local power density. More detail of thermal via planning is discussed in section 4.2

Since IO pads are assumed to exist only on the top tier, several hundred signal vias are used to connect power/ground rails on different tiers, which provide better power distribution and hence reduce IR drop. In this flow, power/ground signal vias are placed along the location of vertical power/ground stripes to save area.

The last step in this flow is signal via and IO pad placement. We will first perform a trial placement on each tier to get the appropriate location for each inter-tier signal via and IO pad. IO pads are placed at core boundaries and a trial placement can determine the locations for most signal pads, though the location of the signal pads are adjustable in later steps. The pad pitch of the MIT technology we used is 100 μm. The determination of locations for inter-tier signal vias is more difficult. If we consider the three-tier MIT process as an example, we use equations (8) – (11) to determine the locations of those inter-tier vias. These equations basically set the new via location on each tier to be the centroid of the previous two locations. Inter-tier vias are then placed to these new locations.

For VIA_AB:

$$new\_location = \frac{TierA\_connections \times TierA\_coordinate + TierB\_connections \times TierB\_coordinate}{TierA\_connections + TierB\_connections} \quad (8)$$

For VIA_BC:

$$new\_location = \frac{TierB\_connections \times TierB\_coordinate + TierC\_connections \times TierC\_coordinate}{TierB\_connections + TierC\_connections} \quad (9)$$

For VIA_AC:

$$new\_location = \frac{TierA\_connections \times TierA\_coordinate + TierC\_connections \times TierB\_coordinate}{TierA\_connections + TierC\_connections} \quad (10)$$

For VIA_ABC:

$$new\_location = \frac{(TierA\_connections \times TierA\_coordinate + TierB\_connections \times TierB\_coordinate + TierC\_connections \times TierC\_coordinate)}{TierA\_connections + TierB\_connections + TierC\_connections} \quad (11)$$

3.4 Placement, CTS, and Route

Directly following the result of floorplanning, the Place & Route (P&R) sub-flow will be invoked. The entire P&R flow is show in Figure 3.7. In this sub-flow, we will do placement, Clock Tree Synthesis (CTS), routing and tier specific Signal Integrity (SI) checks. Besides the normal P&R steps, a thermal refinement (to modify the locations of thermal vias) step during the analysis later in the flow is defined as an option for designers if they identify any thermal violation during the analysis later in the flow.

Starting from the floorplaned designs, a placement and a 3D CTS are performed and saved. This procedure is iterative, because it is very hard to control the skew on different tiers concurrently. And more details about 3D CTS are discussed in section 3.4.2. When CTS is done, a trial route is performed and parasitics of the clock tree are extracted, merged, and back-annotated for a SPICE simulation. The RC values of clock trees on all tiers are extracted in the format of π models and then joined to form an entire clock tree network with resistors,

capacitors and transistors. A SPICE simulation is invoked when the parasitic merge is done. Then we calculate clock skew in the form of (12). If skew meets our specification, we go to next step, otherwise we will load the placed design and change CTS specs to do another synthesis until the skew meets the specification.

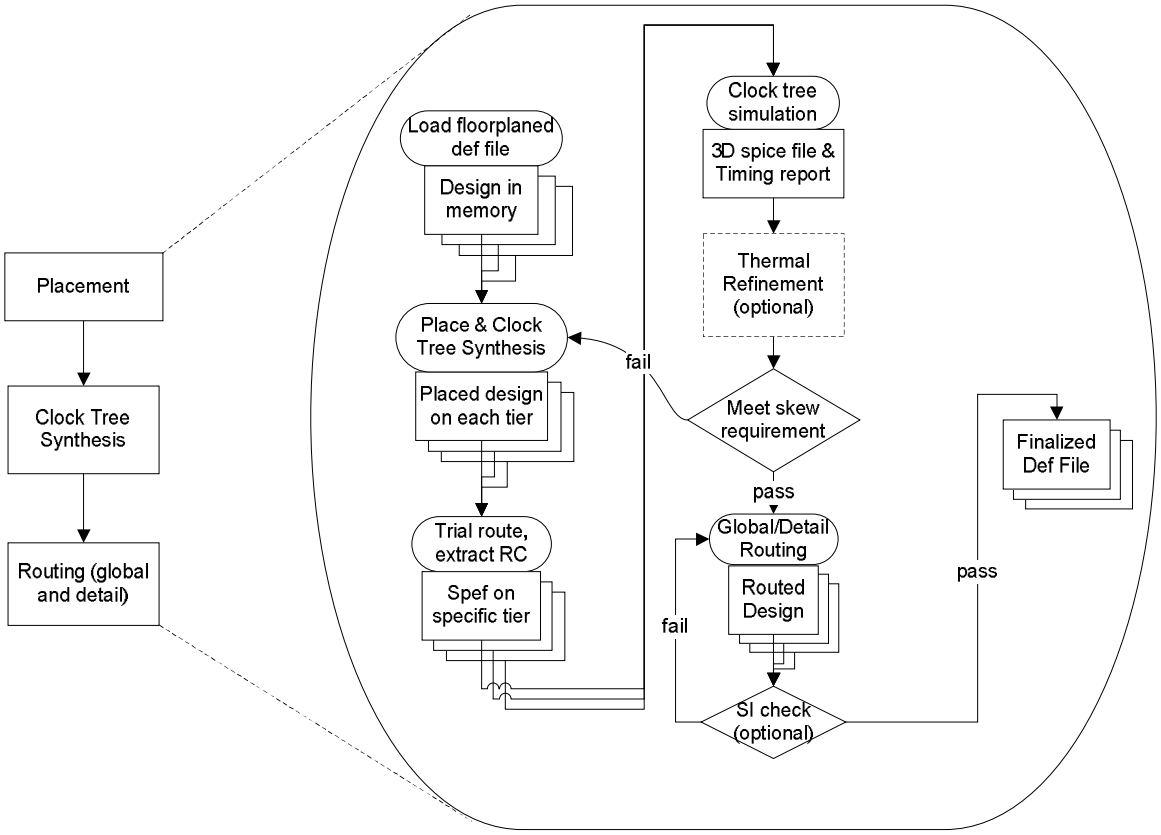

Figure 3.7: Place, clock tree synthesis, and routing.

$$skew = Maximum\_insertion\_delay - Minimum\_insertion\_delay \quad (12)$$

When the CTS is done, an optional thermal refinement process is presented to the designer. This is due to the requirement to eliminate hot spots in the system to avoid timing and/or power violation. Since the leakage power is exponentially dependent on the operating temperature, devices in hot spots are more likely to break down due to a positive feedback between leakage

and temperature. In this case, thermal via cells are moved from the regions with low temperatures to the hot spots to make the temperature profile more uniform. Since the thermal conductivity of thermal via (Tungsten) is two orders of magnitude larger than the Silicon-Dioxide, a slight increase of thermal via significantly improves the equivalent thermal conductivity in the hot region and hence reduces the temperature a lot. Besides, delay is also temperature dependent and we may have clock skew violation if the clock buffers are operating at significant different temperatures. If such violations are identified during analysis, thermal vias have to be moved closer to clock buffers operating at higher temperatures to bring heat away. On the other hand, stacking of clock buffers should be avoided if possible. However, this process is optional because a designer may find the temperatures at hot spots are below threshold so that no further action is required.

Routing is another important issue in this sub-flow. 3D routing is affected by the congestion brought by 3D signal and thermal vias. Usually a trial route is performed to analyze the congestion before the detail route. If the trial route reports very high congestion number, a placement refinement has to be enforced to make the design routable. If the congestion analysis says OK, a crosstalk avoidance routing is carried out to complete detail routing. Tier specific signal integrity (SI) is checked after the detail routing to ensure no large noise exists, especially on clock and reset nets. If the SI check passes, then final tier specific Design Exchange Format (DEF) and SPEF files are exported for subsequent steps. In this 3DIC design flow, *Celtic* [54] is used to perform the SI check.

3.4.1 Parasitics of 3D Inter-tier Via

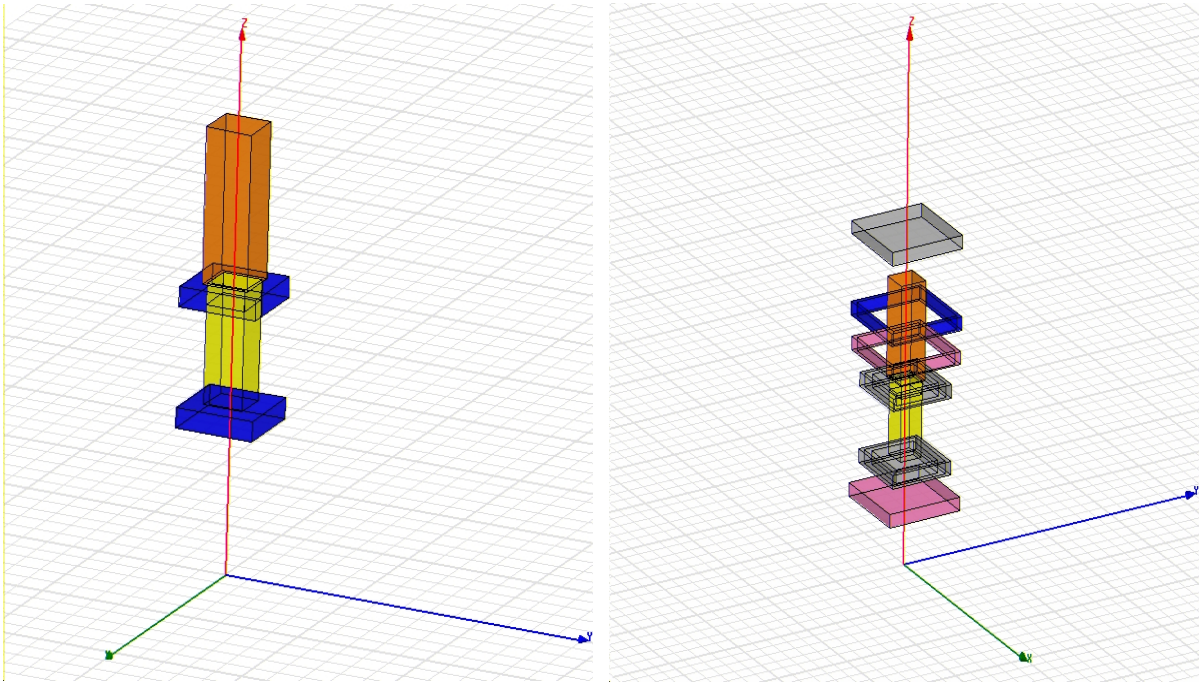

(a). VIA_AB isolated (b). VIA_AB Shield

Figure 3.8: Inter-tier via modeled in Q3D.[2]

The parasitics of a 3D via is a very important parameter for both clock tree verification and final performance verification. Because this data was not provided by the foundry, an electromagnetic field solver *Q3D* [50] was used to simulate the parasitic data for these vias by solving capacitance matrices. Shown in Figure 3.8 [49], two configurations of a 3D via connecting tier A (bottom tier) and tier B (one tier up) are simulated to find the parasitic capacitance bounds. Figure 3.8 (a) shows the isolated case that there is no metal around the 3D via, which gives the lower bound of the parasitic capacitance. Figure 3.8 (b) shows the shielded

[2] Thanks to Chris Mineo for creating this picture.

case, in which the 3D via was surrounded by all possible metal layers with minimum distance, which gives the upper bound.

Though parasitic resistance seems identical for different 3D vias and shielding, *Q3D* simulations of MIT Lincoln Lab technology show a slight difference on parasitic capacitance of via AB and via BC. The parasitic data is listed in Table 3.1. The parasitic capacitance of shielded 3D via is roughly equivalent to a 20 μm minimum width minimum distance shielded metal2 wire. The resistance is very small when compared with either a 20 μm metal2 wire or a normal via (4Ω), which makes it negligible during later parasitic extraction and merging. In the MIT 3D technology, via AC (and ABC) constitutes of one via AB and on via BC.

Table 3.1: 3D via parasitic.[3]

Capacitance					Resistance
		Isolated		Shielded	
	Tier(s)	Handle-silicon only	With back-metal		
Inter-Tier Via	AB	0.82 fF	0.99 fF	4.34 fF	0.115 Ω
	BC	0.89 fF	1.01 fF	4.15 fF	0.115 Ω
Metal 2 (20 μm)	A	1.01 fF	3.69 fF	4.43 fF	6.4 Ω
	B	0.95 fF	3.27 fF	4.42 fF	6.4 Ω
	C	0.89 fF	3.66 fF	4.41 fF	6.4 Ω

3.4.2 3D Clock Tree Insertion

We have investigated two clock tree insertion schemes within this flow. They are illustrated in Figure 3.9. The scheme shown in Figure 3.9 (a) is defined as "Root-Pin CTS" because the clock tree on each tier is inserted from boundary pins that are aligned to the same

[3] Thanks to Chris Mineo for creating this table.

X-Y position on each tier. 3D vias are later inserted just outside the boundary to connect these nets. This synthesis scheme is a straightforward extension of a conventional 2DIC CTS. However, if we use this scheme in 3DIC design, we find the clock skew is very hard to control, because the loads on each branch of the clock tree are difficult to balance. According to the Elmore model, the clock-to-first-level-buffer delay and slew rate of the nodes $N_1 \sim N_n$ in Figure 3.9 varies significantly tier-to-tier due to the complex RC topology, especially when tier number is large. So another scheme defined as "Root-Cell CTS", which is shown in Figure 3.9 (b), is developed to reduce the RC complexity at clock pad output. Rather than use the boundary pin as the clock source, Root-Cell CTS creates a clock source buffer on each tier and aligns them to the same X-Y position on each tier, somewhere close to the off-chip clock pad. Then tier specific clock source buffer acts as the source of the clock tree. The parasitics on the first level of the clock tree are more predictable with this method, so the delays and slew rate from clock pad to node $N_1 \sim N_n$ are more controllable in this scheme. To insert the clock tree under the "Root-Cell CTS" scheme, one set of clock tree constraints is created for each tier. Clock trees are inserted and verified respectively and will be joined for final verification if clock trees on each tier have similar insertion delay and transition time.

As shown in Figure 3.10, the parasitics of both CTS schemes are demonstrated. In the Root-pin CTS scheme, "RC parasitics tier n" represents the RC network between a 3D via and the first level clock buffers inputs. These parasitics are subject to change during different iterations of the design flow and are very difficult to control. Therefore, the delay at node $N_1 \sim N_n$ and the clock skew at each sink are also very difficult to control. Inserting a clock source buffer on each tier totally changes the picture. The RC network connected to clock pad remains

constant as long as the number of tiers is determined and hence the delay at node $N_1 \sim N_n$ remains unchanged. Then the 3D CTS problem is reduced to a 2D CTS problem that is well defined in any commercial tool.

A brief performance comparison between these two clock tree insertion schemes is listed in Table 3.2. The two benchmark circuits are implemented in the original MIT Lincoln Lab process (three tiers, three metal layers) in both 2D and 3D integration using manual tweaking to help reduce the clock skew. As can be seen from the table, the clock skew is significantly reduced with Root-Cell CTS by paying a tiny price of power. In this experiment, the same skew target 400*ps* was set to each CTS scheme. The final skew values were obtained from the CTS verification flow after detail route.

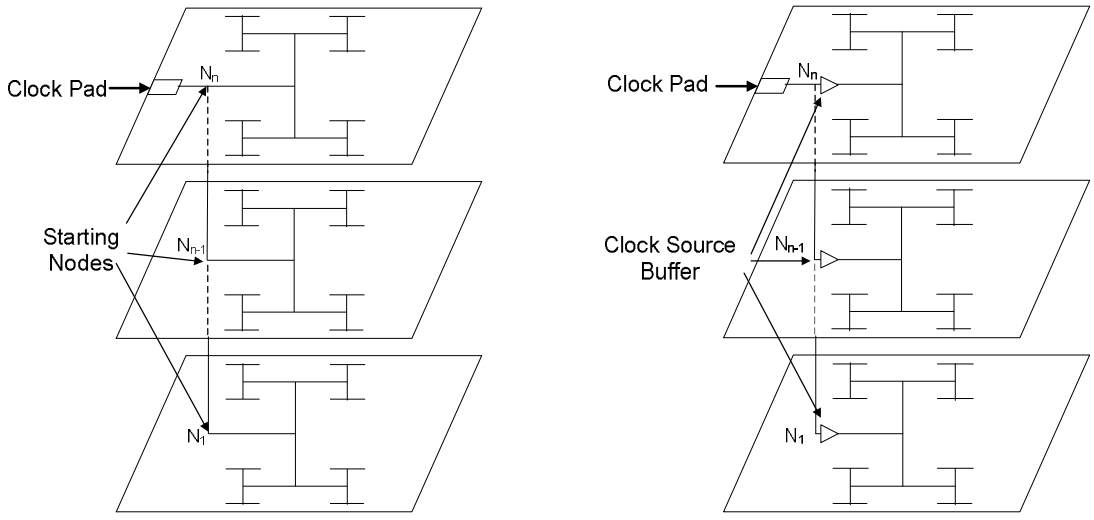

(a) (b)

Figure 3.9: Clock tree insertion schemes. (a) Root-Pin CTS, clock tree is inserted on all tiers simultaneously; and (b) Root-Cell CTS, clock trees are inserted on different tier respectively and joined together.

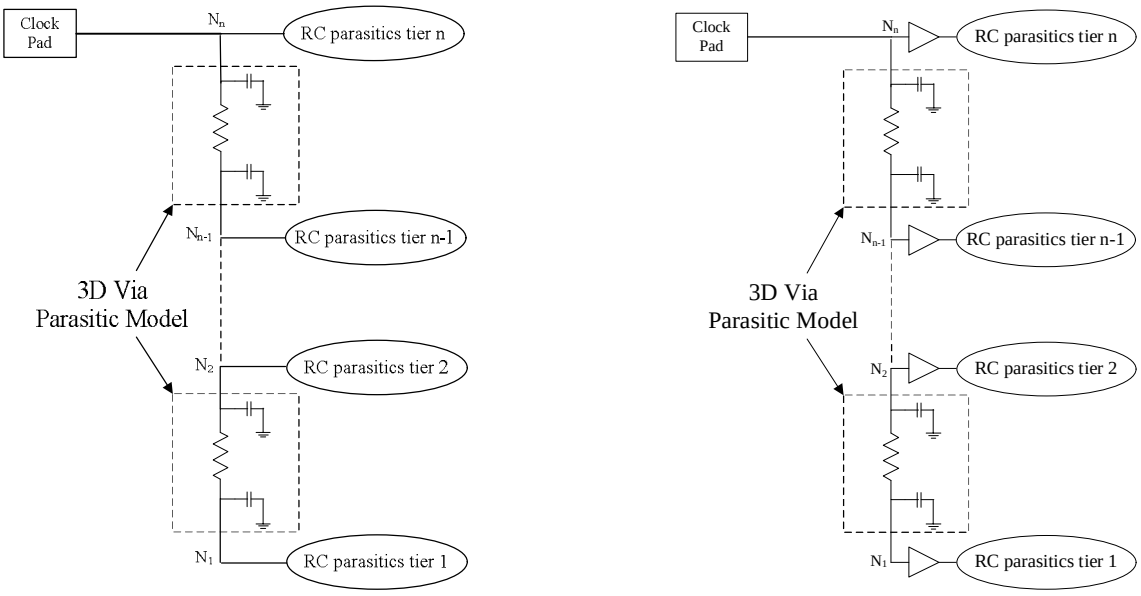

Figure 3.10: Clock tree parasitics. (a) Root-Pin CTS; and (b) Root-Cell CTS.

Table 3.2: Performance Comparison between Root-Pin CTS and Root-Cell CTS

	2D CTS		Root-Pin CTS		Root-Cell CTS	
	Clock skew	Power	Clock skew	Power	Clock skew	Power
FFT	489 *ps*	233 *mW*	446 *ps*	195 *mW*	196 *ps*	202 *mW*
ORPSOC	532 *ps*	794 *mW*	457 *ps*	467 *mW*	180 *ps*	483 *mW*

3.5 Timing and Power Simulation at Nominal Temeperature

After the final tier specific DEF and SPEF files are ready, the verification flow is executed to make sure that timing and power under nominal temperature will match the specification. The verification flow is demonstrated in Figure 3.11. Verifications will be performed on timing, power, and final routed clock tree. When the verification step passed, a timing report and a

power/material density report file will be generated for electro-thermal delay-power coupling analysis[4].

