Design of Energy Efficient Future CMPs with On-Chip Wireless Interconnects

Student Name: Gade Narayana Sri Harsha

IIIT-D-MTech-EC-12-094
June 23, 2014

Indraprastha Institute of Information Technology
New Delhi

Thesis Committee
Prof. Ashwin Srinivasan (Chair)
Dr. Turbo Majumder
Dr. Sujay Deb
Dr. Shobha Sundar Ram
Dr. Subhasis Banerjee

Submitted in partial fulfillment of the requirements
for the Degree of M.Tech. in Electronics & Communication,
with specialization in VLSI & Embedded Systems

Keywords: Low Power, DVFS, Power Gating, Wireless, Channel Modeling, FDTD

Certificate

This is to certify that the thesis titled **"Design of Energy Efficient Future CMPs with On-Chip Wireless Interconnects"** submitted by **Gade Narayana Sri Harsha** for the partial fulfillment of the requirements for the degree of *Master of Technology* in *Electronics & Communication Engineering* is a record of the bonafide work carried out by her / him under my / our guidance and supervision in the Security and Privacy group at Indraprastha Institute of Information Technology, Delhi. This work has not been submitted anywhere else for the reward of any other degree.

Dr. Sujay Deb
Dr. Shobha Sundar Ram
Indraprastha Institute of Information Technology, New Delhi

Abstract

Power density and interconnect delay have emerged as the biggest challenges for Ultra Large Scale Integration (ULSI) and System-on-Chip (SoC) designs particularly beyond $65nm$ generation. Leakage power has been traditionally non-critical in CMOS circuits, but with extensive scaling, even it is increasing significantly. Dynamic Voltage/Frequency Scaling (DVFS) has been demonstrated to be one of the effective ways to reduce the power consumption and power density across the chip. Network-on-Chip (NoC) architectures improve the performance over the traditional bus based architectures. But with increasing chip sizes and long interconnects, the delay due to wired interconnects extend to multiple hops. Long range wireless links in NoC have been proven to improve the latency and energy performance tremendously.

In this work, we have designed and implemented a centralized controller that applies DVFS to the processing cores. DVFS techniques reduce power consumption by scaling down voltage and frequency when possible with a little impact on performance. The proposed controller observes current state and utilization of the core and based on past state transitions, predicts the next state to set the voltage and frequency. To further reduce the power consumption, the controller also applies power gating method to the wireless interfaces used in the system. All wireless interfaces that are not in any active communication are put in idle state. The biggest advantage of centralized controller is the less overhead it adds to the system. But the delay associated with control signal transmission, particularly to remote corners of the chip is very high and so affects the performance of controller. To reduce this delay, we propose the use of wireless interface for the same and a dual band transceiver is used for this purpose.

The use of wireless interfaces definitely reduces the delay significantly, but the delay values used assume ideal operating conditions. Previous works have shown that the wave propagation on chip deviates largely from ideal scenario and multiple propagation paths and wave components exist. The delay in strongest component is much more than the delay of free space direct wave. Hence the second contribution of the work is analyzing and modeling intra-chip wave propagation mechanisms. A 2D model for on-chip components is developed and using FDTD simulations, different propagation paths and modes are identified. It is observed that the free space direct wave is canceled out and reflections from interconnect layers are the dominant component of the signal. The delay in this component is almost twice the free space delay and is dependent on materials used. Finally it is shown that even with increased delay, wireless interfaces still can outperform the wired interconnects and the delay is within single cycle limits.

Acknowledgments

The work presented in this thesis would not have been possible without the help and support of several individuals who assisted me in many ways.

First and foremost I would like to express my sincere gratitude to my advisers Dr. Sujay Deb and Dr. Shobha Sundar Ram for their expert guidance and supervision. They inspired me to do good work and supported me all along the way.

I am grateful to my thesis committee, Prof. Ashwin Srinivasan, Dr. Turbo Majumder and Dr. Subhasis Banerjee for their valuable feedback. I also would like to thank Dr. Pankaj Jalote for enabling such a wonderful work environment at IIIT Delhi.

I owe special thanks to my family for their endearing support and encouragement in my pursuits. Without them none of this would not have been possible. Lastly, I thank my friends and PhD students for providing needful assistance and fruitful discussions, all of which helped me greatly.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Background and Motivation

CMOS devices have been continuously scaled to achieve high performance, packaging densities and low power consumption. The number of transistors on a single chip has doubled for every 18 months in accordance with Moore's Law, allowing multiple cores to be embedded on a single chip and enabling advancements in Chip Multiprocessors (CMPs). This has improved the performance tremendously and reduced the cost of devices considerably.

As per the ideal scaling model, the power density in a chip remains constant with scaling. But the observed trends show that the power density indeed increases with scaling and sub threshold leakage power equals the dynamic power beyond 20nm technology generation as shown in Figure 1.1 [37]. The increased power density leads to higher temperatures in the chip. Higher temperatures degrade the functionality of the device. It reduces the reliability of the transistor and the life time of the device. Hence in Deep Submicron (DSM) and Ultra Submicron (USM) technologies, power becomes the major limiter of system performance.

Dynamic Voltage/Frequency Scaling (DVFS) methods are proposed to tackle large power densities and to reduce the energy consumption of the chip. DVFS techniques exploit the process insensitive idle phases of an application/task to reduce the supply voltage V_{DD} and frequency to achieve large reductions in the power with little performance loss. In case of tasks with heavy core usage, DVFS can boost the voltage and frequency to execute the task faster if needed. In this way, DVFS techniques can reduce energy consumption, boost performance and balance workloads according to the specific requirements of the application.