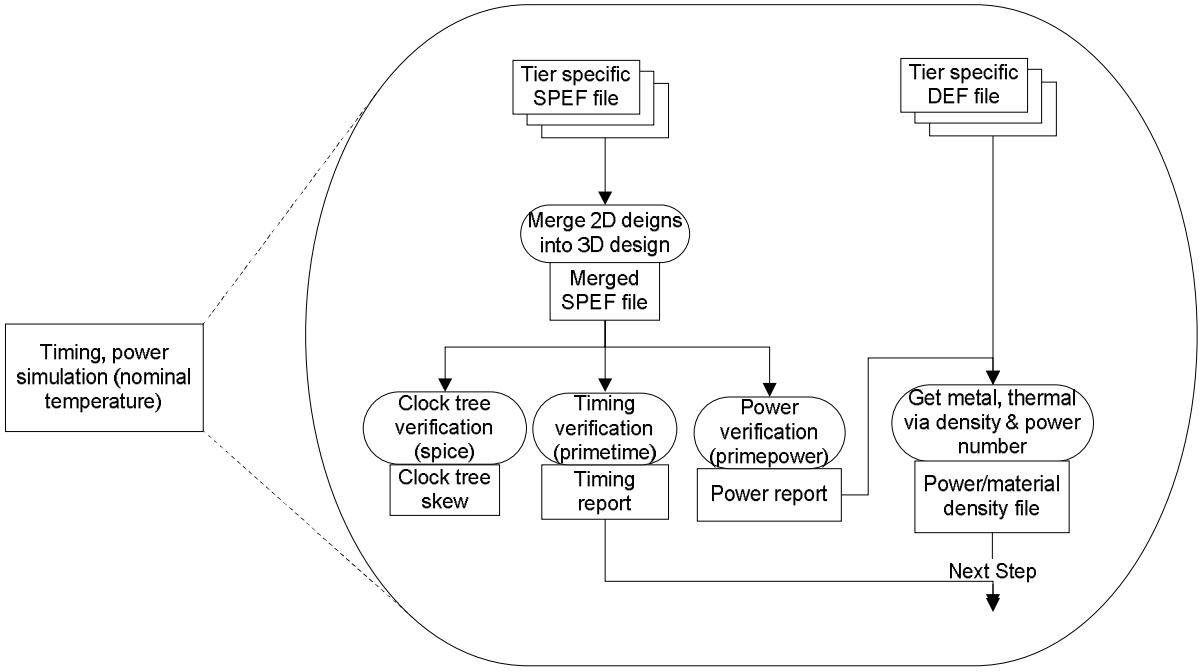

Figure 3.11: Timing and power simulation flow with nominal temperature.

First, the tier specific SPEF files are merged to provide parasitic data for the 3D design[5]. The detailed SPEF merge methodology is described in [49]. The parasitics of the detail routed 3D clock tree will be extracted for a SPICE simulation to update final insertion delay and transition time. The clock skew is calculated with (12) shown in section 3.4. *PrimeTime* reads the top-level verilog netlist, final clock skew, and merged SPEF file, then reports critical path delay and checks for hold violation. The hold time check is most critical for the success of any digital circuit because hold time violation will cause the chip to fail at any clock frequency.

[4] Thanks to Samson Melamed for creating the density report file format.

[5] Thanks to Chris Mineo for creating the SPEF merge script.

PrimePower [64][6] takes the inputs of *PrimeTime*[6] along with a Value Change Dump (VCD) file[7]. It analyzes the circuit and generates a power report for each cell. This power report is analyzed along with the tier specific DEF files to generate a file containing power/material density information[4].

The power/material density report contains following information in sequence: tier-count, row number, column number, metal density on north edge, metal density on south edge, metal density on east edge, metal density on west edge, thermal via density in current region (a region is defined as a rectangular or square block with certain area on the chip) connected to upper tier, thermal via density in current region connected to lower tier, average dynamic power, and average leakage power. Currently, only metal, inter-tier vias (both signal and thermal), and silicon dioxide are considered for heat conduction.

A later power estimation flow adopts 10 tiers uses a different approach. To being, a forward annotation Switching Activity Interchange Format (SAIF) file is created in *Design Compiler*. To write a forward annotation SAIF file in design compiler, the RTL description of the design must first be loaded and linked. Then the "rtl2saif" command is used to interpret the design and finally write a forward annotation SAIF file [47]. When the SAIF file and parasitics are obtained for the design, *Design Compiler* is used to estimate the power consumption.

3.6 Electro-Thermal Power-Delay Analysis

The actual performance of a 3DIC may be significantly different from the value obtained from nominal temperature simulation. Because the temperature, delay, and power are co-

[6] In the extended technology, *Design Compiler* is used instead of PrimePower and PrimeTime

[7] Thanks to Ambarish Sule for creating the VCD file generation flow.

dependent variables, a coupling phenomenon between temperature and performance may occur. This coupling is caused by the device delay and power temperature dependency. In the ideal case, the entire chip has a uniform temperature profile, and this temperature falls below the conventional worst case temperature, which is 125 °C. However, the temperature profile in 3DIC is not uniform and it may have large temperature gradient. The worst case temperature can not be guaranteed to be an upper limit. Studies have shown that SOI devices still function properly up to temperatures of 250 °C [36]. At higher temperatures, the circuit becomes slower but much leakier. If we assume the circuit is running at its maximum speed, the dynamic power goes down at higher temperature and hence the chip becomes cooler. However, leakage power increases exponentially with temperature that makes the chip hotter. The final performance and temperature profile will be quite different from the initial value due to this feedback mechanism. In this section, we will discuss the electro-thermal coupling in terms of performance temperature dependence and this feedback mechanism.

3.6.1 Performance-Temperature Dependence

3DIC performance degrades and reliability is compromised when operating at high temperatures. Though dynamic power is independent of temperature [35] (short circuit power is temperature dependent, but it only consumes a small part of overall power [47]), logic gate delay and leakage power will increase with temperature. If we assume the system is running at its highest speed, longer delays have a negative effect on temperature due to lower frequencies and smaller dynamic power. However, leakage power has positive effect on temperature, because higher temperatures increase leakage power considerably. In evaluating the

performance of a 3DIC, we must consider this relationship at the point at which temperature and performance converge. In practical timing or power analyses, it is impossible to include tens of libraries at different operating temperatures. Power and timing calculation becomes very difficult and computationally intensive when accounting for temperature variation. Many timing arcs span multiple silicon tiers, therefore single timing paths operate with their transistors at significantly different temperatures. Short of the feasible, naïve brute-force approaching making use of tens of separate libraries, it is very difficult for current tools (such as *PrimeTime* and *PrimePower*) to perform a fast yet accurate analysis of system performance. In this section, we developed two models to address the delay-temperature and the leakage-temperature dependencies. The interconnect resistance-temperature dependency is ignored in this work because it is relatively small compared to the transistor output (or holding) resistance.

3.6.1.1 Transistor-Delay Temperature Dependence

The transistor delay-temperature dependence can be expressed as (13) [16] [27] [28]. In this case I_D is the drain current, T is the absolute temperature in Kelvin, α is the velocity saturation index, and μ is the mobility. The typical value of α is 1.5, but it is actually smaller for a short channel MOSFET [28]. The typical β value is 1.5, too. The temperature dependence of threshold voltage V_{TH} and mobility μ can be expressed as (14) and (15) [16], where T_0 is the reference temperature (usually room temperature). As shown in [36], the threshold voltage temperature coefficient, k, is also a weak function of temperature. For fully depleted SOI, the variation in k due to temperature is relatively small [36]; for simplicity, we may treat it as a linear function of temperature shown in (16).

$$delay \propto \frac{CV_{DD}}{I_D} \propto \frac{CV_{DD}}{m(T)(V_{DD} - V_{TH}(T))^a} \tag{13}$$

$$V_{TH}(T) = V_{TH}(T_0) - k(T - T_0) \tag{14}$$

$$m(T) = m(T_0)(\frac{T}{T_0})^{-b} \tag{15}$$

$$k = k_0 + g(T - T_0) \tag{16}$$

Substituting (14) and (15) into (13), we derive:

$$delay(T) \propto \frac{CV_{DD}}{m(T_0)(T / T_0)^{-b}(V_{DD} - V_{TH}(T_0) + k(T - T_0))^a} \tag{17}$$

$$delay(T) = \frac{delay(T_0)(V_{DD} - V_{TH}(T_0))^a T^b}{T_0{}^b (V_{DD} - V_{TH}(T_0) + k(T - T_0))^a} \tag{18}$$

We have varied k_0, α and γ values for a 1X buffer driving 50*fF* and 150*fF* loads, between 0°C and 250°C, to obtain results in line with SPICE simulations. It is necessary to use slightly different k_0 values for falling and rising delays. Table 3.3 shows simulated delay values from SPICE, versus those obtained using equation (18), for a 1X buffer with different load. The table gives the maximum error and the corresponding temperature when the maximum error occurs. Figure 3.12 gives a clearer picture of how the rise delay of a 1X buffer depends on temperature. As can be seen in this figure, the delay model (18) is fairly close to SPICE simulation for this simple case.

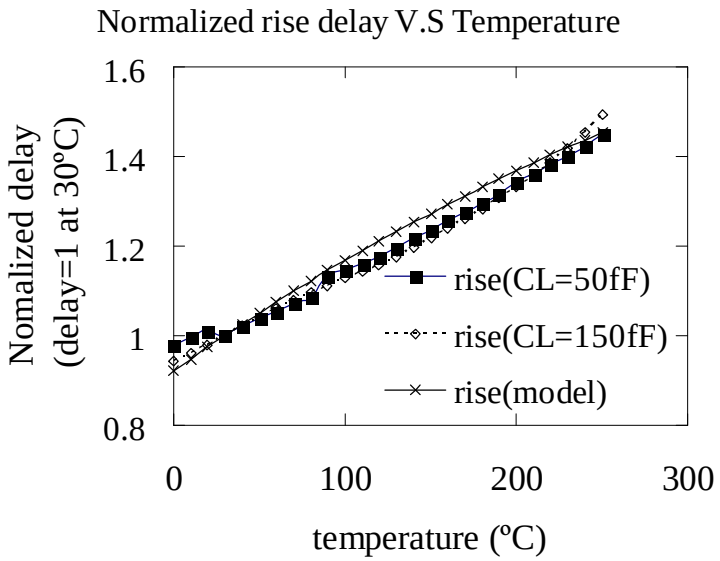

Figure 3.12: SPICE simulation vs delay model on 1X buffer rise delay.

The delay model (18) is the basis for temperature-timing coupling analysis. For the standard cell library used in this work, across the temperature range shown and the range of loads common in the benchmark circuits, the largest disparity between SPICE and the delay model is an XOR gate with a delay error of 11.6%. Further study shows that all delay error falls in the range of [-3.6%, 11.6%].

Table 3.3: Comparison between SPICE simulation and delay model for 1X buffer with temperature as variable (MIT SOI technology).

	k_0 (mV/K)	γ (mV/K)	β	α	Max err	Max err T
Rise	1.2	0.003	1.5	1.3	6.1%	0°C
Fall	1	0.003	1.5	1.4	5.2%	170°C

3.6.1.2 Leakage Power Temperature Dependence

Leakage power drastically varies with operating temperature. Though leakage power in SOI devices is considerably smaller than that of bulk devices, it can still comprise a large portion of the total power when the temperature is very high. To create an accurate leakage model with respect to temperature, extensive SPICE simulations have been run using the BSIMSOI [51] model (Hspice level 57) to determine the temperature-leakage relationship for the standard cells in the library used in this work. The leakage temperature dependency is super-linear, and can be modeled as a polynomial function [37]. For the temperature range between 0°C and 250°C, a third-order polynomial is found to be accurate enough to describe the leakage temperature dependency, with a maximum error of 9.9%. The leakage model is of the form:

$$\frac{I_{leakage}(T)}{I_{leakage}(T_0)} = 1 + a_1 \cdot (T - T_0) + a_2 \cdot (T - T_0)^2 + a_3 \cdot (T - T_0)^3 \qquad (19)$$

Table 3.4: Coefficients of leakage temperature dependency model (MIT SOI technology).

Coefficient	*α1*	*α2*	*α3*
Value	0.0226	0.00033	1.77E-6

Using a curve fitting technique, the values for coefficients *α1*, *α2*, and *α3* in (19) are determined. Because the coefficients are slightly different for each standard cell, the average values are used. These values are listed in Table 3.4.

The model was verified by re-characterizing the stand-cell library for a range of temperatures and using *PowerCompiler* to predict the power on two small benchmark circuits: a

4-bit ripple adder with 19 gates and a 32-bit ripple adder with 1,381 gates. Figure 3.13 shows the leakage model compared to the simulation. It is noticeable that the model (19) becomes more accurate for large circuits because the total leakage approaches the average among all cells. Further simulation shows that the error of individual cell in the entire standard cell library falls in range of [-4.5%, 9.9%]

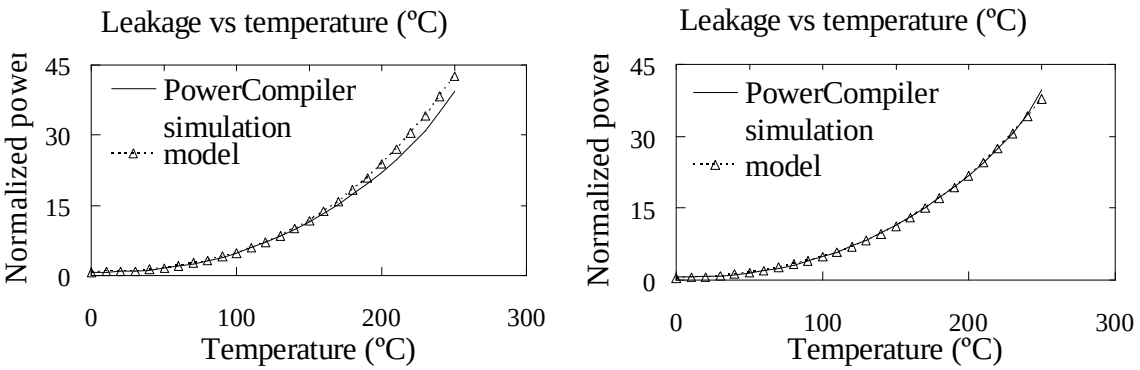

(a) 4-bit adder (b) 32-bit adder

Figure 3.13: Leakage model with power compiler simulation. Normalized Power=1 when T=30ºC

3.6.2 Iterative Calculation of Timing, Power, and Temperature

Power density in 3D integration increases drastically as the number of silicon tiers increases. The maximum temperature can soon exceed the widely accepted nominal worst case temperature. Circuits are often designed for the worst case temperature of 125ºC, but this limit is not guaranteed for 3DICs. Due to the coupling effect, the final delay, power, and temperature of the circuit are determined using an iterative approach illustrated in Figure 3.14. For simplicity, we also assume dynamic power, comprised of the short circuit power and switching

power, is linearly proportional to clock frequency, as shown in (20). The iteration is easily coded into most scripting environments, and the experiments show that timing, power, and temperature tend to converge within 3 or 4 iterations. The Pseudo code of the electro-thermal coupling analysis is shown in Appendix C.

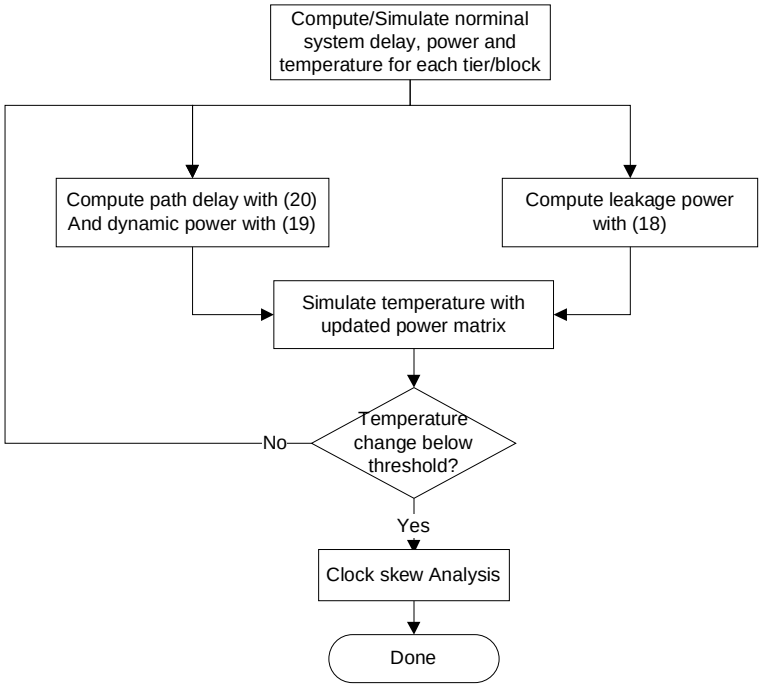

Figure 3.14: Iterative solution of timing, power and temperature.

$$P_{dynamic}(f) = P_{dynamic}(f_0) \cdot f(T) / f(T_0) \tag{20}$$

$$delay_{path}(T) = \sum_{start-cell}^{end-cell} delay_{cell}(T) \tag{21}$$

$$f(T) = \frac{1}{delay_{path}(T)} \tag{22}$$

When the coupling analysis is done, a final clock tree evaluation with the obtained temperature profile will be executed to make sure clock skew still meets the requirement. The Clock skew verification is discussed more in Chapter 4.

3.7 DRC and LVS Check

DRC and LVS checks are not shown in the figures because they have no interaction with any part of the 3D design flow and are only executed when the 3D chip is ready for fabrication. A manual fix is required when an error or violation occurs. Sometimes multiple power and ground names are used in the netlist but are connected together in real layout, which will cause LVS errors. Before DRC and LVS check can be executed, the tier specific layouts have to be loaded into the NCSU 3DIC design environment (refer to Appendix A) and merged into one layout database. Finally, the NCSU 3DIC design environment only supports up to three tiers while our design flow supports up to ten tiers. These exceptions make DRC and LVS hard to automate, and so they form a separate step. When DRC and LVS checks pass, a GDSII file is exported for fabrication.

Chapter 4.

Thermal Management and Reliability Analysis for 3DICs

4.1 Introduction

According to thermal dynamics, the thermal system can be written as (23), where ρ is material density, c is the specific heat, q is the generated power and $k\nabla^2 T$ is the heat conduction due to temperature gradient. In steady state, equation (23) can be simplified into (24) [53] [58], which shows that the temperature of a point (x,y,z) is directly determined by the heat generation and the conduction properties of this point. In an electrical system, the heat generation is equivalent to power consumption, which is modeled as a current source, and the conduction properties can be modeled as six thermal resistors connecting its neighbors [53] [57].

$$r\,c\,\frac{\partial}{\partial t}T = q + k\nabla^2 T \tag{23}$$

$$K\left[\frac{\partial^2 T}{\partial x^2}+\frac{\partial^2 T}{\partial y^2}+\frac{\partial^2 T}{\partial z^2}\right]+q(x,y,z)=0 \tag{24}$$

Thermal effects have been a consideration of electrical system design for many years. A common approach is to transform the material properties to thermal resistance to estimate the system temperature [52]. A simple one-dimensional model of the basic relationship between temperature T, thermal resistance R, and power P is shown in (25), where T_{chip} is the chip

temperature and $T_{ambient}$ is the environmental temperature (usually considered to be 27°C). A more general relationship is given in (26) [24]

$$T_{chip} = T_{ambient} + R \cdot P \tag{25}$$

$$\Delta T_i = R_i \sum_{j=i}^{n} P_j \tag{26}$$

It is useful to simplify the system to one-dimensional, because heat generally flows in the same direction in an integrated circuit towards the heat sink. To understand the thermal model used in this study, it is necessary to assume a package model and certain details of the processing technology. The model used in this study is briefly illustrated in Figure 4.1 [60] (not drawn to scale). The chip is adhered to the heat sink using a thermal paste. The inter-tier signal vias (ISV) act as both electrical and thermal paths and are included during thermal calculation.

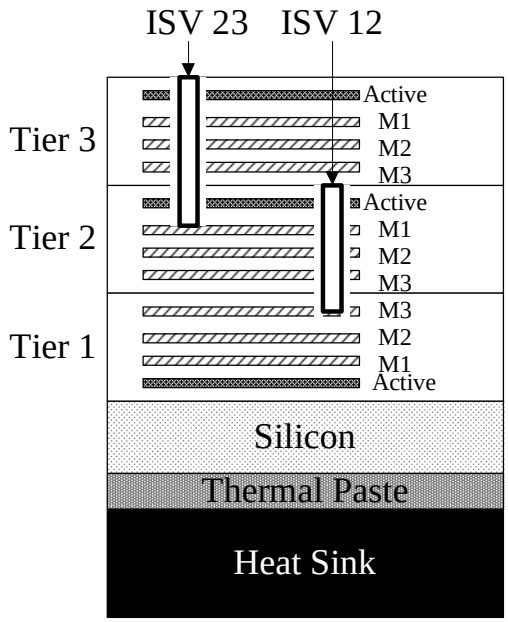

Figure 4.1: Package of the 3D processing technology.

A steady state thermal system can be considered equivalent to an electrical system using the analogies listed in Table 4.1 [30]. To perform the thermal analysis of our system, each tier is evenly divided into m by m blocks, where m can be any integer greater than 1. However, because the fact that run time and memory usage increase exponentially with m, an upper limit of 50 is set for m in this dissertation. Otherwise, a recursive thermal simulation (see Appendix C) is performed to evaluate the temperature profile of the 3D system.

Table 4.1: Thermal-electrical system component analogy.

Thermal <=> Electrical	Thermal <=> Electrical
Temperature T(K) <=> Voltage V[v]	Heat Q[J] <=> Charge Q[C]
Heat transfer q[W] <=> current i[A]	Thermal Res R_T[K/W] <=> Electrical Res R[V/A]
Thermal cap C_T[J/K] <=> Electrical cap C[C/V]	Temperature rise $\Delta T = q \cdot R_T$ <=>Voltage difference $\Delta V = i \cdot R$

Assuming adiabatic boundary conditions on the four sides and the top of the chip, an equivalent electrical network shown in Figure 4.2 (a) is created to represent the thermal network. In steady-state case, power and temperature for any block are treated as constants. Therefore, thermal capacitance can be ignored, and the power dissipation can be modeled as a constant current source. With this assumption, each block is modeled as a constant current source with six resistors connecting to its neighbors. The heat sink is modeled as electrical ground and the ambient temperature is modeled as a constant DC voltage source. The resistors between the bottom tier and the voltage source (the heat sink) represent the overall thermal effects of the heat sink, silicon substrate, and the thermal paste. The lateral resistors model the heat spreading of each block. This model is used to verify a design's temperature profile once the detail routing has been completed. Though such simulation can obtain a relatively accurate

temperature profile [13], it is very timing consuming, especially when electro-thermal coupling has to be considered. Instead, a reduced thermal model shown in Figure 4.2 (b) is used during floorplanning. This model is a simplification of the detailed thermal model in Figure 4.2 (a), in which *m* is set to 1. In this model, an equivalent thermal resistance is incorporated during thermal and performance analysis.

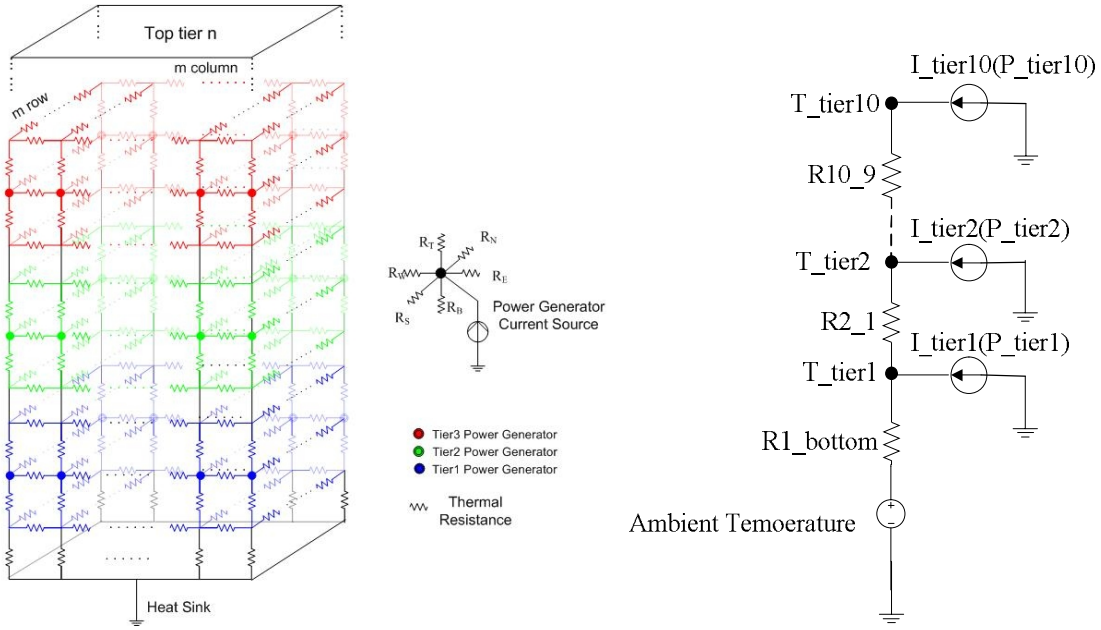

(a) Detail thermal model (b) Simplified thermal model

Figure 4.2: Thermal model.

To better explore the thermal and performance trade-off in 3DICs, a larger temperature gradient in z direction is mandatory [47]. The original 3 tier integration technology is not sufficient because of two reasons: (1). The power density is not high enough to cause a significant electro-thermal coupling phenomenon; and (2) The largest temperature rise is caused by the thermal paste rather than the wafer stacking itself. This methodology extends the tier limit to ten (set *n* in Figure 4.2 (a) and (b) to 10) to overcome the above limitation.

The equivalent thermal resistance of a bulk material, which can be written as (27), is determined by the material's dimensions and thermal conductivity *k*. In this equation, *R* is the thermal resistance of this bulk material, *t* is the thickness, and *A* is the cross-section area. If *n* materials conduct heat in parallel or in series, the equivalent thermal conductivity can be expressed as (28) and (29) respectively, where A_i is the area of bulk material *i*. The equivalent lateral (x, y direction) thermal resistances usually represent the equivalence between two adjacent blocks and the vertical (z direction) thermal resistances are the equivalence between two adjacent active layers. All the equivalent thermal resistances are computed based on (27) with the lateral and vertical equivalent thermal conductivities, which are re-written as (30) and (31). The thermal conductivity's temperature dependency is ignored in this study since the variation is relatively small in the 0~250°C temperature range that we are examining [59].

$$R = \frac{t}{k \cdot A} \tag{27}$$

$$k_{parallel} = \sum_{i=1}^{n} k_i A_i \Big/ A \tag{28}$$

$$k_{serial} = \sum_{i=1}^{n} t_i \Big/ \sum_{i=1}^{n} \frac{t_i}{k_i} \tag{29}$$

$$R_{vertical} = \frac{t}{k_{equ-vertical} A_z} \tag{30}$$

$$R_{lateral} = \frac{l}{k_{equ-lateral} A_l} \tag{31}$$

If many blocks of material are conducting heat in both parallel and series as shown in Figure 4.3, then the equivalent thermal conductivity is much more difficult to calculate. However, the lower and upper bound of thermal resistance for Figure 4.3 (a) can be easily determined by setting the lateral thermal resistances (RL1, RL2, etc.) to 0 or infinite, as shown in Figure 4.3(b) and Figure 4.3(c). The equivalent thermal resistances of case (b) and (c) are given in (32) and (33).

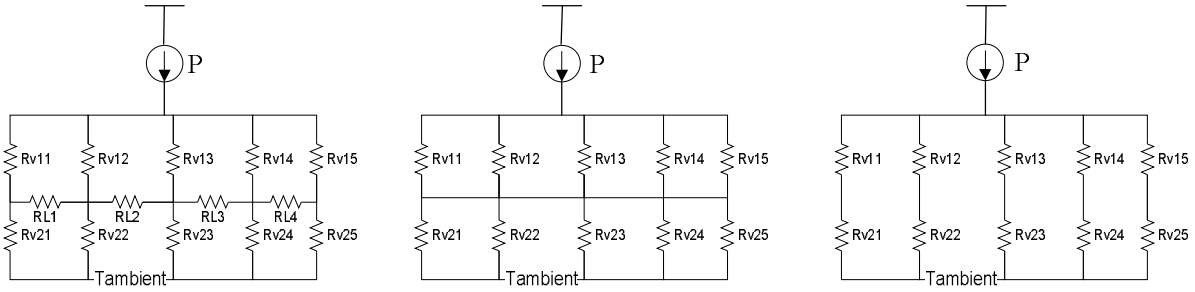

(a) General case (b) Ideal spreading (c) No spreading

Figure 4.3: Upper and lower bound of equivalent thermal resistance. (a) general case; (b) lower bound, ideal spreading; and (c) upper bound, no spreading

$$R_{\min} = (Rv_{11} // Rv_{12} // Rv_{13} // Rv_{14} // Rv_{15}) + (Rv_{21} // Rv_{22} // Rv_{23} // Rv_{24} // Rv_{25}) \quad (32)$$

$$R_{\max} = (Rv_{11} + Rv_{21}) // (Rv_{12} + Rv_{22}) // (Rv_{13} + Rv_{23}) // (Rv_{14} + Rv_{24}) // (Rv_{15} + Rv_{25}) \quad (33)$$

The equivalent thermal resistances of each tier are calculated, using thermal via density as a parameter, for the original three-tier three-metal-layer MITLL technology. Since the material area is unknown, the thermal resistances are calculated with a unit area equal to $1m^2$ (one square meter). As discussed in 5.4.1, the design becomes un-routable when thermal via cell density exceeds 20%, which is equivalent to a 5% thermal via density. So 5% thermal via density is set

as the limit for this calculation. Figure 4.4 shows the thermal resistance bound with no lateral metal considered.

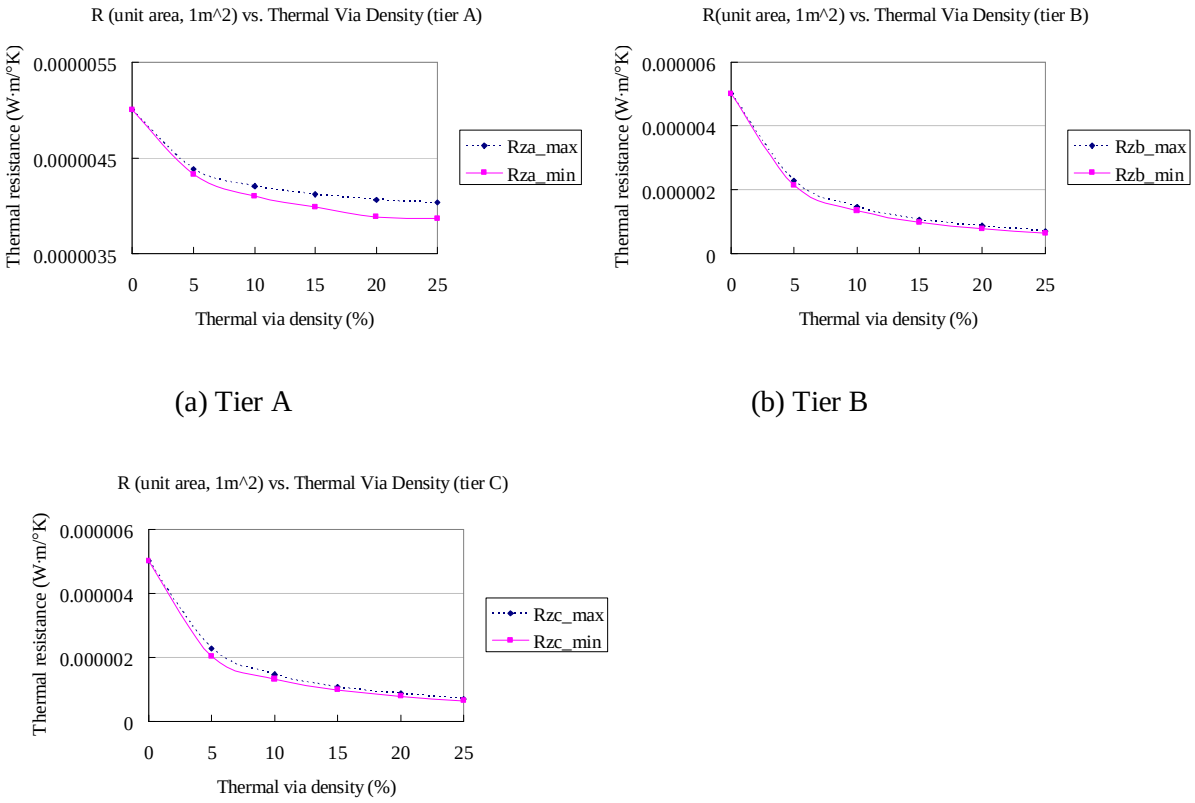

(a) Tier A

(b) Tier B

(c) Tier C

Figure 4.4: Thermal resistance upper and lower bound with different thermal via density.

For the thermal analysis flow used during floorplanning, the maximum thermal resistance between any two particular tiers is calculated using (33), assuming that only silicon dioxide and thermal vias (Tungsten) are present. The minimum thermal resistance is calculated with (32), assuming that the metal and via densities are at the maximum value allowed by design rules. However, after the detail routed design is complete, thermal resistance is calculated using metal densities.

4.2 Thermal Via Pattern Generation

Thermal vias are utilized to facilitate heat removal and has been demonstrated to be efficient [5] [13]. Traditional thermal via pattern formation aims to minimize the temperature. However, this target is vague because it is hard to tell what the right "low" temperature is. If excessive thermal vias are used, both chip area and wire-length will significantly increase, which can actually lead to higher power and temperatures. If no thermal vias are inserted into the system, the hotspots may be so hot that the devices in these hotspots eventually break down. Even if the devices still function correctly, the hotspots are still threats to both timing and reliability of the system.

In this work, via pattern generation is divided into global and detail phases, with different criteria.

1. **Global phase**: Determine total thermal via density and tier-count for optimal performance.

Here, the global phase design aims to optimize our performance if the maximum die temperature is below certain value (for example, we can set 150°C as the limit). In this case, the criterion for optimality is either energy per cycle or critical path delay depending on if the designer is aiming for a low power circuit as opposed to a high-performance circuit. The thermal vias are assumed to be evenly distributed across the chip in the global phase (Figure 4.5 a). Our 3DIC design flow will create a design prototype with different tier-counts and thermal via densities to analyze and compare critical path delay, energy per cycle, and temperature. If several combinations result in similar performance with no temperature violation, we opt for the

pattern with lowest tier-count and the smallest thermal via density. The optimal performance obtained in this phase will be set as the performance target for the detail phase.

2. **Detail phase**: Determine the via-patterns for optimized power, delay and temperature once the tier-count and via density are provided.

Once the tier-count, via density and optimal performance target are determined, this flow will automatically tweak the via pattern to optimize temperature with three constraints. (a) no sub-region that contains clock buffers will be hot enough to cause hold violations; (b) no sub-region that contains the critical path will be hot enough to cause setup violations (a tolerance is allowed for, so the new critical path delay will be $delay_{(new)} = delay_{(old)}$ +/- tolerance); and (c) no sub-region is hot enough to increase its leakage to a considerably higher value (a tolerance like (b) is allowed for) so that the positive feedback between leakage and temperature will eventually cause the devices in this region to breakdown. We set the limit to be the point at which the leakage power becomes 15 times its room-temperature leakage. In this MITLL technology, this temperature limit is approximately 180°C. To facilitate physical design, the thermal via is designed as a via cell (a standard cell). If any of above problem is identified during detailed phase, thermal via cells and some of the standard cells will be relocated to eliminate the problem (Figure 4.5 b). Actually the detail thermal via design only needs to be performance once since we care for optimal performance. During the floorplan design of the detail phase, this methodology optimizes for wire-length, critical path delay and maximum temperature. A thermal design step will then evenly assign thermal via to each sub-region, but these vias cells are subject to move. After placement and CTS, A SAIF file is used to analyze power consumption of each cell. Hence, relatively accurate power consumption for each sub-

region can be obtained. Via cells and standard cells can be moved horizontally at this time to meet the three previously mentioned constraints. When the via patterns are determined, a electro-thermal coupling analysis will be performed to verify final performance and temperature.

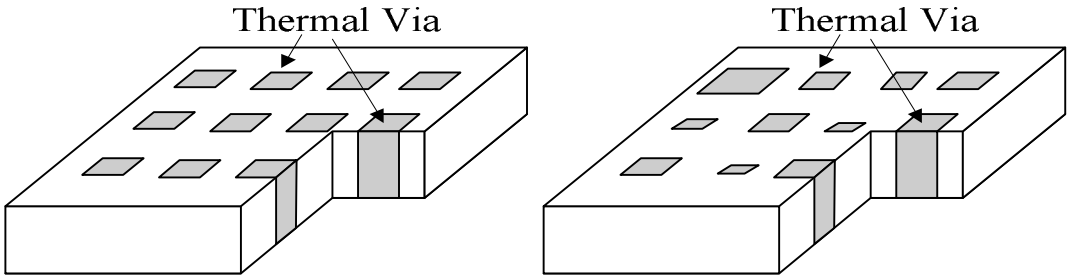

(a) Global phase (b) Detail phase

Figure 4.5: Thermal via Formation.

The two-phase thermal via pattern determination provides advantages on account of 1) simplicity of implementation (ease of writing the code); 2) absence of redundant thermal vias, which reduces chip area and improves system performance; and 3) quick prototyping which provides realistic targets to the detail phase for faster design time.

The 3DIC design methodology can be used for both floorplanning and detail design without any modification. The floorplanning phase is highly automated while the detail phase still needs some designer input to move blocks, insert thermal via, and synthesize the clock tree to avoid violations.

4.3 Reliability Analysis and Management

The reliability issue in this dissertation only refers to the timing reliability in the 3D system. Due to lateral and vertical temperature gradients, clock skew over the 3D chip may

vary alot compared to the uniform-temperature static timing analysis. Since larger clock skew will destroy the functionality of a system, a reliability analysis must be performed to ensure no race conditions exist.

Clock skew is defined as the difference between the maximum and minimum insertion delays of any clock path. Since all registers in the benchmark circuits use positive edge triggered D Flip-flops, only rise skew is analyzed. Figure 4.6 shows the pseudo code for clock skew analysis. During clock skew analysis, each clock path is traversed. When a clock buffer (a clock buffer can be either a buffer or an inverter) is found in the clock path, its location and the surrounding (operating) temperature are read in. Based on the operating temperature, a new clock buffer delay is calculated with equation (18), and this delay is added to the current clock path delay. The new skew is obtained as soon as all clock paths are traversed. This skew analysis is implemented in Python. It takes two inputs: 1). The temperature profile obtained after thermal analysis with detailed model; and 2). All the clock paths report from ***PrimeTime*** or ***Design Compiler***. The clock skew analysis considering temperature variation can be done within several seconds. The runtime is directly proportional to the number of clock buffers.

To be more robust, we use equations (34) and (35) for a hold time check, where t_{hold} is the D flip flop hold time, $t_{skew}(T_0)$ is the clock skew assuming a uniform temperature T_0, $t_{skew}(T_{matrix})$ is the clock skew considering the actual temperature profile, $t_{ck\text{-}q\text{-}min}$ is the minimum D flip-flop clock to output delay, and $t_{logic\text{-}min}$ is the short path delay reported by ***PrimeTime*** or ***Design Compiler*** at the nominal temperature T_0. This hold time check over constraints our design, but it is simpler and faster to be implemented in any digital system. Equation (34) is used to check the original hold time. If this condition is satisfied, the clock skew is updated considering

temperature gradients and equation (35) is then applied to the final hold time check with the new clock skew.

For each clock path
For each buffer in current clock path
Find the clock buffer operating temperature and calculate delay
Clock path delay = $\sum_{clock-path} clockbuffer\_delay(T)$
Skew = (max clock path delay) – (min clock path delay)

Figure 4.6: Clock skew analysis in presence of temperature gradient.

$$t_{hold} + t_{skew}(T_0) < t_{ck-q-min} + t_{logic-min} \quad (34)$$

$$t_{hold} + t_{skew}(T_{matrix}) < t_{ck-q-min} + t_{logic-min} \quad (35)$$

Chapter 5.

Case Study and Design Trade-offs

5.1 Introduction

The main advantage of 3D integration is it can significantly reduce wire-length. As discussed in Chapter 2, Zhang *et al* [33] showed roughly a 40% reduction in the lengths of the longest wires with a 30% reduction for the average wire. Das *et al* [22] showed an average wire-length reduction of 11% when minimizing inter-tier vias and 40% when minimizing wire-length. However, none of these studies included the effect of inter-tier vias. An inter-tier via and consumes all routing tracks around it, which can cause routability problems. It is essential to add this effect when considering wire-length reduction.

System timing and power are two other concerns. Though Zhang *et al* [33] and Das *et al* [22] showed 3D integration reduces power consumption as tier count increases, their analysis did not include the temperature effect. However, this temperature effect cannot be ignored because the temperature and leakage power form a positive feedback loop, which would possibly result in a "self-heating" effect as described by [62]. The temperature effect is captured in this work and will be discussed in the following sections in this chapter.

This chapter explains how tier count and thermal via density were set as variables to explore the compromises among timing, power and temperature. The 3DIC prototyping design

flow was used to implement many combinations of tier count and thermal via density to find the fastest clock cycle and best energy efficiency.

5.2 Benchmark Circuits

Two benchmark circuits were used in this thesis. The first is a Fast Fourier Transform (FFT) [25] implemented with the Winograd algorithm [26] that represents the low-power VLSI category. The second is a design based on the OpenRISC Platform System-on-Chip (ORPSOC) [47] [48] that represents the high-performance VLSI category.

The FFT design is an 8-point FFT calculation engine based on the Winograd algorithm [25]. The design works on complex numbers, which have 32-bit real and imaginary data in the floating point format. As per the Winograd algorithm, the design has 52 dedicated adders and 4 multipliers to calculate the 8 real and imaginary values of constituent frequencies. The inputs and outputs of this engine are registered and the rest of the design is combinational logic. This design is pipelined to improve throughput in the extended technology. The floating-point arithmetic units used in the design have been downloaded from the Free Floating-Point Madness website [74]. These are IEEE 754 standard compliant units. We use the single-precision version; all the floating point numbers thus have one sign bit, 8 exponent bits and 23 mantissa bits. Figure 5.1 shows the schematic of the design. An *R* in the signal name indicates the real component of the signal, whereas an *I* indicates the imaginary component. Signal names beginning with *x* are the inputs to the FFT, while those beginning with *X* are the outputs. Correspondingly, signals beginning with *k*, *s* and *t* are the outputs of Stages 1, 2 and 3

respectively. A Linear Feedback Shift Register (LFSR) test vector generator creates random test vectors as inputs of the FFT core. The FFT core processes the inputs and sends the outputs to the Build-in Self Test (BIST) block. The BIST block includes hash generator and comparator. It will trigger the "correct" signal if the FFT core outputs match what is expected, otherwise, an "incorrect" will be triggered.