Many DVFS algorithms have been proposed in the literature but they tend to be either software based implementations or the actual hardware overheads and their performance have

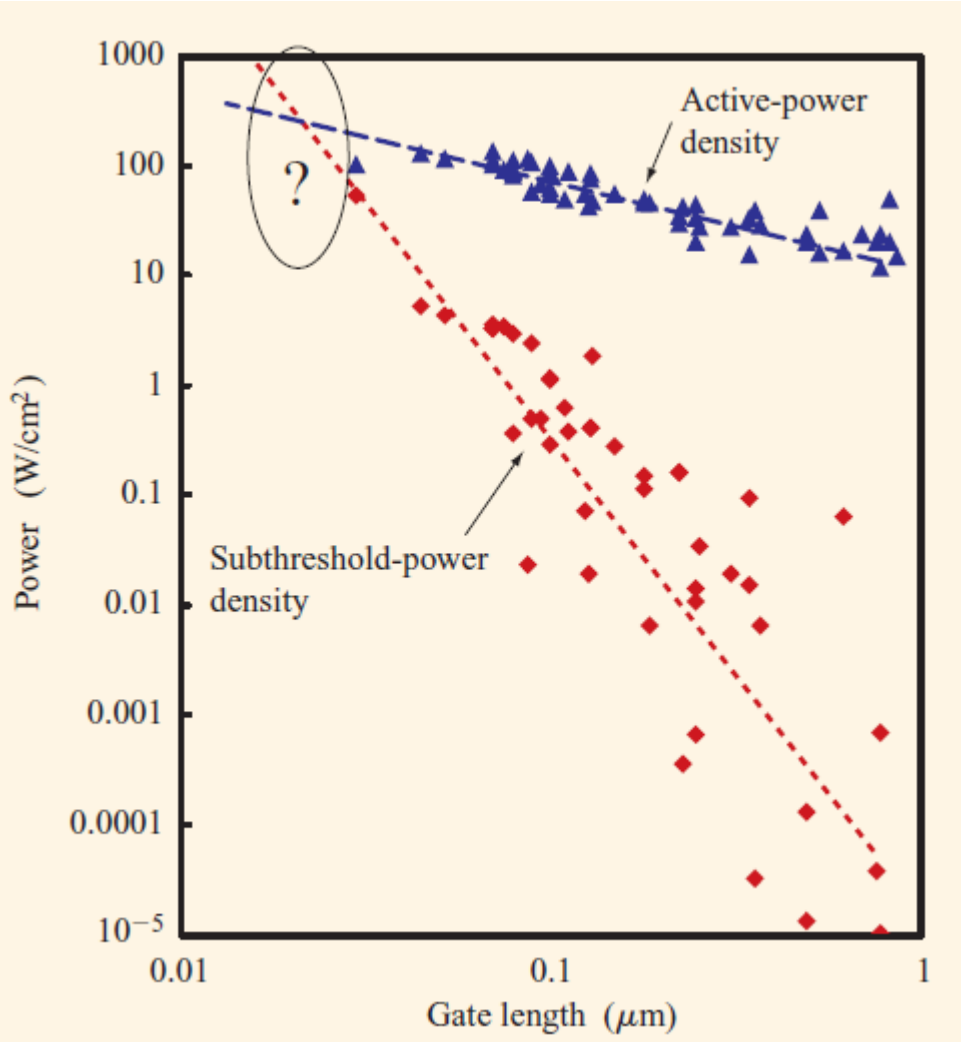

Figure 1.1: Power Density vs Scaling

not been extensively studied. The DVFS hardware implementations can be either per core, clustered or centralized approaches. Centralized and clustered methods, even though can be complex, add less overhead to the total system. In a centralized system, a chip level controller observes all required parameters and controls different parts of the system.

But, with scaling and increasing chip sizes, the interconnect delay increases as opposed to the gate delay. The variation of delay with each technology generation is shown in Figure 1.2 [10]. With each scaling generation, the gate delay has decreased by 30%, whereas the interconnect delay increases by 40%. These long delays become bottleneck for the communication between controller and various modules of the chip. Large interconnect delays also increases the complexity of communication architecture. So, reducing this delay can improve the overall performance.

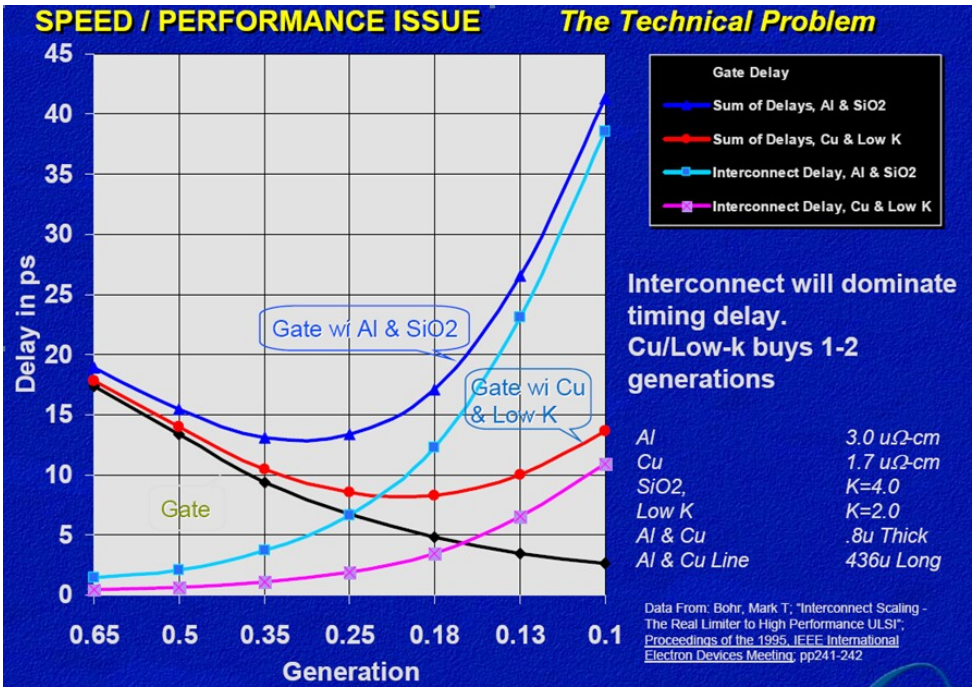

Figure 1.2: Delay vs Scaling

Network-on-Chip (NoC) with planar metal interconnects offer significant improvements in on chip communication architecture and performance compared to the traditional bus based architectures and are extensively used in current day systems. But with large chip areas, the latency and power consumption in NoC is increasing due to the multi hop delays of long metal interconnects. One possible solution to decrease the overall delay in NoC is to use emerging interconnects like G-Lines in place of metal interconnects. But they consume significantly large power compared to metal interconnects and are not ideal for most applications. Single hop on-chip wireless links with high bandwidth provide a feasible solution in terms of both delay and power.

The use of wireless links offer great improvements in delay and energy for NoC architectures. But, the delay values used for evaluation purposes generally assume ideal free space transmission. But previous works have shown that the delay in significant component of the signal is much larger that that with free space case. Also the on-chip wireless interface is still an emerging technology and faces significant challenges in design and implementation. There is extensive research in development of integrated antennas, but the intra-chip wave propagation mechanisms are not yet fully explored.

The wave propagation between two on-chip antennas is affected by different interference structures present on chip like metallic interconnects and substrates. They add huge antenna losses, reduce the transmission gain and affect the delay of the signal. Multiple propagation paths and modes exist for intra-chip wireless communications. And apart from this, the antenna and

underlying analog hardware themselves have their overheads. All these factors together reduce the reliability of the signal compared to the ideal case. Owing to these challenges, the characterization of the on chip wireless channel and analyzing its performance assumes importance to further the implementation of on chip wireless links.

1.2 Problem Statement

Summarizing the previous discussion, the work presented in this thesis can be described as a twofold problem:

1. To design an efficient centralized DVFS controller in hardware using on chip wireless interconnects for improving the performance and efficiency of multi-core systems.
2. To model the behavior of the on chip channel and to observe the effects of substrate & other chip components on the performance of wireless interconnects used in the system.

Hence, by implementing a centralized DVFS control method, we balance the performance and reduce the overall energy consumption of the chip with little overheads. The impact of the interconnect delay on the controller and system performance is avoided by the use of wireless interconnects. And we model the performance of wireless interconnect in a more realistic scenario to ensure that they indeed improve upon the disadvantages of other interconnects.

1.3 Outline of the Report

The outline of the thesis is as follows

Chapter 2 presents the literature study on existing implementation of DVFS methods and the work done in wave propagation mechanisms for on-chip wireless communications.

Chapter 3 discusses the proposed controller design and its implementation. Some of the issues faced with centralized design approach and possible solutions are also presented.

Chapter 4 presents the experimental setup and performance of the proposed DVFS implementation and analyzes complete thermal profile of the system under normal and DVFS operating conditions.

Chapter 5 details the on-chip model used to characterize the intra-chip wireless channel and presents a brief overview of FDTD simulation method used for this work.

Chapter 6 discusses the simulation setup, results and observations made on different components of the wave propagation in multilayered chip structure.

Chapter 7, finally summarizes the work done and concludes with the contributions of this thesis. The possible directions for extending the work in future is also briefly explored.

Chapter 2

Literature Study

2.1 DVFS

There is a significant work done in the domain of DVFS and many algorithms have been proposed to achieve power savings on chip. Power management policies and algorithms are integrated in current operating systems to make use of processor idle states and chip level power gating techniques implemented into the systems. OS C-state [3] policies for both Windows and Linux operating systems are one such example.

Task scheduling algorithms are developed to allocate tasks to individual cores to save power with little impact on performance. [47] presents a two phase framework that assigns and orders the tasks to maximize opportunities that can exploit lowering voltage levels. [45] developed a learning based dynamic power management framework for multi-core processors to judiciously allocate the tasks to achieve better trade-off between power and performance.

[17] proposes a machine learning prediction model and usage model to better predict and apply processor C-states compared to current reactive OS policies. They used Dynamic Bayesian Networks (DBNs) to predict CPU activity patterns for future time given all the observations upto the present time. The developed model improves power savings by 12% and performance by 2% over existing methods.

Most power management techniques operate by reducing performance capacity during idle/low activity phases, but in multi-core systems aggregate monitoring obscures the underlying phases on individual cores. To address these problems, a core level activity prediction method is proposed and discussed in [8] [9]. A Periodic Phase Power Predictor (PPPP) makes use of table-based prediction structures and repetitive nature of power phases to predict performance demand and appropriate DVFS selection is made. This method predicts core activity level rather react to the activity changes.

[36] [35] demonstrate a dual level DVFS for both processors and NoC to improve the power and thermal profiles without significant impact on execution time. Wireless Network-on-Chip (WiNoC), an emerging technology for low power and high bandwidth multi-core chips reduces hop count between distant communicating cores. This attracts a significant amount of overall traffic. A history based DVFS is implemented, where each router predicts the future link utilization based on previous short-term and long-term utilization.

Some other works include [23] analyzes the benefits of fine grained core level DVFS to multiple VFIs using real workloads for different application classes. They demonstrate that core level granularity offers little advantages. [31] proposed a clustered DVFS approach as intermediate solution between per-chip and per-core DVFS methods to find trade off between flexibility and incurred expenses. [20] presents a framework to compute theoretical bounds on DVFS performance under the impact of technological constraints like reliability, temperature, process variations and inductive noise. [38] models accurately the DVFS transition overheads for both energy consumption and delay.

2.2 Wireless Channel Modeling

The use of on-chip antennas as viable interconnect options is an emerging technology for Ultra Large Scale Integration (ULSI) or System-on-Chip (SoC) architectures. A large amount of research is focused on design and development of integrated antennas for wireless interconnects, but the propagation mechanisms and impact of integrated structures on chip are not yet fully explored.