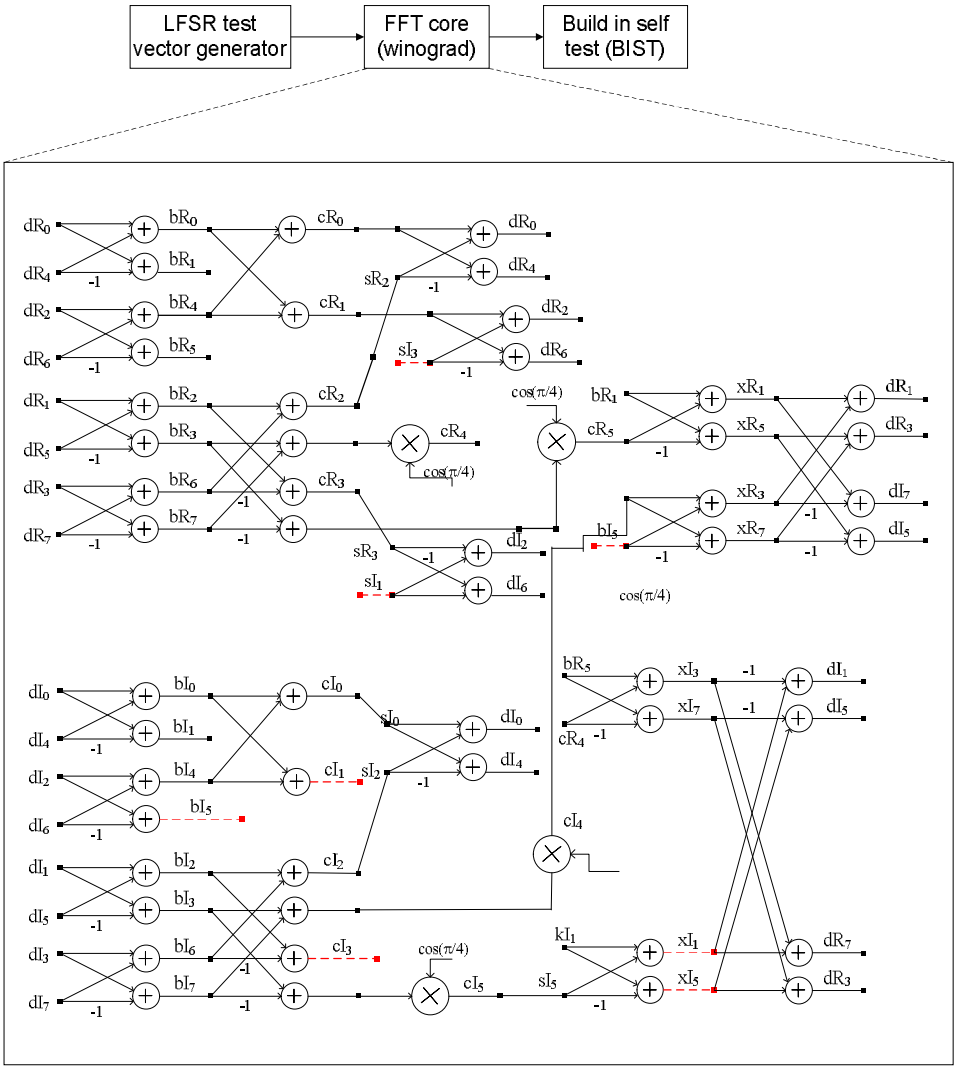

Figure 5.1: FFT architecture. [8]

[8] Thanks to Ambarish Sule for designing the FFT system and supplying this figure

The purpose of the ORPSOC is to allow for the rapid creation of OpenRISC based system-on-a-chip design with reduced verification time. The original ORPSOC uses several IP cores which includes: 32-bit OR1200 processor, Ethernet media access control (MAC), universal asynchronous receiver/transmitter (UART), SRAM interface, Audio interface, PS/2 interface, Simple VGA interface, and JTAG debug interface [75] and a 40*KB* embedded memory. The OpenRISC microprocessor is realized as a 5-stage integer pipeline with virtual memory support, which includes a central processing unit (CPU), instruction and data cashes, instruction and data memory management units (MMUs), and OpenRISC wishbone interfaces. More details of the OpenRISC CPU can be found in [47] and [75].

The original ORPSOC has been modified to include a second OpenRISC core that communicates with the same memories through a bus arbitration unit. Both OpenRISC cores are assumed to have similar workload and hence increases heat dissipation and throughput to form a more interesting test case [47]. The architecture of the OPRSOC benchmark used in this work is shown in Figure 5.2. A "Wishbone Traffic Cop" was used in the ORPSOC to easily interface the microprocessors and peripherals (in this work, memories) [47]. The wishbone traffic cop arbitrates the data transfer between a single wishbone-compliant master and a single wishbone-compliant slave, which serves as a glue between the memories and OR1200 CPUs. When either OR1200 CPU needs to access memory, a request will be sent to the wishbone traffic cop and a response (either accept or deny) will be send back to the OR1200 CPU. The memory controllers in ORPSOC control the memory for reading and writing. These controllers also monitor the activity of the wishbone traffic cop to conform to the wishbone standard. It is important to note that SRAM cells were not designed for this study. Estimates of SRAM area,

delays and power are based on a partial port of an 180nm bulk CMOS SRAM [77]. A total of 40*KB* of SRAM used in this study is composed of five 8*KB* SRAM cells.

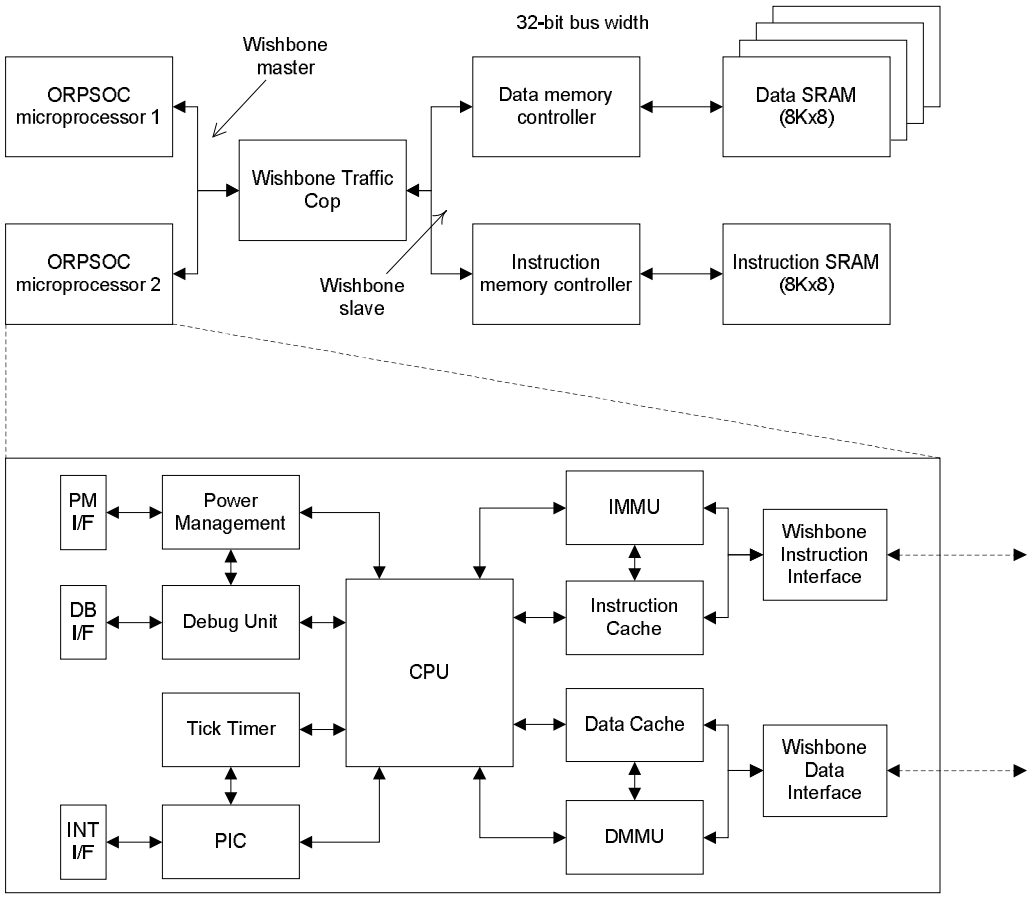

Figure 5.2: ORPSOC architecture. [9]

Both benchmarks are synthesized using design compiler with a standard cell library (see Appendix B) designed for the MIT Lincoln Lab 3D SOI technology. After synthesis, the FFT has 158*K* standard cells and the ORPSOC has 120*K* standard cells. Table 5.1 shows the summary of synthesized benchmark circuits, including the number of standard cells, total net

[9] Thanks to Kory Schoenfliess for creating this figure

count, and the total area of standard cells excluding the SRAM blocks. Table 5.2 shows the performance estimation of the 8*KB* SRAM module.

Table 5.1: Summary of synthesized benchmarks.

	Number of cells	Number of nets	Cell area
FFT	158*K*	243*K*	10.23 mm^2
ORPSOC	120*K*	168*K*	9.76 mm^2

Table 5.2: Area, timing, and power estimation of 8*KB* SRAM block.

	Area	Timing	Dynamic power	Leakage power
Value	1450 × 980 μm^2	4.85 *ns*	41.6 *pJ*/cycle	4.9 *mW*

5.3 Comparisons in The Original MIT Technology

It is interesting to know how much wire-length reduction and performance improvement can be achieved in an existing 3D technology with the proposed design flow. Wire-length is compared in terms of both longest wire-length and total wire-length. The longest wire-length determines the timing performance and the total wire-length determines the switching power. Inter-tier signal vias and thermal vias have great impact on wire-length because they are both placement and routing blockages.

The total power consumption is defined in (36) [63], where α is switching activity, C_L is wire load, f_{clk} is clock frequency and I_{SC} is short circuit current. With the reduction of total wire-length, C_L is also reduced proportionally, which will also reduce the transition time and thus reduce short circuit power [65]. In general, the short circuit power consumes approximately 10% of the total power consumption [66] and is considered to scale with the

dynamic power. Equation (36) is then re-written into (37) and the total power can be calculated once α is determined. However, the derivation of α is neither clear nor trivial. As described in Chapter 3, α in this work is specified in a SAIF file, and *Power Compiler* is used to estimate both dynamic and leakage power at the nominal temperature.

$$P_{total} = P_{switching} + P_{short-circuit} + P_{leakage} = aC_L \cdot Vdd^2 \cdot f_{clk} + I_{SC} \cdot Vdd + I_{leakage} \cdot Vdd \quad (36)$$

$$P_{total} = P_{switching} + P_{leakage} = aC_L \cdot Vdd^2 \cdot f_{clk} + I_{leakage} \cdot Vdd \quad (37)$$

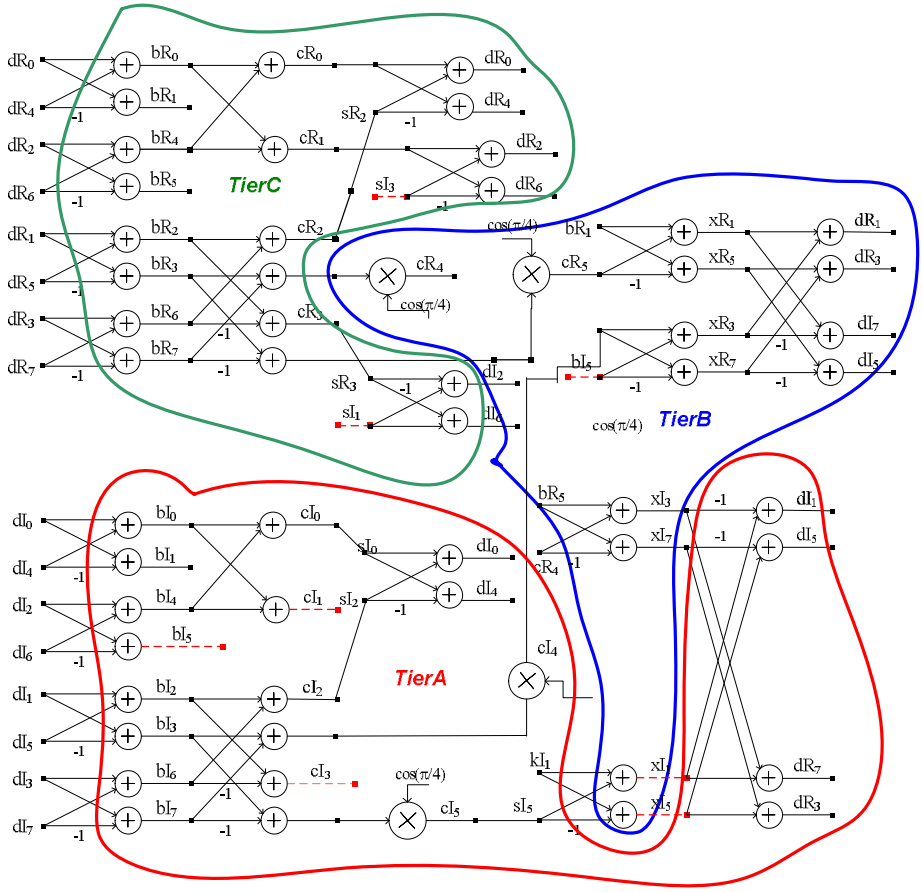

Figure 5.3: Three-tier partition schemes employed for FFT.

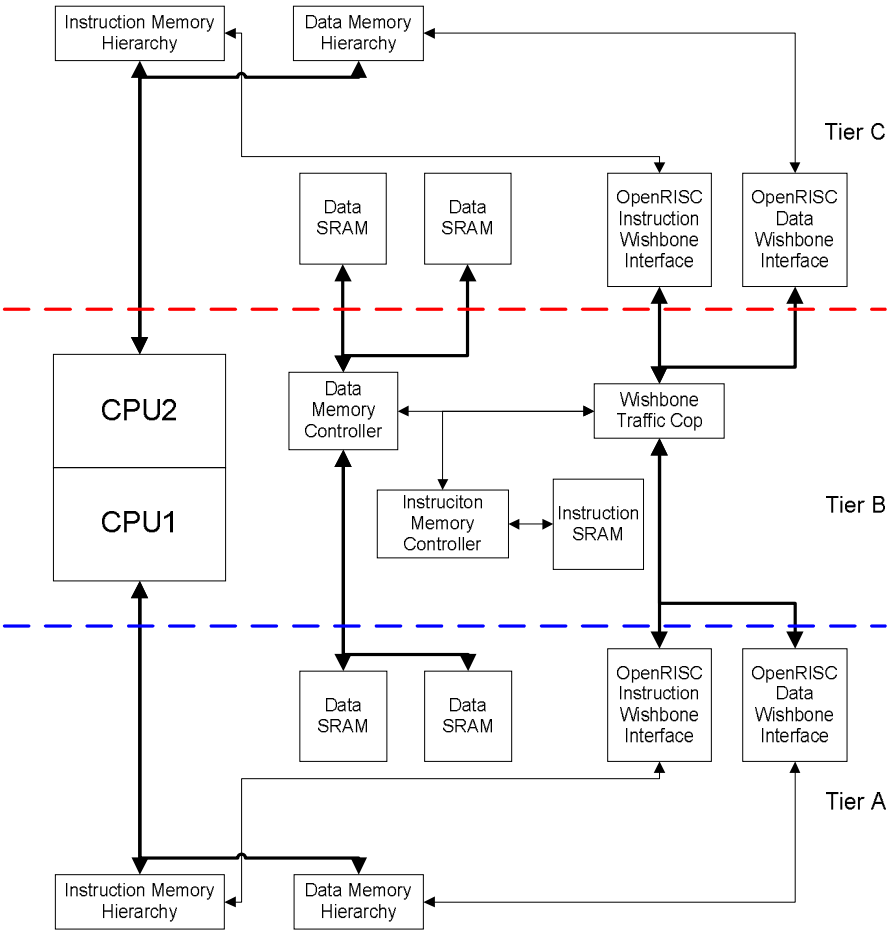

Figure 5.4: Three-tier partition schemes employed for ORPSOC[10].

Figure 5.3 and Figure 5.4 show the three-tier partition determined by *kMETIS* for both benchmark circuits. For the FFT, multipliers are assigned to bottom tiers to minimize temperature. For the ORPSOC, memories are assigned to different tier to optimize the performance.

Wire-length distributions in both 2D and 3D integration are compared. Figure 5.5 shows the histogram of wire-lengths from the 2D and 3D detail route results. Figure 5.5 (a) is the

[10] Thanks to Kory Schoenfliess for creating this partition figure.

histogram for the FFT benchmark circuit, and Figure 5.5(b) is the histogram for the ORPSOC benchmark circuit. One of the obvious similarities is that in both designs, the number of long interconnects is significantly reduced in 3D integration. However, the ORPSOC design has a bigger improvement in terms of wire-length.

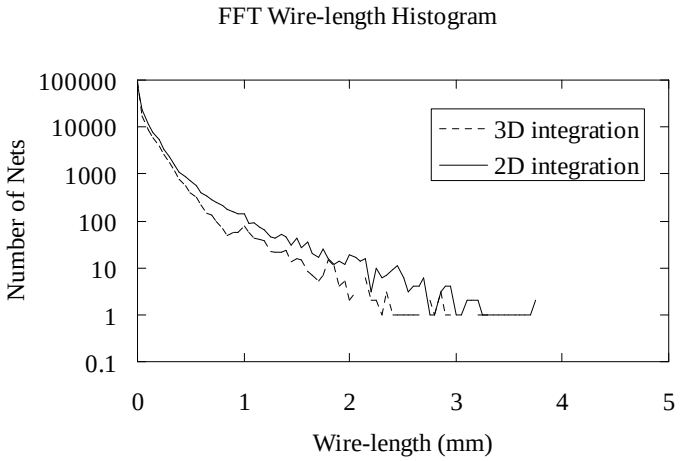

(a) Histogram of FFT wire-length distribution

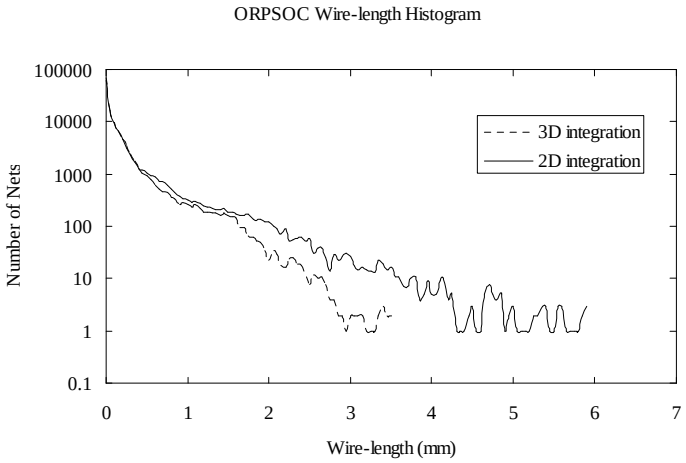

(b) Histogram of ORPSOC wire-length distribution

Figure 5.5: Histogram of horizontal wire-lengths from the 2D and 3D Place-and-Route results (bin size = 50 μm). (a) FFT, (b) ORPSOC

The performance improvement is shown in Table 5.3. Both static delay and power are reduced in 3D integration. 3D integration achieves better energy efficiency compared to its 2D counterpart. In accordance with the wirelength reduction shown in Figure 5.5, the ORPSOC design shows more performance improvement when compared to the FFT design even when temperature effects are considered.

Table 5.3: Comparison of performance in 2D and 3D integration (Original process technology).

Integration scheme	FFT			ORPSOC		
	Static delay	Power	E/cycle	Static delay	Power	E/cycle
2D	90.07 *ns*	294.0 *mw*	26.5 *nJ*	25.54 *ns*	3.86 *W*	98.6 *nJ*
3D	88.28 *ns*	250.4 *mw*	22.1 *nJ*	18.75 *ns*	3.74 *W*	70.1 *nJ*

5.4 Comparisons in The Extended MIT Technology

The original MIT 3D process only supports three tiers and three metal layers. This makes it less interesting because current state-of-the-art processing technology for 180 *nm* technologies use up to seven metal layers [11]. A smaller number of metal layers will result in a much smaller core utilization and longer wire-length, which will possibly make the performance of three-tier three-metal-layer 3D integration worse than the one-tier seven-metal-layer 2D integration. To overcome this limitation, this 3D design flow is extended to accommodate up to ten silicon tiers and seven metal layers. With the extended 3D integration, performance is characterized again and design trade-offs are examined. To make a more interesting electro-thermal test case, the FFT benchmark circuit is also pipelined to increase throughput for this study. In this section, FFT refers to the pipelined FFT.

Table 5.4 shows the core utilization and performance value for different number of metal layers implemented in the extended MIT technology with one tier (a 2D implementation). The core utility and performance do not change when more than 5 metal layers are used. So the optimal metal layer number is considered to be 5 in the modified technology. We will use the 5 metal layers per tier technology since it gives the optimum in a sense of performance and processing complexity.

Table 5.4: Core utilization and performance with different metal layers in 2D integration.