The characteristics of integrated antennas on bulk, SOI and SOS substrates are studied and characterized in [30]. The measurements confirm signal transmission with reasonable gain at high frequencies ($> 15GHz$). The gain with SOS substrate is higher than that with bulk and SOI substrates. [22] and [42] investigate the impact of metal structures such as power grid and data lines on on-chip antenna performance. It has been observed that input impedance and phase of S_{12} is significantly changed and $|S_{12}|$ is reduced for a pair of antennas. Based on the evaluation, a set of guidelines have been developed to reduce the impact of the structures.

A plane wave model to understand propagation for intra-chip communications is proposed in [29]. They measured gain from $6GHz$ to $18GHz$ using two integrated dipole antenna and results show that the gain can be improved by inserting a dielectric layer between silicon wafer and ground plane. [28] shows that transmission gain can be improved using high resistivity Si substrates.

The propagation mechanisms of radio waves over intra-chip channels in the frequency range of $10GHz$ to $110GHz$ were studied in [48]. By measuring the S-parameters, they found that the

Path Loss Exponent (PLE) of the channel is significantly lower that the free space channel PLE. The time of first arrival of the signal is much later than that by free space transmission and surface waves dominate the intra-chip channel propagation. Inter-chip wireless communications are studied in [12].

Fundamental wave propagation mechanisms for intra-chip communications using multilayered chip structures are studied in [44]. Green's functions are used to represent different wave components for different antenna structures and it was found that surface waves may or may not provide the dominant propagation channel depending on frequency and chip structure. A guiding layer placed below or above the silicon substrate can enhance single channel transmission and thereby reduce dispersion. A comprehensive overview of on-chip antenna, their implementation benefits and challenges are presented in [11]. The authors conclude that current CMOS technologies and foundry specific rules are not best suited for on-chip antennas and innovative design techniques are required to overcome shortcomings and achieve true SoC RF solutions.

Chapter 3

Proposed Centralized Controller Design

In this chapter, the controller design and its hardware implementation is presented. The proposed design is a centralized controller which operates at the top level of the system. The controller applies the DVFS mechanism to the processing cores and power gating to the wireless interfaces used in the system. The DVFS algorithm makes use of core utilization information, system thermal state and user inputs to achieve runtime voltage and frequency scaling. When the wireless interfaces are not in use, power gating is applied to achieve energy savings. For this purpose, we have considered a multi-core system with a Wireless NoC (WiNoC) architecture.

A brief overview of power consumption on chip and DVFS is first presented in the Section 3.1. The DVFS and power gating algorithms are discussed in Section 3.2. Section 3.3 describes the hardware implementation of the controller and the issues associated with centralized controller and possible solutions are detailed in Section 3.4. The use of wireless interconnects to alleviate some of the issues with centralized controller implementation and their design is discussed in Section 3.4.1.

3.1 Overview of DVFS

Traditionally, operating frequency and gate delay have been the major limiters of the design performance in CMOS based systems. Through improvements in scaling and advancements in CMOS manufacturing technologies, the gate delays have been considerably reduced and higher performance is achieved. Developments in multi core designs have further improved the performance beyond the normal capacity of multiple single core systems. But, in recent years, particularly beyond the 65nm technology generation, power has emerged as the primary hindrance for the system performance. The power dissipation in the system limits the achievable performance due to the caps on cooling capacities or increase the total cost of the system. Signal

integrity also becomes a major concern due to IR drops, inductance effect, etc.

3.1.1 Power

The total power dissipation in a CMOS design is mainly composed of two components, Dynamic power ($P_{dynamic}$) and Static power (P_{static}). Dynamic power is the power from charging and discharging during the switching activity of a device and is given by $P_{dynamic} = \alpha * C * V_{DD}^2 * f$, where α is switching activity factor, C is the total switching capacitance, V_{DD} is the supply voltage and f is the operating frequency. Static power is due to the leakage in the device when it is idle. It can be due to sub threshold conduction, tunneling in gate oxide and leakage current through reverse biased diodes. It can be represented by $P_{static} = I_{Static} * V_{DD}$, where I_{Static} is total static current in the device. The total power consumed in a chip has been dominated by dynamic power, but there has been a steady increase in the leakage power with each generation.

One solution to reduce both the dynamic and static power is reduce the supply voltage V_{DD}. Keeping the supply voltage constant while scaling other parameters of the CMOS device increases the power density of the chip. The dynamic power varies quadratically with V_{DD} and the static power varies linearly with V_{DD}. Hence reducing the supply voltage provides a logical solution to gain significant reduction in power consumption. The dynamic power can be further reduced by scaling down the frequency. But if the execution time is assumed to be inversely proportional to frequency, the energy consumed remains same.

3.1.2 DVFS

Dynamic Voltage/Frequency Scaling (DVFS) is one of the most widely used techniques to reduce the supply voltage and frequency. DVFS methods reduce the voltage and frequency during the idle phases of an application to reduce the power consumption. Since the frequency is reduced, the application throughput decreases linearly with it. But as discussed earlier, the dynamic power is quadratic in voltage and linear in frequency. Therefore, DVFS achieves cubic reduction in dynamic power with a linear decrease in throughput. The static power is also reduced since it is linear in voltage.

Any increase in the voltage/frequency (V/F) first increases the supply voltage V_{DD} accompanied by the increase in frequency level. On the other hand, when the V/F level is to be reduced, first the frequency is first locked to lower value and then the V_{DD} is reduced. During both transitions, the processor operation is paused when the system is being locked to the new frequency to prevent inconsistencies in the data.

The DVFS algorithms can be broadly classified into two categories, Offline approaches and Online approaches. Offline approaches are basically software implementations where the DVFS

algorithms are built into the operating system. Online approaches are hardware implementations designed into the chip hardware that performs the DVFS operations. Methods like task scheduling, managing C-states, future workload predictions can be categorized under offline approaches. Online methods include Voltage/Frequency Islands (VFI), power gating techniques, etc. The centralized controller proposed in this work falls under online methods and monitors various system parameters to make decisions regarding VF levels.

3.2 Controller Algorithm

3.2.1 DVFS Algorithm

The proposed DVFS algorithm uses a time slice based approach, wherein each core utilization for a given time slice is observed and decisions are made to predict the state of the core for next time slice. The algorithm also takes into consideration the runtime temperature of the system and adjusts the voltage/frequency levels accordingly to prevent any damage to the system under heavy usage. The duration of the time slice can be chosen depending upon the general run time behavior of the system. Choosing a short slice duration may lead to increased operational overheads of the controller whereas choosing a very long duration may result in missing significant state changes in the system.

In our design, in any given time slice, a core can be operating in one of the N states, where each state represents a fixed operating voltage/frequency pair. At the end of each time slice, the algorithm observes the utilization of each core for that time slice and the temperature of the system. The core utilization is again divided into multiple levels. Using the past utilization levels and corresponding state changes of a core for each time slice, a probabilistic state change model assigns a probability to all possible changes between core states for each core. Using the core busy/idle periods, the average duration for which a core is busy performing any task is determined. Based on the state change probabilities for the current core state and mean busy duration period of the core, the core state and utilization for the next time slice is predicted. The corresponding operating voltage/frequency is then set for that core. At the next time slice, the actual utilization level is observed and the state change probabilities and mean duration values are updated accordingly at each slice to account for deviations from predicted values. This defines the normal operation of the algorithm.

But based on the observed temperature of the system, the algorithm may deviate from this operation. For each core state, we have defined three tolerable levels for the run time temperature of the system; acceptable, moderately high and very high levels. If the observed system temperature is within the acceptable levels for the particular core state, the predicted state is applied for the next slice. If the temperature is in moderately high level, then any transition to a higher powered state is prevented and core continues operating in same state for next slice.

But if the temperature is very high, then the core is operated at low powered state for next slice irrespective of the predicted state. This is done to prevent damage to system from cores continually operating at higher voltage/frequency levels. Finally, a manual user level input is also incorporated which forces the system to operate either in high performance state or a low power state. The normal algorithm operation is suspended as long as this input is active.

For implementation, we considered four core states, a low power state (S_{LP}), normal state (S_N) and two high performance states (S_{HP_1} and S_{HP_2}). The normal state corresponds to the system operating at rated voltage/frequency. The core utilization for a time slice is also divided into four levels as L_1 $(0 - 30\%)$, L_2 $(30 - 50\%)$, L_3 $(50 - 80\%)$ and L_4 $(> 80\%)$. The state machine diagram representing the possible state changes and corresponding utilization levels is shown in the Figure 3.1. On powering up, all the cores start in normal state, S_N. Now, if a core is currently operating in S_N state, the predicted utilization is L_1 and the probability for core continue its normal operation is high, then the state S_{LP} is assigned to the core for next time slice. Similarly the general possible state transitions and utilization cases are shown in the

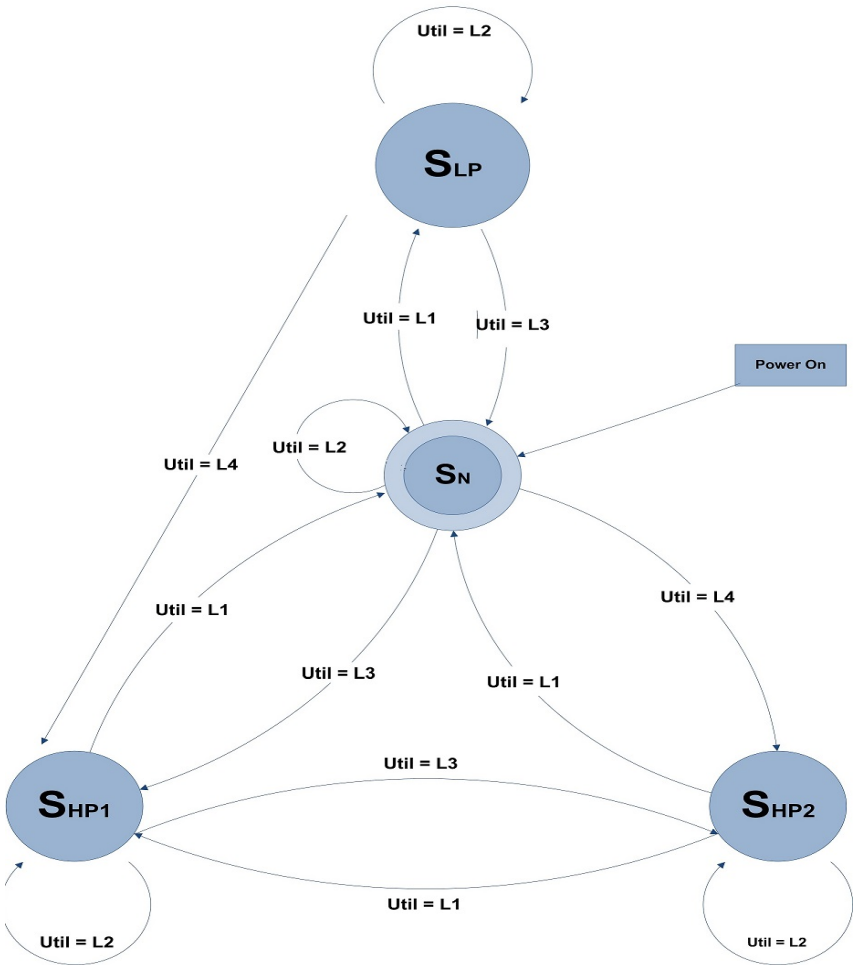

Figure 3.1: State Machine Diagram for Core DVFS Algorithm

figure. We have prevented direct transitions from high performance states to low power state.