	Metal layers	Core utilization	Power	Timing	E/cycle
FFT	3	0.75	890 *mw*	33.9 *ns*	30.2 *nJ*
	5	0.95	809 *mW*	26.1 *ns*	21.1 *nJ*
	7	0.95	809 *mW*	26.1 *ns*	21.1 *nJ*
ORPSOC	3	0.72	3860 *mW*	25.5 *ns*	98.6 *nJ*
	5	0.95	3298 *mW*	17.8 *ns*	58.7 *nJ*
	7	0.95	3298 *mW*	17.8 *ns*	58.7 *nJ*

Figure 5.6 shows the histogram of the benchmark circuits implemented with different tier count but same metal layers (5 metal layers), where tier=1 represents the 2D integration. In this implementation, no thermal vias are used. As tier count increases, the number of long wires becomes smaller. However, 10 tiers and 7 tiers do not show much difference, which means 3D integration has a limit regarding horizontal wirelength reduction.

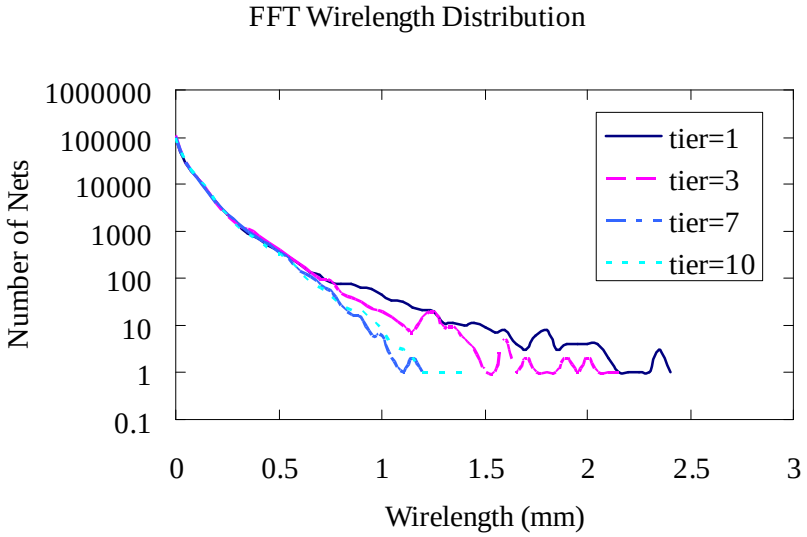

(a) FFT wirelength distribution

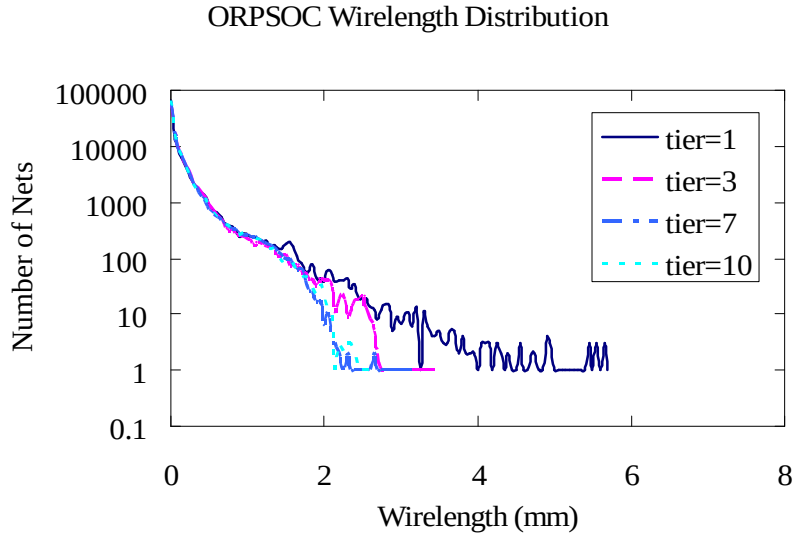

(b) ORPSOC wirelength distribution

Figure 5.6: Histogram of horizontal wirelength distribution with different tier count.

We compare the actual horizontal wirelength shown in Figure 5.6 with the theoretical horizontal wirelength predicted by [33] (which is shown in Figure 5.7). When tier count is low, the reduction is significant, regardless the metal layer used in the specific technology. So Figure 5.5 and Figure 5.6 tier=1 and tier=3 show a similar trend. This agreement remains true

till tier count reaches 7 in the actual experiment. Beyond that, horizontal reduction is no longer significant. This is because of the inter-tier vias. Since inter-tier vias are both placement and routing blockages, they create localized congestion. When the tier count is low, the total area occupied by inter-tier vias in a design is small compared to the standard cell area. However, as tier count increases, the inter-tier cuts increase as well, and they create a large amount of congestion. As congestion increases, the wirelength also increases, because we have to increase the chip area to make the design routable. On the other hand, higher congestion will require more detours for a net to reach its destination, which also increases the wirelength. For a more accurate wirelength prediction, the effect of inter-tier vias has to be considered.

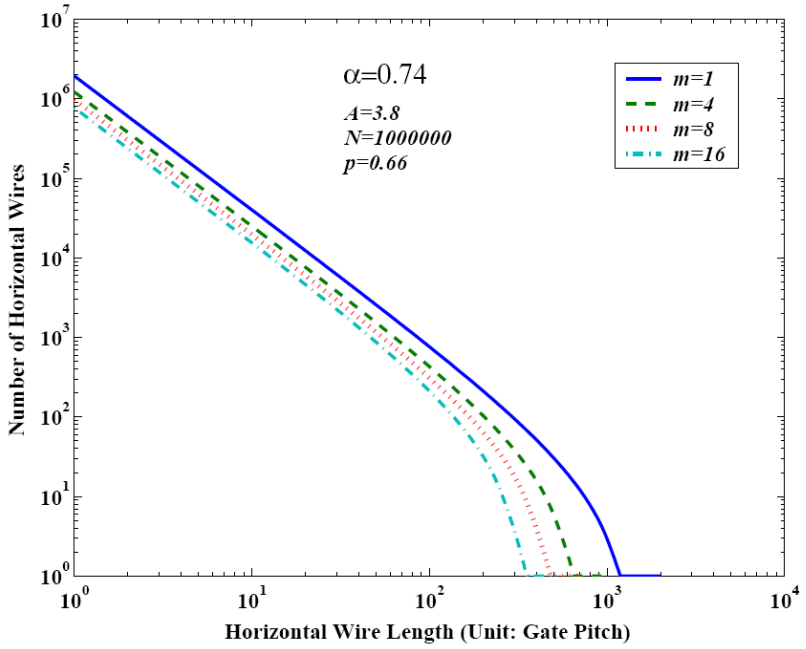

Figure 5.7: Horizontal wirelength distribution of R. Zhang's work.[11]

[11] This figure is from R. Zhang's work [33] for comparison between the theoretical analysis and actual experiments.

5.4.1 Design Trade-off Exploration

Stochastic estimates based on Rent's rule predict that total interconnect capacitance for a system will decrease as the number of tiers increases, up to a maximum of a 40% reduction for six tiers [33][32]. However, this research has overlooked the heat dissipation and the fact that thermal vias increase routing congestion. In this work, the trade-off between dynamic power and leakage power was explored by varying tier count and thermal via density.

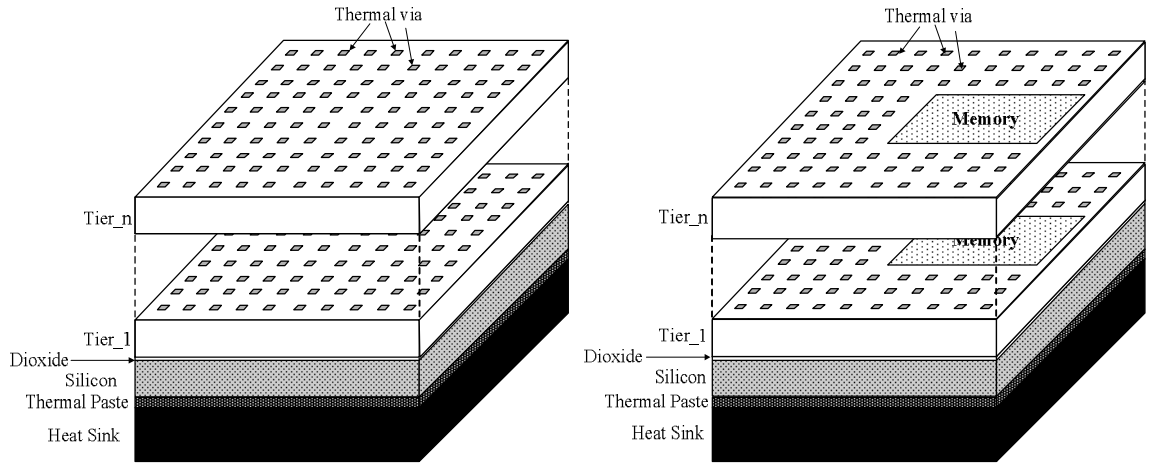

(a) FFT (b) ORPSOC

Figure 5.8: Thermal via insertion in (a) FFT, (b) ORPSOC.

Thermal via is implemented as via cells with same height of standard cells, and only 26% of the via cell area is actually filled by thermal via. The "density" in the following study refers to the ratio of total via-cell area to the die area. For simplicity, thermal cells are uniformly distributed across all tiers. An example of thermal via insertion during global phase for the benchmark circuits is shown in Figure 5.8. The maximum thermal via cell density is restricted to 20% of the total chip area. Figure 5.9 and Figure 5.10 show an analysis of the most congested tier for each design in the extended three-tier, five-metal-layer technology. Each

histogram shows the number of G-cells with various number of tracks assigned. G-cells with more than 40 tracks assigned are more than 100% congested. As shown in the figures, the designs become unroutable at higher via-cell densities [45].

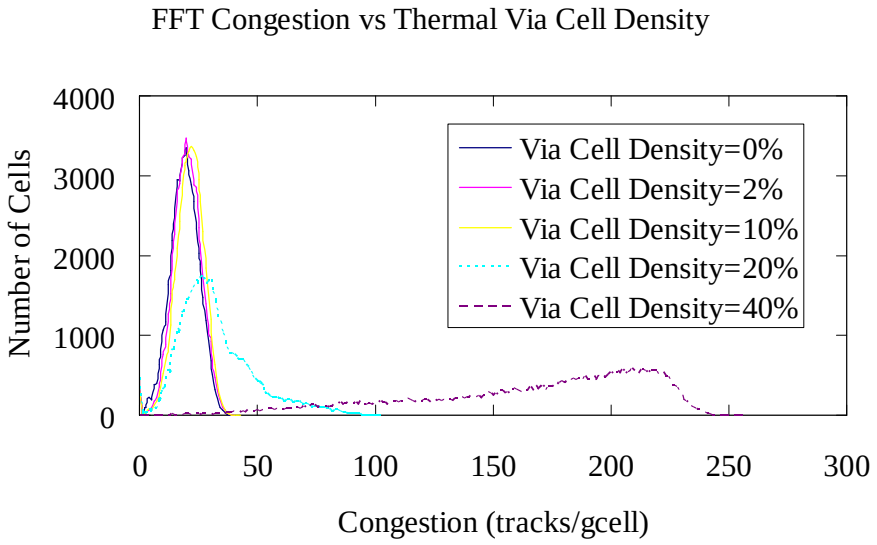

Figure 5.9: Congestion analyses with different thermal via density for FFT.

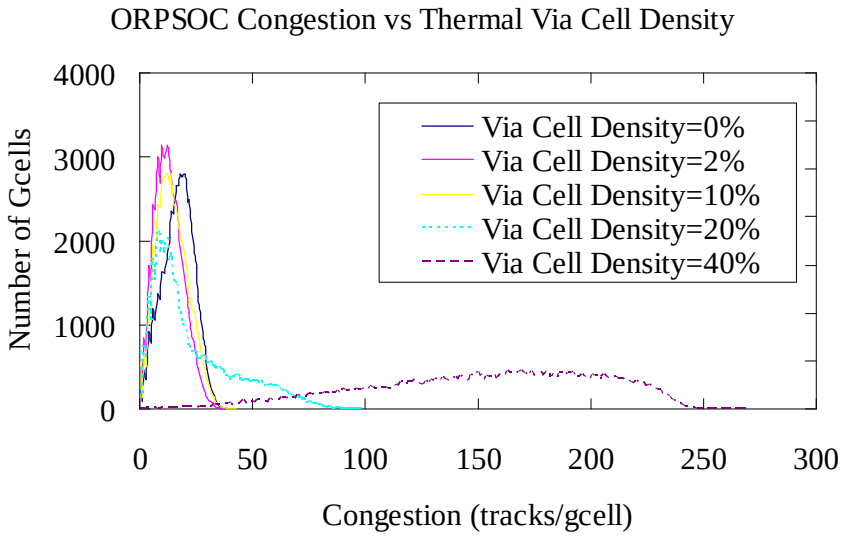

Figure 5.10: Congestion analyses with different thermal via density for ORPSOC.

In this section, we present our experiments with the pipelined FFT and ORPSOC circuits. All results are based on the flow illustrated in Figure 3.2, which is applied to designs with 1 to 10 tiers and thermal-via densities ranging from 0%-20%, using the 83-cell standard-cell library that was characterized with nominal device parameters for this work. As described in section 5.2, the memories used in the ORPSOC system were not completed in time for this study, and so an estimate based on an 8KB SRAM block is used. A total of five of these SRAM blocks are used, each having an estimated read/write delay of 4.8*ns*, dynamic power of 41.6*nJ*/cycle, and leakage power of 4.9*mW*. Since the memory delay is relatively small compared to the path delay, we place these memories in the upper tiers so that the most power-hungry and timing-critical blocks can be placed close to the heat-sink. To fairly compare the improvement brought by 3D integration, the benchmark circuits are implemented with 5 metal layers in 2D and the design summary of both circuits is repeated in Table 5.5.

Table 5.5: Design summary of benchmark circuits with 5 metal layers.

Design	Path delay	Power	Number of cells	Die area
FFT	26.1 *ns*	0.809 *W*	158 *K*	11.6 mm^2
ORPSOC	17.8 *ns*	3.298 *W*	120*K*+40KB SRAM	18.8 mm^2

The first comparison shows the impact of our iterative timing/power/temperature calculation methodology. Figure 5.11 shows the delay values for FFT and ORPSOC both considering and neglecting the temperature effect. This figure shows that delay-temperature coupling is less critical in the low-power (FFT) application, because it is cooler. Note that the increase in delay for the FFT with 6 tiers (Figure 5.11 (a)). This is due to a difficulty in partitioning the design into 6 tiers. As Figure 5.11 (b) shows, the delays of the ORPSOC

increase rapidly when the tier count exceeds eight, because the excessively high temperature begins to dominate.

The next comparison is to see how much improvement in energy-per-operation and path delay could be achieved with 3D integration compared to a conventional 2D approach. Table 5.6 shows the results with the corresponding number of tiers and thermal via density. The FFT design achieved the most improvement in both energy and delay (27% and 20% respectively) when using 5 tiers and no thermal vias. The ORPSOC design shows the most improvement with 4 tiers, and different thermal-via densities were shown to minimize energy and delay.

Table 5.6: Best energy/cycle and timing of 2D/3D integration.

Design		2D	3D	Tiers	Via density	improvement
FFT	E/cycle	21.2*nJ*	15.5*nJ*	5	0%	26.9%
	Delay	26.1*ns*	20.9*ns*	5	0%	19.9%
ORPSOC	E/cycle	58.7*nJ*	47.9*nJ*	4	2%	18.4%
	Delay	17.8*ns*	14.8*ns*	4	5%	16.9%

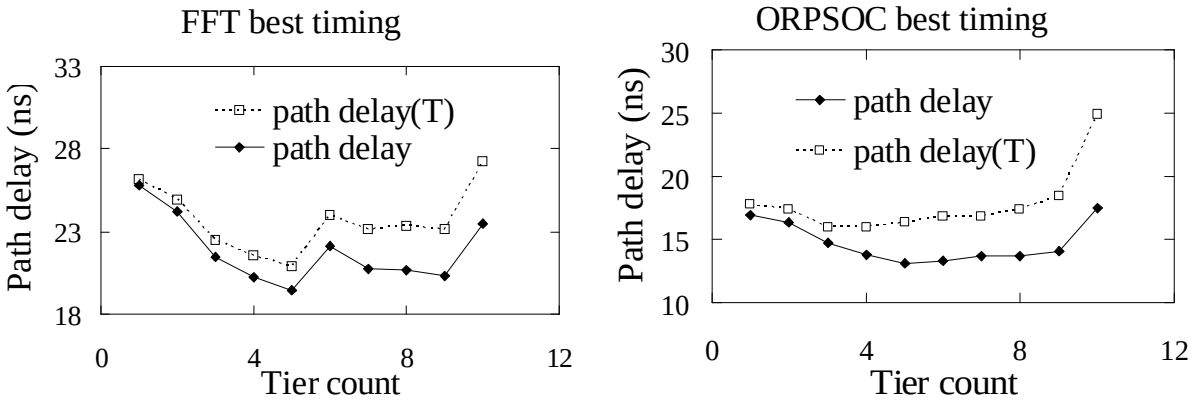

(a) FFT timing (b) ORPSOC timing

Figure 5.11: Best timing value with/without delay temperature coupling.

Figure 5.12 gives a more detailed picture of the design-space with plots showing how maximum temperature, total wire-length, path delays, and power of the benchmark circuits vary with the two variables (tier number and thermal via density). The darker areas represent smaller values, and the brighter areas represent large values. The energy-per-cycle numbers in Table 5.6 are determined by the product of power and delay. Figure 5.12 (a) and (b) give the maximum temperature trend of the FFT and ORPSOC. As can be seen from the figure, the maximum temperature is monotonic as tier number increases and thermal via density decreases. However, as thermal via area increases, the temperature drop in the ORPSOC is more pronounced than that in the FFT. This can be explained by the polynomial leakage power-temperature dependency. Because the ORPSOC design is hotter, the leakage power is more sensitive to an increase in thermal-via density. Though higher temperature usually causes longer path delays and in turn reduces dynamic power, it has a different effect in low-power and high-performance applications. Figure 5.12 (c) and (d) give the total wirelength of each combination of tier count and thermal via density. With the same thermal via density, wirelength decreases as tier count increases due to the proximity of blocks. Note that the figure is somewhat misleading, because it does not include vertical length, which is difficult to quantify in a meaningful way. Figure 5.12 (e) and (f) show the power trend for the FFT and ORPSOC. In the FFT design, the upper left portion of the graph is where the most power is consumed, where tier count is low and thermal via density is high. This is because the design consumes relatively little power, and the total temperature rise above ambient is small. Therefore, the power is mainly determined by interconnect wire-length (dynamic power), rather than leakage. Therefore, the total power consumption of the FFT varies with the number of tiers. However, when tier count is large, the total wirelength no longer decreases, and so the

power reduction is less pronounced. The case with the ORPSOC is different. Due to the relatively high temperature, the leakage tends to dominate, and the highest power is in the lower right of the graph. Therefore, the optimum energy-per-operation occurs with fewer tiers. For each tier count, as thermal via density increases, the total power consumption first decreases and the increases. This is because the leakage power dominates for low via densities, but dynamic power dominates for higher densities. In our experiments, the leakage power of ORPSOC varies from 2% to 30% of total power. Designers are faced with the trade-off between dynamic power and leakage power to achieve optimal system performance. One reason for the severe leakage power effect in the ORPSOC design is that the memories were placed on the top tiers, which has the highest temperature. Without using any thermal vias, these memories would consume approximately 25X more leakage power and run 60% slower with 10 tiers. If our critical path delay is close to the memory latency, we must be very careful to find the location for the memories. In such a case, we would not opt for a design with a large number of tiers, because of the additional difficulty associated with the heat removal. Figure 5.12 (g) and (h) show the timing trend of the FFT and ORPSOC. The optimal timing region of the FFT is when the tier number is equal to 5 and thermal via density is equal to 0, while the optimal timing region of the ORPSOC is when the number of tiers is approximately equal to 4 and thermal via density is equal to 5%.

Based on the observations made above, timing and power is mainly restricted by interconnect wirelength in a low power design. In a high performance design, both wirelength and temperature have a large impact.