3.2.2 WI Power Gating Algorithm

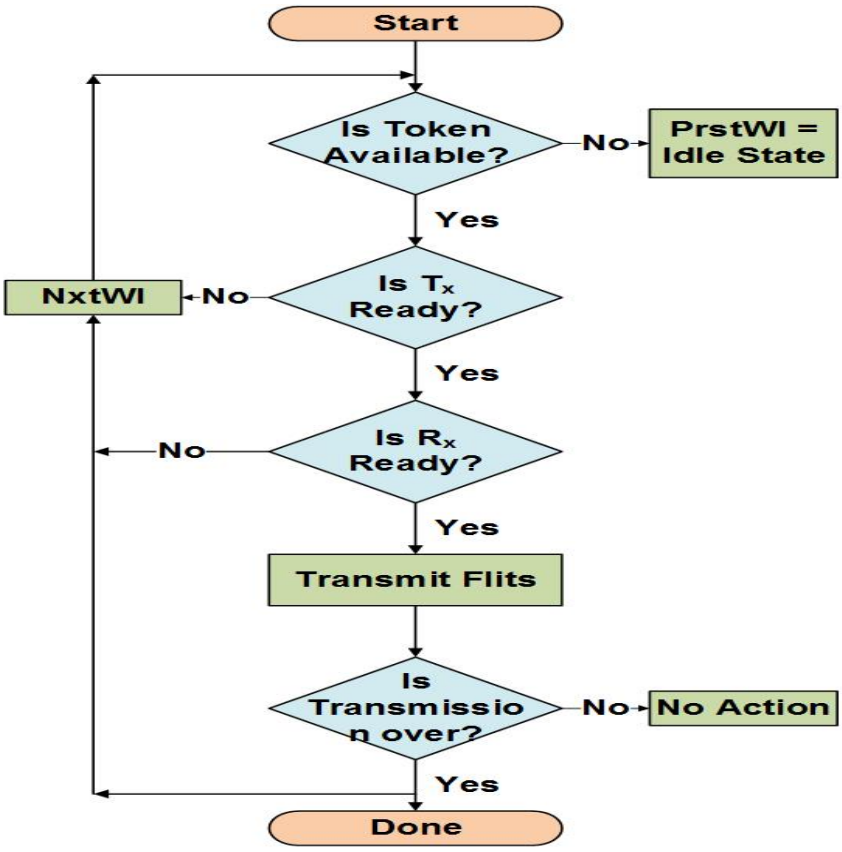

Figure 3.2: Flow Diagram for WI Power Gating Algorithm

The system architecture we have considered uses a hierarchical network architecture with wireless links as long rage shortcuts. The wireless links are strategically placed to reduce latency in long distance communication across the chip. It has been shown in [16] that wireless links provide energy savings along with latency improvement over conventional wired links. We further reduce the energy in wireless interfaces by incorporating power gating technique.

The flow diagram for power gating algorithm for WIs is shown in the Figure 3.2. The system has N WIs, named WI_1 to WI_N. At the system power up, all WIs are kept in sleep state initially. The WIs are then operated in a round robin fashion. We use a token management system, a token is passed through each WI from WI_1 to WI_N to check for availability of data to be transmitted at any WI. If there is no data to be transmitted at present WI (PrstWI), the the token is passed to next WI (NxtWI) and PrstWI is put in idle state. If data is available for transmission, the transmitter WI sends the receiver address to the controller. If the receiver is

ready to receive the data, the corresponding transmitter and receiver WIs are turned on and the data transmission is initiated. Once all the data flits are transmitted, both the WIs are again put into sleep state and the token is then passed to the NxtWI. During data transmission, the token remain the transmitting WI. Hence at any instant of time, only two WIs are active at most. This method achieves energy savings by turning off WIs that are not in any active data transmission.

3.3 Hardware Implementation

The centralized controller with core control and WI control modules is shown in the Figure 3.3. Core Control module implements the DVFS algorithm for the cores and WI Control unit applies power gating to the WIs and controls their sleep mode operation according to availability of data.

3.3.1 DVFS Controller

The core control unit implementation has four major modules; the Current State, State Change Model, Busy/Idle Pattern and the Temperature Control modules. The core utilization level is calculated as number of cycles for which the core is busy in any given time slice. The Current State module counts the number of busy clock cycles to calculate the utilization of the core in that slice and represents it as a two bit value pertaining to one of the four utilization levels mentioned in the previous section. The Busy/Idle pattern module also reads in the core busy/idle state at each clock cycle and updates the average duration in number of clock cycles for which the core remains busy. The modules for calculation utilization and average duration are all implemented as simple counters. The counters in current state block are all reset at each time slice but busy/idle pattern counters operate independent of time slices.

The State Change Model module reads in the current utilization level and state of the core from Current State module and using the state change probabilities estimates the state and utilization of the core for next time slice. Initially we assume that the probability for a core to continue in the same state to be one. Then as the core operation continues, the core utilization levels are considered and the transition probabilities from one state to other states for different utilization levels are updated at each slice. To prevent the use of floating point numbers for representing transition probabilities, they are represented as number of transitions from one state to another state for every 100 times the core is in a particular state. Using the data from this block and busy/idle block, the state of a core for next time slice is predicted. The Temperature Control block reads the temperature of the system at each time slice and compares with the predefined tolerable limits for the current state. Depending temperature level, a control signal is generated to assign the predict state or current state or the low power state to the core for next slice.