Figure 5.12: Max temperature, power and timing with different silicon tier number and thermal via area

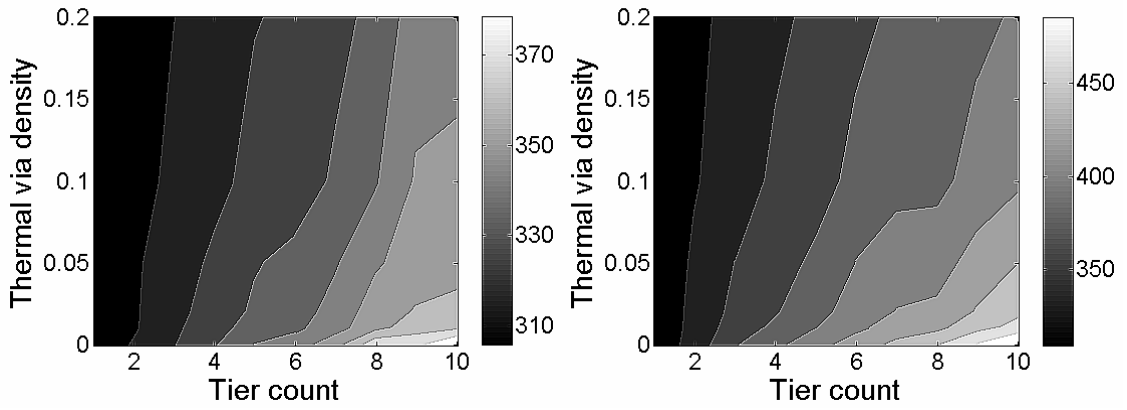

(a) FFT temperature (b) ORPSOC temperature

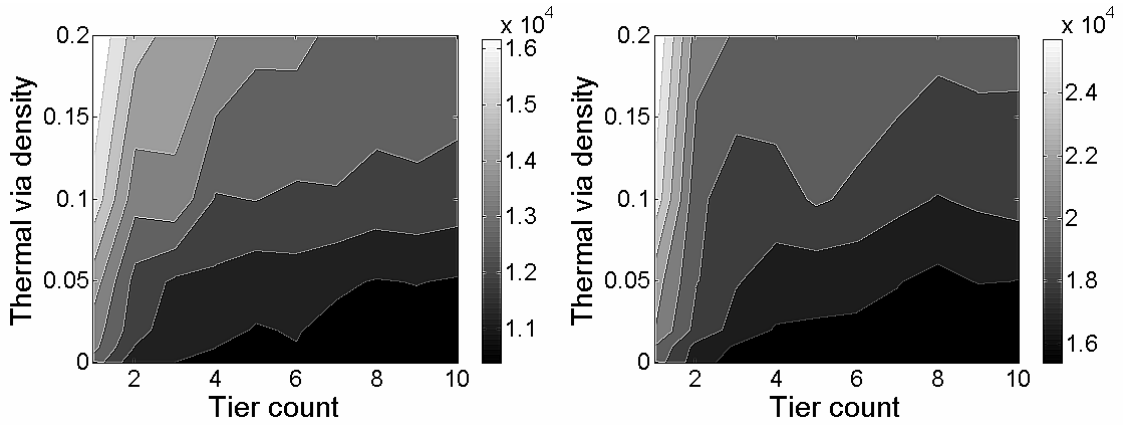

(c) FFT wire-length (d) ORPSOC wire-length

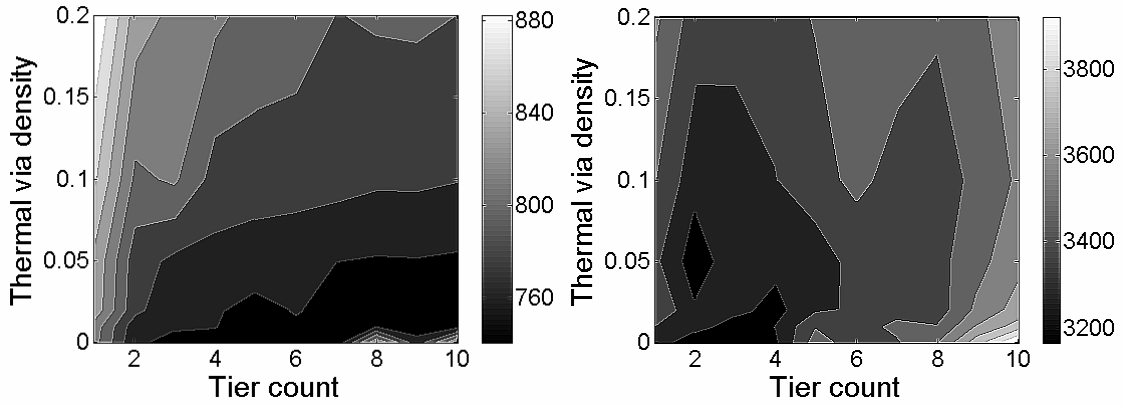

(e) FFT total power (f) ORPSOC total power

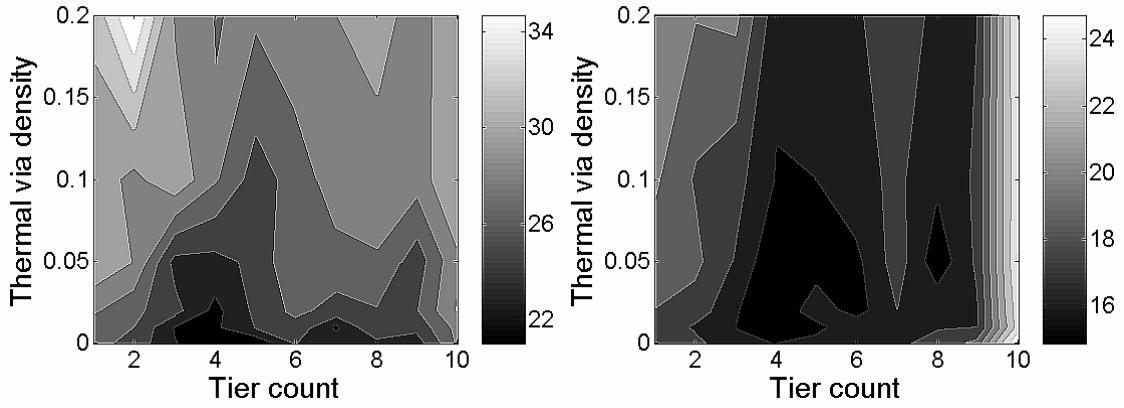

(g) FFT path delay

(h) ORPSOC path delay

5.4.2 Thermally Induced Timing Reliability Problem

In this thermal reliability analysis, different thermal simulation resolution will be used. The number of grid subdivisions on each side (called m in Figure 4.2 (a)) is varied between 1 and 50. To emphasize the clock buffer effects, we intentionally set an aggressive target with clock skew equal to 100*ps* and clock sink transition time equal to 100*ps*. In this case, a lot more clock buffers are inserted to match the target, and stacking of clock buffers on different tiers becomes inevitable. An analysis performed with a fine resolution shows that the clock buffers are the most significant heat sources, especially when these buffers on different silicon tiers are above one another. Under such circumstances, thermal vias can help to reduce the temperature of hotspots and improve reliability. Results also show that the coarse and fine analyses report similar delay and energy per cycle. However, the accuracy in measuring clock skew depends on the analysis resolution. In the thermal analysis, SPICE is used to simulate the equivalent electrical network shown in Figure 4.2 (a).

The two benchmark circuits, the FFT and ORPSOC, are analyzed with coarse and fine resolution simulations. The coarsest resolution treats each tier as a single block. The second coarsest resolution divides the entire design into 10 by 10 blocks per tier. The finest resolution divides the entire design into 50 by 50 blocks per tier, where each block has a size of approximately 40*um* by 40*um* in FFT and 50*um* by 50*um* in ORPSOC. Two methods are used for adding thermal vias: 1) No thermal vias are used, and 2) Thermal vias are inserted when a clock buffer is inserted. In the second scheme, two thermal via cells are inserted adjacent to the left and right side of the clock buffers, and 26% of the via cell area is consist of thermal via. In this analysis, the three-tier five-metal-layer technology is used for demonstration purpose.

Table 5.7 compares the timing, clock skew, maximum temperature and runtime analyzed with different resolutions implemented in the three-tier five-metal-layer technology. These experiments were run on a Sun Ultra-Sparc3 900MHz CPU with 16G memory. The seventh column of Table 5.7 gives the density of thermal via cells (TV density) compared to the die area that was inserted with the clock buffers. The '-' symbol corresponds to our first scheme in which no extra vias were inserted. From Table 5.7, we can see the trade-offs of analyzing a design with various resolutions. The two coarse analyses run two orders of magnitude faster than the fine analysis, however they fail to capture the hotspots and the hold-time violation in the case where no thermal vias are inserted. The fine analysis runs much slower and gives similar performance values, but it captures the hotspots and the hold time violation. In the FFT, the maximum error between coarse and fine simulation for delay and energy per cycle are 1.1% and 5.0%, respectively. When comparing the resolution, the clock skew varies by 10.4% without thermal vias but only varies 5% if thermal vias are inserted. In the ORPSOC, the maximum errors between the coarse and fine simulations for delay and energy per cycle are 8.2% and 4.0% respectively. The variation of clock skew simulated with different resolutions is reduced from 14.5% to 5.5% if thermal vias are inserted. The larger variation in delay and clock skew in the ORPSOC design is due to the memory. The temperature in memory regions is much lower than other regions filled with logic circuits. The coarsest analysis averages this effect and underestimates the path delay and clock skew.

Further investigation shows that the electro-thermal analysis threshold value defined in Figure 3.14 also affects the accuracy of our final performance. The result is shown in Figure 5.13. The horizontal axis is the temperature change threshold below which the iterative

calculation ceases. The solid line is the runtime and the dashed line is the maximum energy per cycle change between the final two consecutive simulations. As temperature threshold is below 0.1%, the accuracy only improves slightly but the runtime increases considerably. For this reason, the temperature threshold has been set to 0.1% for this study.

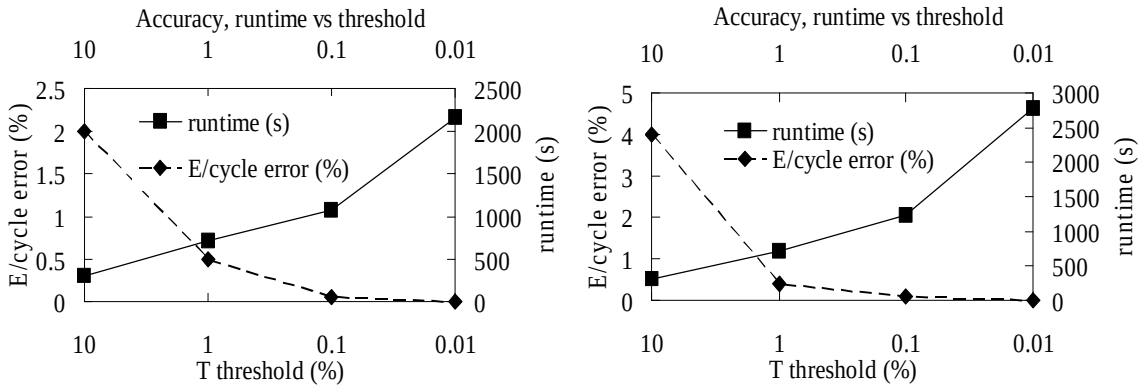

(a) FFT (b) ORPSOC

Figure 5.13: Accuracy, runtime vs temperature threshold.

Table 5.7: System performance verified with different resolutions.

Design	Resolution	Delay	E/cycle	Clock skew	Max T °*K*	TV cell density	Hold Violation?	Runtime
FFT	1×1	22.3 *ns*	21.0 *nJ*	183 *ps*	331	-	No	<1*s*
		22.5 *ns*	21.1 *nJ*	180 *ps*	328	0.9%	No	<1*s*
	10×10	22.4 *ns*	21.3 *nJ*	191 *ps*	340	-	No	3*s*
		22.5 *ns*	21.2 *nJ*	186 *ps*	336	0.9%	No	3*s*
	50×50	22.6 *ns*	22.1 *nJ*	202 *ps*	351	-	**Yes**	1078*s*
		22.5 *ns*	21.4 *nJ*	189 *ps*	341	0.9%	No	1078*s*
ORPSOC	1×1	15.2 *ns*	57.9 *nJ*	186 *ps*	365	-	No	< 1*s*
		15.1 *ns*	56.4 *nJ*	182 *ps*	357	1.5%	No	<1*s*
	10×10	16.3 *ns*	59.3 *nJ*	197 *ps*	392	-	No	3*s*
		16.1 *ns*	57.2 *nJ*	187 *ps*	374	1.5%	No	3*s*
	50×50	16.6 *ns*	60.3 *nJ*	213 *ps*	404	-	**Yes**	1078*s*
		16.2 *ns*	57.5 *nJ*	192 *ps*	381	1.5%	No	1078*s*

Figure 5.14 shows the maximum temperature profile of the FFT and ORPSOC when analysed with the finest resolution. The vertical axis is temperature in Kelvin. (0, 0) corresponds to the lower-left corner of each design and (1, 1) corresponds to the upper-right corner. Note that the spikes are gone in Figure 5.14 (d) compared toFigure 5.14 (c), it is evident that inserting thermal vias beside clock buffers have effectively eliminated the hotspots.

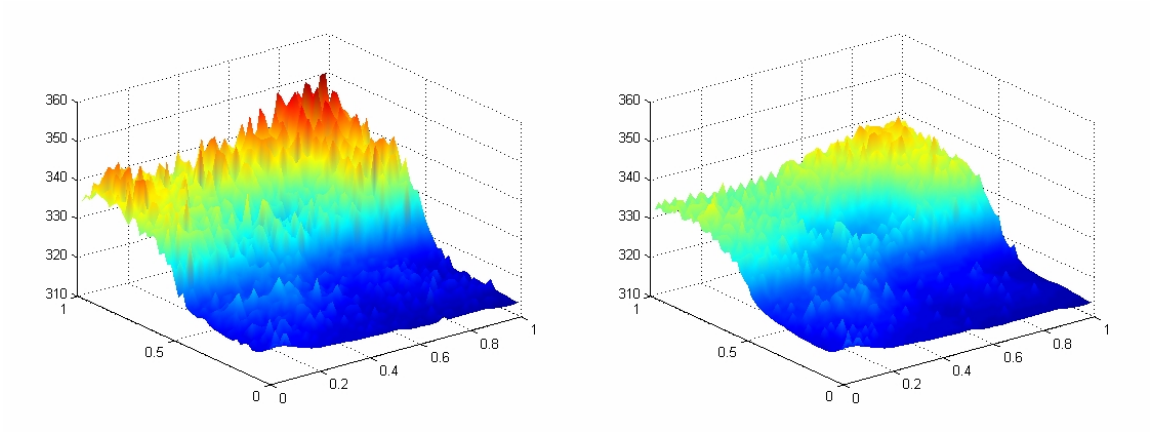

(a) FFT no thermal via

(b) FFT with thermal via

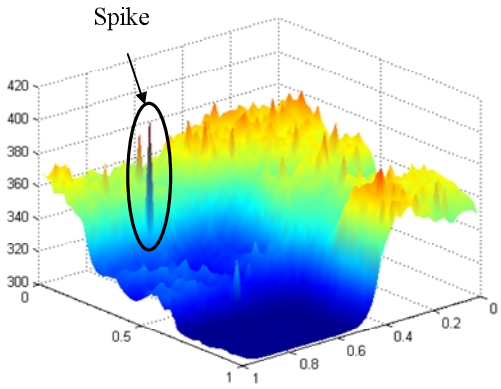

(c) ORPSOC no thermal via

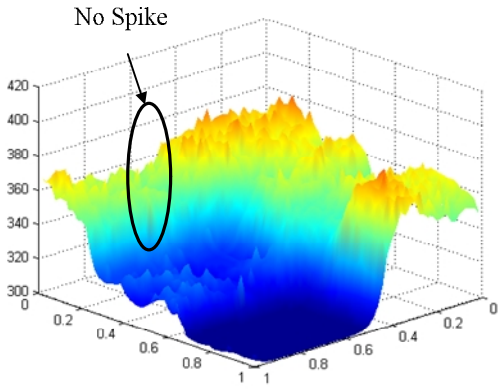

(d) ORPSOC with thermal via

Figure 5.14: Full chip temperature profile with three-tier, five-metal-layer technology; 50x50 blocks per tier.

5.5 Performance Trends of 3DICs

System performance improves as technology advances. One technology generation usually results in a 17%-25% improvement in clock frequency [11] with conventional IC processing technology. The performance trend in 3DIC also requires investigation because it gives a roadmap of future 3DICs. However, due to the lack of real technology, a Predictive Technology Model (PTM) [67] is used to create the standard cell library in 90*nm* and 45*nm* FDSOI technology nodes. Previous work has predicted the performance trend of 3DICs [33], but have not taken into account the increased temperatures and leakage power. As technology advances, the power density increases considerably and hence the leakage-thermal coupling becomes more significant.

Based on the thermal model shown in Figure 4.2 (b) and the same design flow discussed in Chapter 3, a performance trend of 3DIC is studied. However, due to the higher power density in future technologies, the temperature in 3DICs will quickly rise to a point that deep sub-micron (DSM) devices will break. In this case, fewer tiers can be integrated into a 3D system.

5.5.1 Temperature Dependency in New Technology

Using the same delay and leakage model derived in (18) and (19), we derive the new parameters for 90*nm* and 45*nm* technology node in the temperature range of 0ºC~150ºC. The highest temperature is reduced to 150ºC because simulation shows that the 45*nm* devices break down when the temperature exceeds this limit. For temperatures exceeding 150ºC, we will assume that the circuit fails to operate properly.

Table 5.8 shows k_0, γ, β, and α values for the standard cell library implemented in different technologies. We have performed extensive simulations on the standard cell library used in this study to test the accuracy of our model (18). In section 3.6.1.1, we showed that the largest disparity between SPICE and our model in the 180*nm* FDSOI technology node is an XOR gate with an overestimated delay error of 11.6%. In 90*nm* and 45*nm* FDSOI technologies, the largest error between SPICE and our model is still the XOR date, with overestimated delay errors of 11.3% and 12.5% respectively.

Table 5.8: Coefficients of delay temperature dependency model (18).

Technology	Edge	k_0 (*mV/K*)	γ (*mV/K*)	β	α
180*nm*	Rise	1.2	0.003	1.5	1.3
	Fall	1	0.003	1.5	1.4
90*nm*	Rise	1.5	0.003	1.5	1.3
	Fall	1.5	0.003	1.5	1.3
45*nm*	Rise	1.6	0.0035	1.5	1.3
	Fall	1.6	0.0035	1.5	1.3

Still using a curve fitting technique, we find values for the coefficients a_1, a_2 and a_3 in (19), for each standard cell and derive the average values for our library to estimate the leakage power with temperature as a variable. The average values for each technology are given in Table 5.9.

Table 5.9: Coefficients of leakage model (19).

coefficient	a_1	a_2	a_3
value (180*nm*)	0.0226	0.00033	1.77E-6
value (90*nm*)	0.03147	0.00038	1.87E-6
value (45*nm*)	0.02728	0.0018	3.0E-5

5.5.2 Performance Trend in New Technologies

We have implemented the benchmark circuits FFT and ORPSOC with three-tier, five-metal-layer 3D integration processing technology in the 3DIC design flow discussed in this work, and carried out a performance analysis with RC parasitics extracted with the wire width and space from the PTM and the ITRS roadmap. Through a proper iteration technique, the final performance values are found for both designs. The reason we only use three tiers is that the maximum temperature of ORPSOC exceeds the 150°C limit in 45*nm* technology. Figure 5.15 and Figure 5.16 show how the timing and clock skew of the FFT and the ORPSOC change as the feature size shrinks. Table 5.10 and Table 5.11 give more detailed information of the FFT and the ORPSOC performance in various technology nodes.

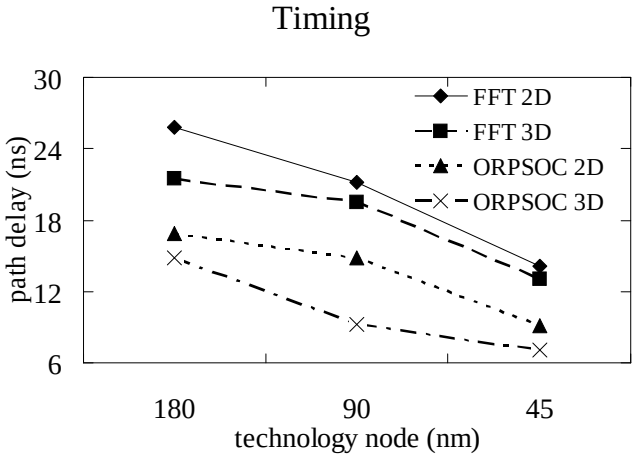

Figure 5.15: Timing of benchmark circuits vs. feature size.

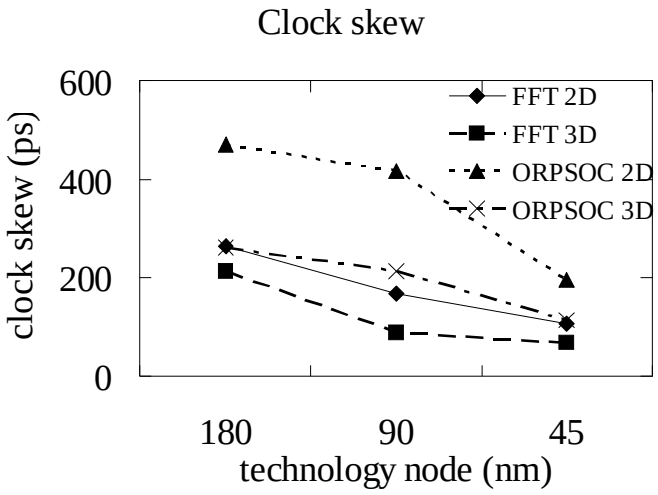

Figure 5.16: Clock skew of benchmark circuits vs. feature size.

Table 5.10: Performance summary of (a) 2D, and (b) 3D

tech	Design	Timing	Power	A(mm^2)	Max T
180nm	FFT	25.81 *ns*	1050 *mW*	11.61	39 ºC
	ORPSOC	17.80 *ns*	3298 *mW*	18.66	62 ºC
90nm	FFT	21.14 *ns*	439 *mW*	4.18	41 ºC
	ORPSOC	14.82 *ns*	1944 *mW*	6.76	82 ºC
45nm	FFT	14.16 *ns*	334 *mW*	2.32	45 ºC
	ORPSOC	9.16 *ns*	1215 *mW*	3.26	98 ºC

(a) 2D integration performance summary

tech	Design	Timing	Power	A(mm^2)	Max T
180nm	FFT	21.49 *ns*	952 *mW*	11.61	39 ºC
	ORPSOC	14.78 *ns*	2933 *mW*	18.66	62 ºC
90nm	FFT	19.54 *ns*	394 *mW*	4.18	41 ºC
	ORPSOC	10.02 *ns*	1551 *mW*	6.76	82 ºC
45nm	FFT	13.01 *ns*	256 *mW*	2.32	45 ºC
	ORPSOC	8.33 *ns*	1029 *mW*	3.26	98 ºC

(b) 3D integration performance summary

Table 5.11: Clock skew and power in 2D/3D integration.

technology	Design	Clock skew (*ps*)			Clock power (*mW*)		
		2D	3D	Improvement	2D	3D	Improvement
180*nm*	FFT	264.3	213.4	19.25%	240.5	201.9	16.05%
	ORPSOC	469.6	260.6	44.51%	838.9	482.9	42.44%
90 *nm*	FFT	167	88.8	46.83%	61.0	44.6	26.89%
	ORPSOC	317.6	213.0	32.93%	694.2	419.1	39.63%
45 *nm*	FFT	106.3	68.1	35.94%	66.4	48.3	27.26%
	ORPSOC	194.9	115.1	40.94%	798.1	487.9	38.87%

As can be seen, the total power consumption of 3D integration is significantly reduced due to the power saving on clock tree and switching power. In the 45*nm* technology node, the ORPSOC design consumes 13X more leakage power than its nominal value with only three tiers. Even when taking this increase in temperature and leakage power into account, however, the delay reduction is in the range of 8%-32%, and total power reduction is in the range of 9%-23%. Therefore, we can say that 3D integration with 3 tiers can improve the system timing by up to the equivalent of two technology generations. However, due to the dramatic temperature increase in 45*nm* technology, we do not analyze more than 3 tiers. In this performance trend study, thermal vias were not utilized, as further experiments showed that thermal vias did not help improve performance.

5.6 Impact of New Processing Technology

The 3D inter-tier signal vias in the original technology create undesirable exclusion regions for devices (Figure 5.17 (a)). This exclusion region consumes a lot of routing resources and hence reduces core utilization and increases wirelength. A new backside metal scheme

shown in Figure 5.17(b) has been proposed to overcome the inter-tier via exclusion region disadvantage. In this new scheme, 3D vias land on the backside metal rather than metal 1. Then the backside metal is connected to metal 1 with a normal sized via. Though this intra-tier via cannot penetrate active islands, it can be put within a standard cell between the gaps in the active layer. In general, the size of the intra-tier via is much smaller than the inter-tier via, and the intra-tier via does not require an exclusion region for metal layers. The primary advantage of this scheme is to allow the stacking of vias to improve vertical interconnect. However, the benchmark circuits used in this study were not limited by vertical interconnect, so we would expect this change to yield only a slight improvement.

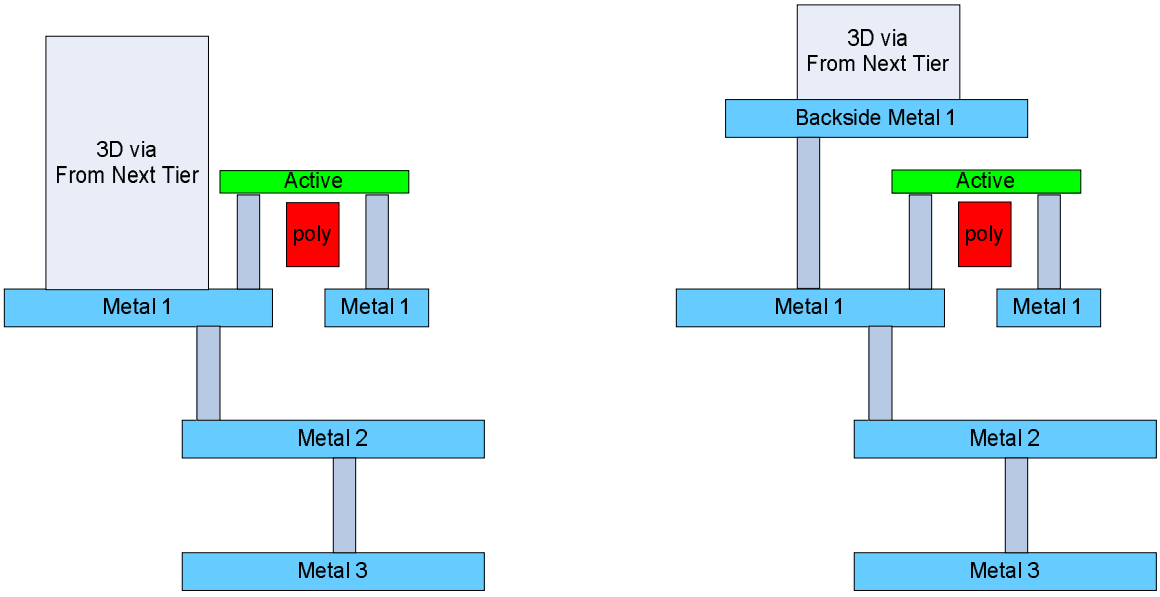

(a) Original scheme (b) Backside metal scheme

Figure 5.17: 3D via formation. (a). original scheme, 3D vias land metal1, creates exclusion region for devices; (b). backside metal scheme, 3D vias land on backside metal 1 rather than on the underside of metal 1.

A brief study has been carried out to evaluate the improvement brought by the new 3D via formation scheme. We implemented the two benchmark circuits in the three-tier five-metal-layer technology. Table 5.12 shows the area, total wirelength, and longest wirelength reduction in the new 3D via formation scheme compared to the original scheme. As expected, the new scheme shows a small improvement.

Table 5.12: Area and wirelength improvement with new 3D via formation scheme.

Design	Area reduction	Total wirelength reduction	Longest wirelength reduction
FFT	1.88%	1.93%	2.31%
ORPSOC	1.28%	0.83%	1.24%

Chapter 6.

Conclusion and Future Consideration

6.1 Conclusion

This work has successfully demonstrated the feasibility of 3DIC implementations with existing technology and tools. This work described a design flow with commercial tools glued together with a set of Python and Tcl scripts to support novel 3DIC designs. Using this flow, we designed two chips representing low-power and high-performance circuits respectively.

Design trade-offs of 3DIC are studied based on the design methodology we developed. As expected, heat is the fundamental limitation of performance improvement in 3DICs. Thermal vias can be used to remove heat, but they create routing congestion, which leads to a trade-off between leakage power and dynamic power. The extended technology explores the trade-offs based on the two benchmark circuits. Contrary to some research that emphasizes temperature control in 3DIC, we have pointed out that the overuse of thermal vias cannot benefit 3DIC system performance due to the increase of wirelength. In low power design, the temperature gradient across tiers is not significant. For these designs, thermal vias do not help much. In the high performance design, the interaction among timing, power, and temperature is more pronounced. In such cases, careful thermal via placement can effectively reduce the temperature rise within the 3D chip. Since the leakage power is exponentially dependent on temperature, thermal vias help reduce leakage power as well as maximum temperature rise.

This work has also shown that 3DIC is an attractive way to improve system performance. For the FFT example, a maximum of 27% reduction in energy per clock cycle was achieved along with a 20% improvement in speed. In the ORPSOC design, 3D integration achieves a maximum of 18% reduction in energy per clock cycle when we optimize for energy efficiency. The maximum delay reduction achieved was 17% in the ORPSOC design when we optimized for performance (timing).

A performance trend study showed that the thermal issue will become more and more severe as the technology node shrinks. Smaller devices (smaller channel length) result in higher power density and are more likely to break down at same temperature compared to a larger device (larger channel length). In a word, the 3D integration discussed in this work is limited by heat dissipation. Not may tiers can be integrated into one chip if there is no innovative heat removal technology.

Based on the exploration of this study, the most significant power saving of 3DIC comes from three parts: (1). clock tree power saving due to smaller chip area; (2). repeater power saving due to shorter wirelength; and (3). switching power saving due to shorter wirelength. With more efficient design methodology and design tools, more power saving can be expected from 3D integration and hence alleviate the heat removal difficulty.

6.2 Future Consideration

The quality of 3DIC design can be further improved by introducing advanced packaging technology and more powerful CAD tools. With current packaging technology, only one heat

sink is available and the major heat removal is due to conduction. Novel packaging methods that allow multiple heat sinks or other heat removal techniques (for example, convection) will further boost the 3DIC performance.

Incorporating emerging thermally aware 3DIC design tools may further improve the design quality of 3DICs. According to an internal discussion with researchers from University of Minnesota, a 3D placer is able to achieve another 11% reduction in total wirelength of the FFT benchmark circuit compared to our result. It is desirable to incorporate these new design tools in a mature design flow to verify their efficiency and reliability. The 3DIC design and verification methodology developed in this work provides a suitable vehicle to test the new design tools.

Another reliability issue, signal integrity, is more prominent in 3DIC due to the larger temperature gradient in both vertical and horizontal directions compared to conventional 2D integration. The driving capability of the cells in high temperature regions degrades and glitches are more likely to occur on the nets driven by these "hot" cells. A more detailed noise analysis technique needs to be developed to address this problem.

References

[1] Cadence Design System. “SignalStorm ® Library Characterizer User Guide”

[2] Metis website, available at http://www-users.cs.umn.edu/~karypis/metis/

[3] A. Rahman, and R. Reif, “Thermal analysis of three-dimensional (3-D) integrated circuits (ICs)”, *Proceedings of the IEEE International Interconnect Technology*, June 2001 Page(s):157 – 159

[4] B. Black, D.W. Nelson, C. Webb, and N. Samra, “3D processing technology and its impact on iA32 microprocessors”, *VLSI in Computers and Processors*, 2004, Oct. 2004, Page(s):316 – 318

[5] B. Goplen, and S. Sapatnekar, “Thermal Via Placement in 3D ICs”, *International Symposium on Physical Design*, 2005, Page(s): 167 – 174

[6] Ching-Te Chuang, Pong-Fe Lu, and C. J. Anderson, “SOI for digital CMOS VLSI: design considerations and advances”, *Proceedings of the IEEE*, April 1998, Page(s):689 – 720

[7] W. R. Davis, “A Hierarchical, Automated Design Flow for Low-Power, High-Throughput Digital Signal Processing IC’s”, *PhD thesis*, spring 2002, UNIVERSITY OF CALIFORNIA, BERKELEY

[8] G. Chandra, P. Kapur, and K. C. Saraswat, “A methodology for the interconnect performance evaluation of 2D and 3D processors with memory”, *Proceedings of Interconnect Technology Conference*, June 2002, Page(s):164 – 166

[9] B. Goplen, and S. Sapatnekar, “Efficient Thermal Placement of Standard Cells in 3D ICs using a Force Directed Approach”, *International Conference on Computer-Aided Design*, Nov. 2003, Page(s): 86-89.

[10] I. Meikei, K.W. Guarini, V. Chan, K. Bernstein, R. Joshi, J. Kedzierski, and W. Haensch, “Three dimensional CMOS devices and integrated circuits”, *Proceedings of Custom Integrated Circuits Conference*, Sept. 2003, Page(s):207 – 213

[11] ITRS Roadmap, available at http://www.itrs.net

[12] J. Burns, L. McIlrath, C. Keast, C. Lewis, A. Loomis, K. Warner, and P. Wyatt; "Three-dimensional integrated circuits for low-power, high-bandwidth systems on a chip", *IEEE International Solid-State Circuits Conference*, Feb. 2001, Page(s):268 - 269

[13] J. Cong, W. Jie, and Z. Yan, "A Thermal-Driven Floorplanning Algorithm for 3D ICs", *IEEE International Conference on Computer Aided Design*, 2004. Nov. 2004, Page(s):306 – 313

[14] J. Stine, "The Illinois Institute of Technology System-on-Chip Design Flow". Available at http://www.ece.iit.edu/vlsi/scells

[15] K. Banerjee, S. J. Souri, P. Kapur, 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*, May 2001, Page(s):602 – 633

[16] A. Bellaouar, A. Fridi, M. J. Elmasry, and K. Itoh, "Supply voltage scaling for temperature insensitive CMOS circuit operation", *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, Mar 1998, Page(s): 415-417

[17] K. Warner, J. Burns, C. Keast, R. Kunz, D. Lennon, A. Loomis, W. Mowers, and D. Yost, "Low-temperature oxide-bonded three-dimensional integrated circuits", *IEEE International SOI Conference*, Oct. 2002, Page(s):123 – 124

[18] R.S. Li, "Optimization of Thermal Via Design Parameters Based on an Analytical Thermal Resistance Model," *Thermal and Thermomechanical Phenomena in Electronic Systems*, 1998, Page(s): 475-480

[19] MIT Lincoln Labs 3D Design Notes

[20] P. Saxena, N. Menezes, P. Cocchini, and D.A. Kirkpatrick, "Repeater scaling and its impact on CAD", *IEEE Transactions on Computer Aided Design of Integrated Circuits and Systems*, April 2004, Page(s):451 – 463

[21] R. Reif, A. Fan, Kuan-Neng Chen, S. Das, "Fabrication technologies for three-dimensional integrated circuits", *IEEE International Symposium on Quality Electronic Design*, March 2002 Page(s):33 – 37

[22] S. Das, A. Chandrakasan, and R. Reif, "Design tools for 3-D integrated circuits", IEEE Proceedings of the ASP-DAC 2003. *Asia and South Pacific Design Automation Conference*, Jan. 2003 Page(s):53 – 56

[23] S. Das, A. Chandrakasan, and R. Reif, "Timing, energy, and thermal performance of three-dimensional integrated circuits", *IEEE Great Lakes Symposum on VLSI*, 2004

[24] S. Im, and Banerjee, "Full chip thermal analysis of planar (2-D) and vertically integrated (3-D) high performance ICs", *IEDM 2000*, Page(s): 727-730

[25] S. Winograd, "On Computing the Discrete Fourier Transform", *Mathematics of Computation*, vol. 32, no. 141, 1978, Page(s):175-199

[26] S. M. Alam, D. E. Troxel, and C. V. Thompson, "A comprehensive layout methodology and layout-specific circuit analyses for three-dimensional integrated circuits", *IEEE International Symposium on Quality Electronic Design*, March 2002 Page(s):246 – 251

[27] S. W. Sun, and P. G. Y. Tsui, "Limitation of CMOS supply-voltage scaling by MOSFET threshold voltage variation", *IEEE Journal of Solid-State Circuits*, Aug. 1995, Page(s) 947-949

[28] K. Kanda, K. Nose, H. Kawaguchi, and T. Sakurai, "Design impact of positive temperature dependence on drain current in sub-1-V CMOS VLSIs", *IEEE Journal of Solid-State Circuits*, Oct. 2001, Page(s): 1559-1564

[29] UCLA ASIC Design Flow, available at http://eda.ee.ucla.edu/EE201A-04Spring/

[30] T. Y. Chiang, K. Banerjee and K. C. Saraswat, "Compact Modeling and SPICE-Based Simulation for Electrothermal Analysis of Multilevel ULSI Interconnects", *IEEE International Conf of Computer Aided Design*, Nov. 2001, Page(s): 165-172

[31] M. Y. Lanzerotti, G. F. Fiorenza, and R. A. Rand, "Interpretation of Rent's Rule for Ultralarge-Scale Integrated Circuit Designs, With an Application to Wirelength Distribution Models", *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, Dec. 2004, Page(s):1330 – 1347

[32] Rongtian Zhang, K. Roy, Cheng-Kok Koh, and D.B. Janes, "Exploring SOI device structures and interconnect architectures for 3-dimensional integration", *IEEE Design Automation Conference, 2001*, Page(s):846 – 851

[33] Rongtian Zhang, K. Roy, Cheng-Kok Koh, and D. B. Janes, "Power trends and performance characterization of 3-dimensional integration for future technology generations", *IEEE International Symposium on Quality Electronic Design*, March 2001 Page(s):217 – 222

[34] D. Edelstein et al, "Full copper wiring in a sub-0.25um CMOS ULSI technology", in IEDM Tech Dig. 1997, Page(s): 773-776

[35] W. Liao, L. He, and K. M. Lepak, "Temperature and supply voltage aware performance and power modeling at microarchitecture level", *IEEE transaction on computer-aided design of integrated circuits and systems*, vol. 24, issue 7. July 2005, Page(s): 1042-1053

[36] G. Groeseneken, J. P. Colinge, H. E. Maes, J. C. Alderman, and S. Holt "Temperature Dependence of Threshold Voltage in Thin-Film SOI MOSFET's", *IEEE Electron Device Letters*, Vol. 11, No. 8, Aug. 1990, Page(s): 329-331

[37] H. Su, F. Liu, A. Devgan, E. Acar, and S. Nassif, "Full Chip Leakage Estimation Considering Power Supply and Temperature Variations", *IEEE International Symposium on Low Power Electronics and Design*, Aug. 2003 Page(s): 78 – 83

[38] Python homepage, available at http://www.python.org

[39] Tcl documentation, available at http://www.tcl.tk

[40] H. Hua, C. Mineo, K. Schoenfliess, A. Sule, S. Melamed, R. Jenkal, and W. R. Davis, "Exploring Compromises among Timing, Power and Temperature in Three-Dimensional Integrated Circuits", *IEEE Design Automation Conference*, Jul. 2006.

[41] S. J. Souri, K. Banerjee, A. Mehrotra, and K. C. Saraswat, "Multiple Si layer ICs: motivation, performance analysis, and design implications", *IEEE Design Automation Conference*, Jun. 2000. Page(s): 213 - 220

[42] W. R. Davis, J. Wilson, et al. "Demystifying 3D ICs: The pros and cons of going vertical", *IEEE Transaction on Design & Test of Computers*, Nov. 2005, Page(s): 498-510

[43] R. Ho, K. Mai, and M. Horowitz, "The Future of Wires", *Proceedings of the IEEE*, April, 2001, Page(s): 491-506

[44] K. Balakrishnan, V. Nanda, S. Easwar, and S. K. Lim, "Wire Congestion And Thermal Aware 3D Global Placement," *IEEE/ACM Asia South Pacific Design Automation Conference*, 2005, Page(s): 1131-1134

[45] W. R. Davis, H. Hua, et al. "Wire-delay Reduction Analysis of a 3-tier, 8-point Fast Fourier Transform 3D-IC", *International VLSI Multilevel Interconnection Conference*. Oct, 2005, Page(s): 474-479

[46] S. Das, "Design Automation and Analysis of Three-Dimensional Integrated Circuits", *PhD thesis*, May 2004, MASSACHUSETTS INSTITUTE OF TECHNOLOGY

[47] K. M. Schoenfliess, "Performance Analysis of System-on-Chip Application of Three-Dimensional Integrated Circuits", *Master's Thesis*, North Carolina State University, Dec. 2005.

[48] OpenRISC Reference Platform System-on-a-Chip and OpenRISC 1200 IP Core Specification, available online at http://www.opencores.org/projects.cgi/web/or1k/orpso

[49] C. Mineo, "Clock Tree Insertion and Verification for 3D Integrated Circuits", *Master's Thesis*, North Carolina State University, Raleigh, NC, September 2005, pp. 30 – 33.

[50] Ansoft Q3D Extractor product information, available from Ansoft at: http://www.synopsys.com/products/analysis/primetime_ds.html

[51] BSIMSOI model, available online at: http://www-device.eecs.berkeley.edu/~bsimsoi/

[52] A. Minichiello, and C. Belady, "Thermal design methodology for electronic systems", *Thermal and Thermomechanical Phenomena in Electronic Systems*, Jun. 2002, Page(s): 696 - 704.

[53] P. Wilkerson, A. Raman, and M. Turowski, "Fast, automated thermal simulation of three-dimensional integrated circuits", *Thermal and Thermomechanical Phenomena in Electronic Systems*, Jun. 2004, Page(s): 706 - 713.

[54] Celtic User Guider, *Cadence*. Available with Cadence Celtic or Encounter product.

[55] Design Compiler User Guider, *Synopsys*. Available with Synopsys design compiler product.

[56] First Encounter User Guider, *Cadence*. Available with Cadence Encounter product.

[57] A. Akturk, N. Goldsman, and G. Metze, "Coupled Simulation of Device Performance and Heating of Vertically Stacked Three-Dimensional Integrated Circuits", *International Conference on Simulation of Semiconductor Processes and Devices*, Sep. 2005, Page(s): 115-118.

[58] C. N. Chu and D. F. Wong, "A Matrix Synthesis Approach to Thermal Placement", *IEEE Proceeding of International Symposium on Physical Design*, 1997, Page(s): 163-168

[59] Material thermal prosperities database, available online at: http://www.efunda.com/materials/elements/element_list.cfm

[60] H. Hua et al "Performance Trend in Three-Dimensional Integrated Circuits". *IEEE International Interconnect Technology Conference*, Jun. 2006.

[61] Abstract User Guider, *Cadence*, available with Cadence Abstract product.