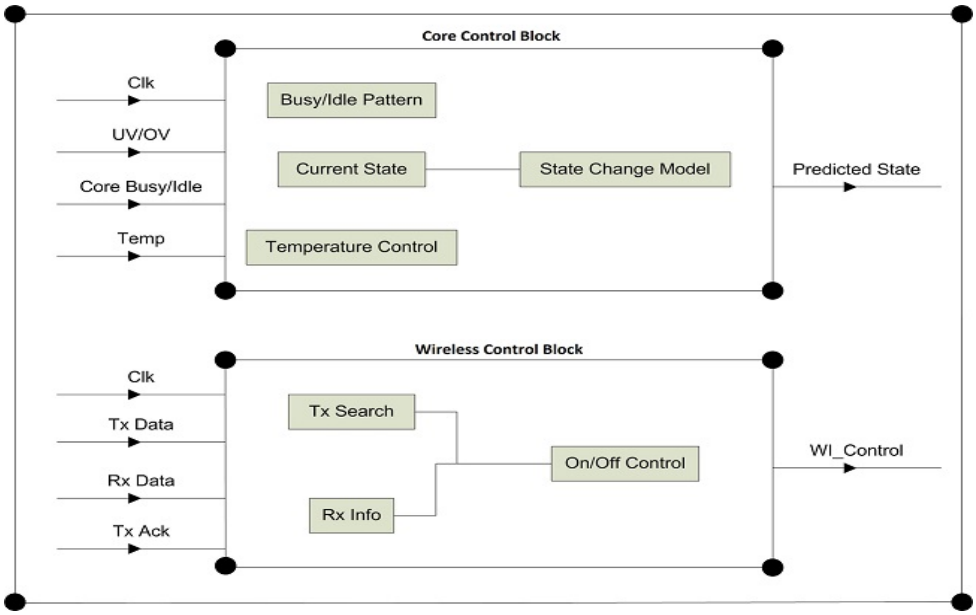

Figure 3.3: DVFS Controller with Core Control and WI Control Modules

3.3.2 WI Controller

For WI control module, the controller sends a token signal to one of WIs in the system at each clock cycle. If there is no response from the WI, the controller moves onto next WI. If it receives an acknowledgement signal from the WI, the receiver info available with the transmitter WI (T_X) is read in. The controller then sends a handshake signal to the corresponding receiver WI (R_X) and if it receives an acknowledgement signal, the controller disables the transmitter search operation. Then it sends control signals to the power gating modules at T_X and R_X to turn on the WIs. Once the data transmission is completed, the T_X and R_X send done signals to the controller and the controller resumes the T_X search operation.

3.3.3 Controller Size

Since the core utilization is measured in terms of clock cycles, the status of each core needs to be transmitted to the controller at each clock cycle. This adds a significant traffic overhead to the system. To prevent this, the counters for calculating the utilization are placed with the corresponding core and the utilization represented as one of the four levels is transmitted to the controller at the end of each time slice. This reduces both the communication overhead and the amount of data to be transmitted. Another aspect to be taken care of is scaling of controller hardware with system size. In case of WI controller, the hardware does not change significantly with the number of WIs. But the DVFS controller hardware scales up as the number of cores in the system increases. For systems with large number of cores, the system is divided into multiple clusters each containing a group of neighboring cores. Then we propose to use time multiplexing of these clusters to keep the controller hardware in check. Instead of updating all

the cores at the same time, cores from different clusters are updated at different times. The slice duration for each cluster remains the same, but each cluster is updated serially one after the other. The time duration between each update is chosen to accommodate data transmission delay from and to the controller and internal controller delay. This way the hardware required for the controller is equivalent to the hardware required for maximum number of cores present in any cluster.

3.4 Issues with Centralized Controller

One of the major issues with the centralized controller implementation is the delay in communication between the controller and various cores/clusters in the system. As discussed in the previous chapters, the delay spans multiple hops even with NoC architectures using wired interconnects. The delay increases as the length of the interconnect increases. Hence, the farther away, a core is from the controller, the longer the length of interconnect and so is the delay in data transmission. Reducing this delay is needed for efficient implementation of a centralized controller.

One of the solutions to reduce the interconnect delay is to use emerging interconnects like G-lines [32] for control signal transmission in conjunction with conventional wired interconnects. Global interconnect lines (G-lines) are multi-drop, broadcast capable and ultra low latency communication lines that reduce power and improve latency. A capacitive feed-forward circuit with voltage mode signaling method for global lines is proposed in [24] [25] that achieves single cycle delay for long RC wires. G-lines improve the latency performance of the NoC considerable but consume a lot more energy than metal interconnects and hence are generally not a viable solution for many applications. Another alternative to reduce the delay is to use the wireless interfaces for transmitting the control signals. Inserting long range wireless links in NoCs have been shown to provide improvements in both delay and energy savings [16] [19].

3.4.1 Dual Band Wireless Interfaces

The system architecture considered already uses wireless interfaces to transmit data signals between various components on chip. We extend the use of these WIs for transmitting the control signals between the controller and different cores/clusters in the system. To achieve this goal, we use two different frequencies for data and control signal transmission.

In this section, we present a brief description of the WIs used with controller and associated with the cores/clusters. The WI at the controller is only used for control signal transmission and reception. So, a single band transceiver is used with a *zigzag* antenna for this purpose. The *zigzag* antenna has near omni-directional radiation pattern and is ideal for controller WI since it needs to communicate with different parts of the chip. It is designed with $60\mu m$ length, $10\mu m$

trace width, and 30° bend angle. The single band transceiver design is adopted from [14]. The controller WI is operated at $44\,GHz$ for control signals.