[62] B. M. Tenbroek, M. S. L. Lee, W. Redman-White, C. F. Edwards, R. J. T. Bunyan, and M. J. Uren, "Measurement and simulation of self-heating in SOI CMOS analogue circuits", *IEEE SOI Conference*, Oct 1997, Page(s): 156-157

[63] A. P. Chandrakasan, and R. W. Brodersen, "Minimizing Power Consumption in digital CMOS circuits", *Proceeding of the IEEE*, Apr. 1995, Page(s): 498-523.

[64] Primepower User Guider, *Synopsys*. Available with Synopsys Primepower product.

[65] Q. Wang,; S. B. K. Vrudhula, "On short circuit power estimation of CMOS inverters", IEEE ICCD, Oct. 1998, Page(s): 70-75

[66] K. Nose, and T. Sakurai, "Analysis and future trend of short-circuit power", *IEEE Transactions on Computer Aided Design of Integrated Circuits*, Sep. 2000, Page(s): 1023-1030

[67] Predictive Technology Model, available online at: http://www.eas.asu.edu/~ptm/

[68] P. Van Zant, "A Practical Guide to Semiconductor Processing", *New York, McGraw-Hill*, 1990.

[69] Ziptronix 3D integration technology, available online at http://www.ziptronix.com/

[70] Tezzaron 3D integration technology, available online at http://www.tezzaron.com/

[71] C. Ababei, Y. Feng, B. Goplen, H. Mogal, T. Zhang, K. Bazargan, and S. Sapatnekar, "Placement and Routing in 3D Integrated Circuits", *IEEE Transaction on Design & Test of Computers*, Nov. 2005, Page(s):520-531

[72] J. W. Joyner, P. Zarkesh-Ha, and J. D. Meindl, "Global Interconnect Design in a Three-Dimensional System-on-a-Chip", *IEEE Transaction on Very Large Scale Integration (VLSI) System*, Apr. 2004, Page(s): 367-372

[73] "Reconfigurable Computing", *University of Massachusetts Amherst*, spring, 2006. Available online at http://www.ecs.umass.edu/ece/tessier/courses/697ff/index.html.

[74] Free Floating-point arithmetic units, available online at: http://www.hmc.edu/chips/index.html

[75] *OpenRISC Reference Platform System-on-a-Chip and OpenRISC 1200 IP Core Specification*, available online at http://www.opencores.org/projects.cgi/web/or1k/orpsoc

[76] P. Benkart, A. Heittmann, H. Huebner, and U. Ramacher, “3D Chip Stack Technology using Through-Chip Interconnects”, *IEEE Transaction on Design & Test of Computers*, Nov. 2005, Page(s): 512-517.

[77] N. Walker, “Integrated Circuit Module Generation”, *Master’s Thesis*, 1999, University of California, Berkeley, CA.

Appendix

Appendix A.

An Introduction to NCSU 3DIC Design Environment

A physical design environment known as NCSU 3DIC PDK (North Carolina State University 3DIC Circuits Physical Design Kit) was developed for cadence users to design their 3DIC in cadence virtuoso. This PDK provides a user friendly environment and eliminates the manual label and alignment of inter-tier vias that is required by another published methodology [26]. This PDK provides verification features including DRC, LVS, and parasitic extraction. The technology file and DRC/LVS rules support multiple tiers simultaneously, allowing designers to implement a complete 3D design within one layout database. Designers can design a 3DIC in the same way they design a conventional IC.

The NCSU 3DIC PDK currently supports MIT Lincoln Lab 0.18 μm Three-tier Three-metal-layer Low-power Fully Depleted SOI technology. In this PDK, three tiers are name Tier A (Tier 1), Tier B (Tier 2), and Tier C (Tier 3) respectively. Figure A.1 is an example showing an NAND2 gate implemented in NCSU PDK. The NMOS is implemented on tier A and PMOS is implemented on tier B. They are connected by the inter-tier via (three large squares). Tier A and B are electrically connected by the inter-tier vias and form a NAND2 gate in 3D integration. The inter-tier via is large compare to the NAND2 gate itself and it is also placement and routing blockage. Besides the above mentioned tiers, this PDK also supports a generic tier to support

an automated design flow. A skill script is created to load tier specific layouts and then merge them in this PDK for DRC/LVS check. The layer names associated with a tier name defines which tier a specific layer belongs to. For example, M1 (metal 1) on generic tier, tier A, tier B and tier C are named M1, M1_A, M1_B and M1_C respectively. Cadence virtuoso supports up to 255 layers. Currently, this 3D PDK supports a total layer number of 240.

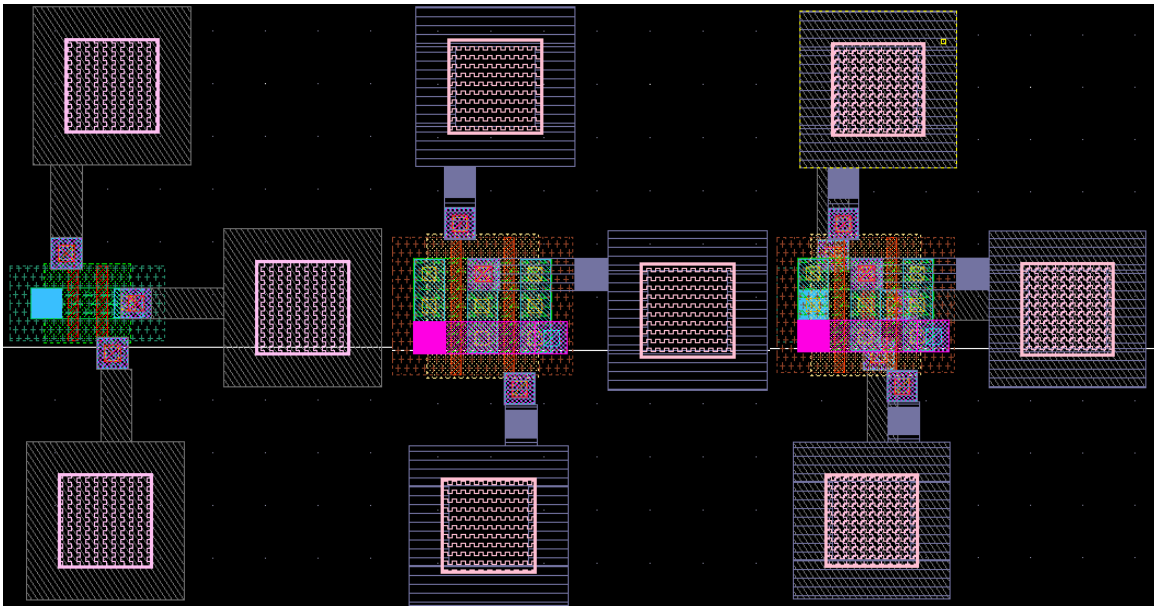

(a) NMOS (tier A) (b) PMOS (tier B) (c) CMOS (tier A and B)

Figure A.1: 3D PDK example. [12]

[12] Thank Dr. Rhett Davis for creating the layout and let me use it.

Appendix B.

Standard Cell Library Development

Table B.1: Standard cell functions.

Cell	Function	Cell	Function
AND2X(1,2,4,LP)	Y=(A&B)	AOI21X(1,LP)	Y=!((A&B)|C)
AOI22X(1,LP)	Y=!((A&B)|(C&D))	BUFX(1,2,4,LP)	Y=A
CLKBUF(1,2,4)	Y=A	CLKBUFX(1,2,4,8,LP)	Y=A
CLKBUFX(1,2,4,8,LP)	Y=A	DFFNEGX(1,2,LP)	Data=D, Clock=!CLK, Q=DS0000
DFFPOSX(1,2,LP)	Data=D, Clock=CLK, Q=DS0000	DFFSR(1,X1,X2,LP)	Data=D, Clock=CLK, preset=!S, clear=!R, Q=P0001
DFFPOSSCAN	Data=(TE TI) + (!TE D), Clock=CLK, Q= DS0000,	DFFSRSCAN	DFF with set and reset
FAX(1,LP)	YC=(A&B)|(B&C)|(C&A) YS=(A^B^C)	HAX(1,LP)	YC=(A&B) YS=(A^B)
INVX(1,2,4,8,LP)	Y=!A	LATCH(1,LP)	Data=D, Clock=CLK, Q=DS0000
MUX2X(1,LP)	Y=!(S?(A:B))	NAND2X(1,LP)	Y=!(A&B)
NAND3X(1,LP)	Y=!(A&B&C)	NOR2X(1,LP)	Y=!(A|B)
NOR3X(1,LP)	Y=!(A|B|C)	OAI21X(1,LP)	Y=!((A|B)&C)
OAI22X(1,LP)	Y=!((A|B)&(C|D))	OR2X(1,2,4,LP)	Y=(A|B)
TBUFX(1,2,LP)	Y=(EN?!A:'BZ)	XNOR2X(1,LP)	Y=!(A^B)
XOR2X(1,LP)	Y=(A^B)	Inter-tier Via Cells	Via cells that provide tier to tier connection
minimum routing pitch = 1*um*		manufacture grid = 0.025*um*	
maximum routing layer = 3		*Wp/Wn* ratio = 2	

A standard cell library was created for the implementation of our automated 3DIC design flow. This library is base on IIT standard cell library [14] but has more cells with different driving strength. The functions of each cell are listed in Table B.1. Inter-tier vias are designed as standard cells for automated placement and routing.

It is desirable to have the minimum delay for the standard cells. Experiments on an inverter an FO4 load for various *Wp/Wn* ratios show that when *Wp/Wn*=2, the inverter has the smallest delay (Figure B.1). Another important criterion to evaluate a cell is energy delay product (EDP). Simulations show that when *Wn*=0.5 μm, EDP has the smallest value with *Wp/Wn* equals 2. Based on this result, all cells' *Wp/Wn* ratio are set to 2 and all low power cells (end with LP) have *Wn*=0.5 μm. The model used for this characterization is BSIMSOI 3.2 [51], which corresponds to Hspice level 57. Timing and power values for this cell library is characterized by Signalstorm [1]. Power supply for this technology is 1.5V. Clock buffers are designed specifically for a fast clock tree synthesis. The input capacitance of each clock buffer is identical so that change one clock buffer in a sub-routine does not affect any other branches. Clock buffers with certain driving strength have different delay stages, which is helpful to reduce skew during clock tree synthesis by replacing a buffer with same driving strength but different delay value.

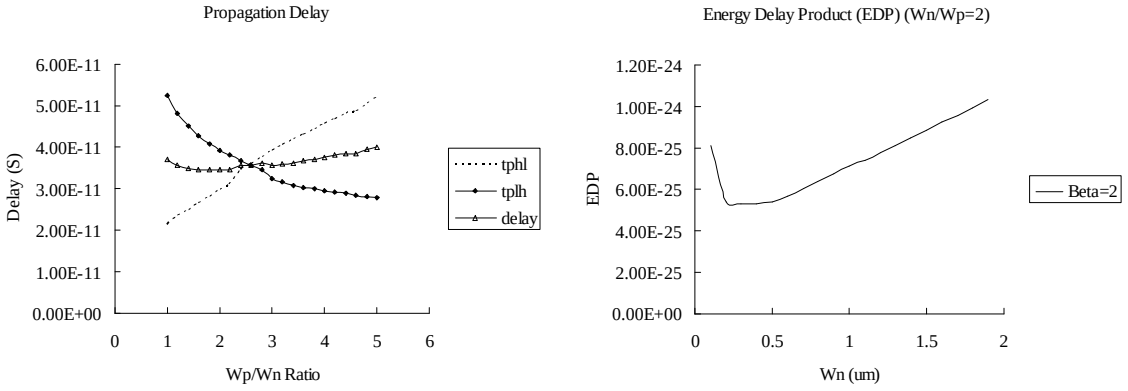

Figure B.1: Propagation delay and energy delay product for an inverter.

The Library Exchange Format (LEF) file is created by cadence tool "abstract" (version 5.5) [61]. Antenna damage is reduced in SOI because the buried oxide greatly suppresses the

current through gate oxide during etch [19]. Therefore, we did not implement antenna rule check in our LEF file.

Appendix C.

Electro-Thermal Delay-Power Analysis

C.1 Pseudo Code of Electro-Thermal Coupling Analysis

Data Preparation Inputs:
DEF Files: These files are detail routed tier specific design.
Timing Report: 10% most timing critical path delay (this percentage is subject to change).
Power Report: *PrimePower* or *Design Compiler* cell based power report.

Data Preparation Outputs:
Density Report: This file contains **metal** density and **power** density.
Timing Report: This file annotates the **location** of each cell in the timing critical paths.

Pseudo Code for Thermal-Performance Coupling:

```
Read in the "Density Report" and "Timing Report"
Create compact resistive network as shown in figure 4.2 (a) according to "Density Report"
Create dynamic power matrix according to "Density Report"
Create leakage power matrix according to "Density Report"
Create temperature matrix and initialize all elements in temperature matrix to be 0.
temperacture_change = 100%
While temperacture_change > threshold {
   annotate the dynamic and leakage power matrices to the thermal system
   simulate with Hspice
   update temperature matrix
   update delay of each cell in timing critical paths with temperature matrix (equation 18)
   calculate new critical path delay
   update dynamic power matrix with new critical path delay (equation 20)
   update leakage power matrix with temperature matrix
   temperature_change = max [temperature change of same matrix element on two consecutive runs]
}
Report new critical path delay
Report new dynamic and leakage power
Report final temperature
```

Figure C.1: Electro-thermal Delay-Power Coupling Analysis Pseudo Code

The pseudo code of electro-thermal coupling analysis is shown in Figure C.1. To begin, two files must be ready. The first file is a density report, which includes block ID, metal densities in X, Y, and Z direction of that block, and power density (both dynamic and leakage). The block ID is in the form of "tier x y", where tier tells which tier is the current block on and x, y gives the horizontal coordinates of that block. If a tier is divided into *m* by *m* blocks, x and y can be any integer number between 1 and *m*. The metal densities in X, Y directions are the report for the lateral metal layers Metal1 to Metal5, while the metal density in Z direction is the inter-tier via density that includes both thermal and signal vias. The second file is the timing report, which is in the form of "cell-name delay edge x y". The cell-name gives the name of a cell in the timing critical path. In our method, the cell-name also contains the tier name. Delay gives the delay value of the cell that is simulated with nominal temperature. Edge indicates if it is a rising or falling edge. x and y tell the block coordinates where the current cell locates. These two files are obtained by extracting information from the tier-specific DEF files, *PrimeTime* or *Design Compiler* timing report, and *PrimePower* or *Design Compiler* cell based power report. A Python script named "extract_power.py" is used to extract this information. For any specific design with a fixed tier count and resolution (*m* value), this extraction only needs to be done once.

A Python script named "thermal_closure.py" performs electro-thermal coupling analysis based on the "density report" and the "timing report". A three-dimensional array is created to represent dynamic power matrix, leakage power matrix, and temperature matrix respectively. A two-dimensional list is created to store the timing critical paths. A compact resistive network is created for thermal simulation. The dynamic power and leakage power of the elements with

same indices in the dynamic and leakage power metrics are added and annotated to the constant current source (shown in Figure 4.2 a) at the corresponding locations. Then Hspice is used to simulate the temperature profile of the compact resistive network. Timing is updated using equation (18) and power matrix is updated using equation (20) with the simulated temperature matrix. The coupling analysis loops until the maximum temperature change of same block is below threshold on two consecutive runs. In this study, the thermal conductivity temperature dependence is ignored for simplicity.

If we limit the maximum temperature change to be 0.1%, the runtime for different simulation resolutions in the three-tier five-metal-layer technology is shown in Table C.1. Column 2 denotes how many iterations are executed before temperature change converges. Column 3 gives the final maximum temperature change of any block on two consecutive simulations. The runtime increases exponentially with resolution. When resolution exceeds 50, the runtime becomes too long to be applicable. A recursive thermal simulation need to be employed to reduce runtime.

Table C.1: Runtime of electro-thermal coupling analysis vs simulation resultion.

Resolution	iterations	Final max temperature change (%)	Runtime	
			CPU	Real
1×1	2	0%	< 1*s*	1*s*
10×10	3	0.03%	11*s*	12*s*
50×50	3	0.05%	1235*s*	1236*s*

C.2 Recursive Thermal Simulation

The mechanism of recursive thermal simulation is shown in Figure C.2. Initially, the entire system is divided into a small number of blocks for simulation. Thereafter, each simulated sub-block is again divided in to a same number of blocks for simulation, using the adjacent blocks to determine the boundary conditions as shown in Figure C.2 (a). However, in the global view, the system is being simulated with a finer and finer resolution as shown in Figure C.2 (b).

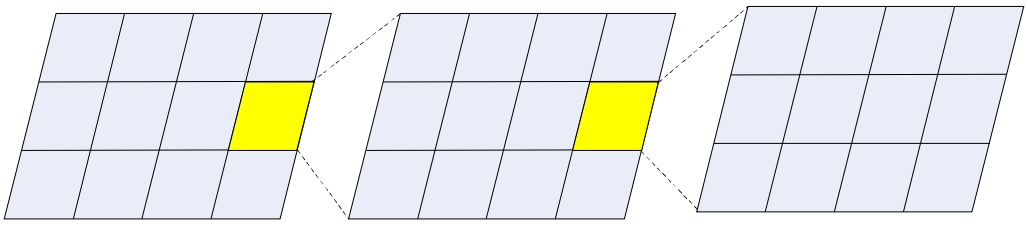

(a) Local simulation

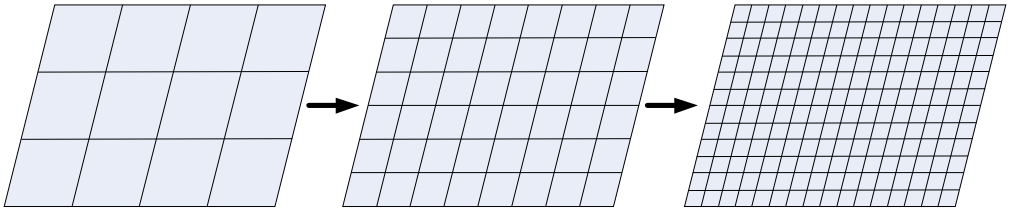

(b) Global simulation

Figure C.2: Recursive thermal simulation.

Figure C.3 shows the pseudo code for recursive thermal simulation. The recursive thermal simulation constitutes the simulation part the refinement part. In the thermal simulation phase, an initial temperature matrix was obtained for the target system. In the refinement phase, more accurate boundary conditions are applied to each sub region. The refinement phase will stop when the two successive temperature simulations converge.

```
Proc: thermal_sim (BCs, thermal-network, level) {
   Simulate temperature with current BCs and thermal-network
   if (current simulation level > 1)
      Foreach block in current system
        Construct new-thermal-network for current block (Figure C.2 (a))
        Using the temperature and thermal conductance (or resistance) of adjacent blocks as new
            boundary conditions (new-BCs)
        thermal_sim (new-BCs, new-thermal-network, level-1)
   else
       return (temperature)
}
Proc: Refine (thermal-network, bin-size) {
   update temperature with bin-size
   if (error>threshold)
     update temperature
     refine(thermal-network, bin-size)
   else
     return (temperature)
}
Main:
temperature = thermal_sim (BCs, thermal-network, level)   #initial temperature
temperature = Refine (thermal-network, bin-size)          # final temperature
```

Figure C.3: Pseudo code for recursive thermal simulation.

The major trade-offs of the recursive simulation are runtime and accuracy. A finer top level simulation usually results in a better accuracy at the price of longer runtime. However, the runtime increases exponentially with the simulation bin size, to obtain reasonable result in a relatively short time, some accuracy have to be sacrificed. If we consider the whole thermal network simulation (equivalent to level 1 simulation) is accurate, the accuracy vs runtime is shown in Table C.2 for a randomly generated three-tier five-metal-layer thermal network and heat generation. The resolution column describes how many grids the system is divided into.

The bin size column gives how many sub-blocks current block is divided into (as shown in Figure C.2 (a)). The simulation level tells how many recursive simulations will be executed. For example, if the level is 1, the entire system is simulated as a whole; if the level is 2, then the simulation illustrated in Figure C.2 (a) will stop at the second figure; and so on so forth. The runtime give the real time (including real time and CPU time only) of how long the simulation takes. The last two columns give the errors. As can be seen from the table, as bin size decreases, the CPU time reduces significantly while the mean error only increases a little. However, real system time increases due to the large amount of time spend on checking out the tool license. A simulation method to avoid this license checking time is necessary to expedite the simulation process.

Table C.2: Runtime vs accuracy in recursive thermal simulation.

Resolution	Bin size	Simulation level	Runtime		Mean error	Error range
			CPU	Real		
64×64	64×64	1	1335*s*	1336*s*	0%	--
64×64	8×8	2	143*s*	212*s*	2.55E-7	[-7.2E-5%~7.5E-5]
64×64	4×4	3	85*s*	413*s*	9.75E-7	[-8.1E-5%~1.0E-4]
81×81	81×81	1	5582*s*	5584*s*	0%	--
81×81	9×9	2	351*s*	449*s*	0.004%	[-1.1%~0.94%]
81×81	3×3	4	181*s*	1247*s*	0.005%	[-1.38%~1.18%]
100×100	100×100	1	Fail to finish		N/A	N/A
100×100	10×10	2	425*s*	526*s*	N/A	N/A