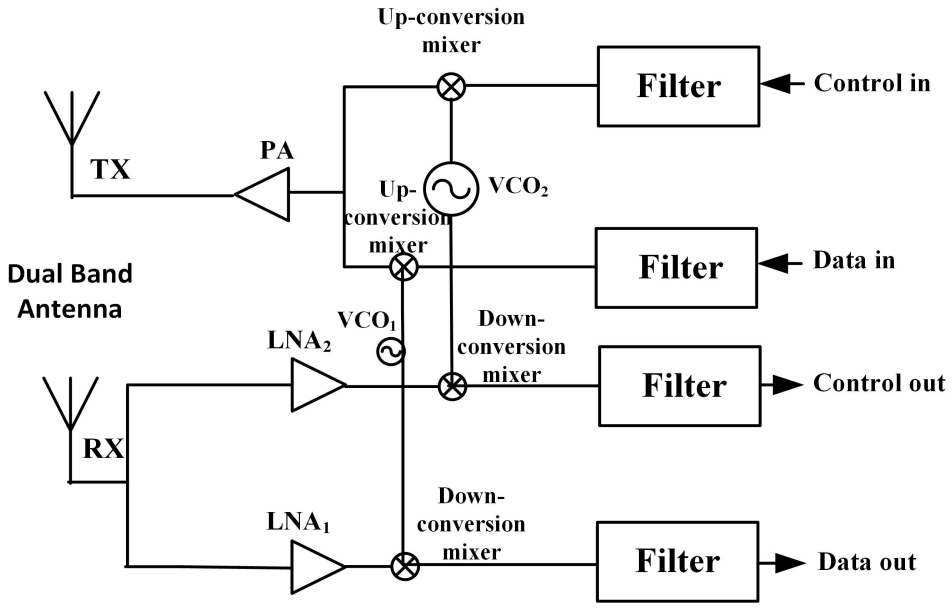

Figure 3.4: Dual Band Wireless Transceiver used for Controller WI

The dual band transceiver used for other WIs has been adopted from [39] [13] [27] as discussed in [34]. Non-coherent On-Off Keying (OOK) modulation scheme is used for low power consumption. The transmitter consists of a pulse shaping filter, up-conversion mixer, and a power amplifier as shown in the Figure 3.4. At the transmitter (T_X) side, the data and control signals are fed into a pulse shape filter and the filtered signal is then amplified by a power amplifier (PA). At the receiver (R_X) end, the received RF signals are fed into the Low Noise Amplifier (LNA) and mixer through pulse shaping filter. Two LNA; LNA_1 and LNA_2 are required for the two different frequencies of operation. Power hungry PLL can be replaced by injection-lock Voltage Controlled Oscillator (VCO) and direct conversion topology. The VCO is used to generate the carrier frequencies; LNA_1, VCO_1 are used for data and LNA_2, VCO_2 are used for control signals. To ensure high performance and energy efficient of WNoC, the transceiver circuit has to offer wide bandwidth and low power consumption. A planar log-periodic antenna [40] at millimeter wave range is used with these WIs. It is operated at $60\,GHz$ for data transmission.

Chapter 4

Performance Evaluation

In this chapter, the performance and overheads of the proposed DVFS controller is characterized using detailed full system evaluations. The controller is synthesized from a RTL level design using Design Compiler and Design Vision tools from Synopsys. The $65nm$ standard library from TSMC is used for synthesis. The design is driven at a clock frequency of $2\,GHz$.

The GEM5 [7] simulator is used for evaluating the performance. GEM5 is a cycle accurate full system simulator, which can simulate a complete system with devices and an operating system in full system (FS) mode. A system of 16 ALPHA cores running Linux operating system is considered for all evaluations. The cores are assumed to be independent of each other and the voltage/frequency for each core can be adjusted separately. This setup provides us with all statistics of the application running on the system including utilization, timing and hardware behavior.

The statistics from GEM5 are then fed into McPAT [33] to obtain core areas, run time power and energy for each run. The technology node in McPAT is set to $65nm$. The area data from McPAT is used to create a floorplan layout and ArchFP [18] tool is used for this purpose. Using the run time power and floorplan layout, HotSpot [26] thermal profiling tool is used to evaluate the thermal profile of the system without and with DVFS incorporated. Three PARSEC [6] benchmarks, BLACKSCHOLES, CANNEAL and DEDUP are considered in GEM5 FS mode [21] to study the system behavior. The benchmarks are run from beginning to end to obtain the statistics.

The latency and energy performance of different interconnects for various interconnect lengths is presented and the advantages of emerging interconnects over the traditional wires is discussed. The wireless interfaces provide better results than other interconnects but the actual delay is more compared to the ideal values. These values are presented in subsequent chapters and it is shown that WIs are still better in terms of delay. The overheads of different components is detailed in the last section.

4.1 Power Consumption

4.1.1 DVFS

The performance of the DVFS algorithms is evaluated under the presence of the three PARSEC benchmarks as discussed earlier. The execution time and the energy consumption with and without DVFS incorporated is shown in the Table 4.1.

Table 4.1: Core DVFS Simulation Results for PARSEC Benchmarks

Benchmark	**Energy Consumption (kJ)**			**Execution Time (sec)**		
	Normal	**DVFS**	**Savings(%)**	**Normal**	**DVFS**	**Penalty(%)**
BLACKSCHOLES	1.57	1.46	6.92	45.39	48.28	6.36
DEDUP	7.69	7.22	6.07	209.27	263.46	25.90
CANNEAL	8.99	8.46	5.82	252.48	324.65	28.60
FREQMINE	1.27	1.19	5.85	391.50	455.85	16.44
SWAPTIONS	2.66	2.46	7.27	81.24	78.40	-3.50

Negative values in savings or penalty fields indicate the opposite i.e., the energy consumption with DVFS is more than that with normal operation and execution time with DVFS is less than that with normal operation. As shown in the table, the proposed DVFS method achieves energy savings and the average reduction in energy is around 6.5% for all the benchmarks. The execution time in case of BLACKSCHOLES benchmark is increased by 6.36% and for DEDUP, the execution time is increased by 26%. In case of CANNEAL benchmark, the execution time with DVFS is 28% more compared to that of normal operation. The energy consumption with all cores operating at low power or high performance states is almost similar in all cases. The increase in power consumption due to higher voltage and frequency is compensated by the reduction in execution time.

4.1.2 Power Gating

The power gating technique is evaluated with 6 and 13 WIs present in the system. The transceivers operate at normal voltage when active. When the WIs are operated without incorporating power gating, the total power consumption for 6 WIs is $440.4mW$. With the proposed power gating method applied, at any instant of time, a maximum of $147.35mW$ power is consumed for single band wireless interfaces. Power savings of 66.54% can be achieved with this method over the normal operation. The hierarchical wireless NoC architecture proposed in [14] consumes $220.2mW$ for 6 WIs. Hence the proposed method improves the power consumption upto 33.085% compared to this method. As the number of cores and the system size increases, the required number of WIs for optimal performance also increases. The optimal number of WIs for different system sizes is discussed in [15]. The savings in power vary from 33.085% for 6 WIs to 69.12% for 13 WIs over the existing method.

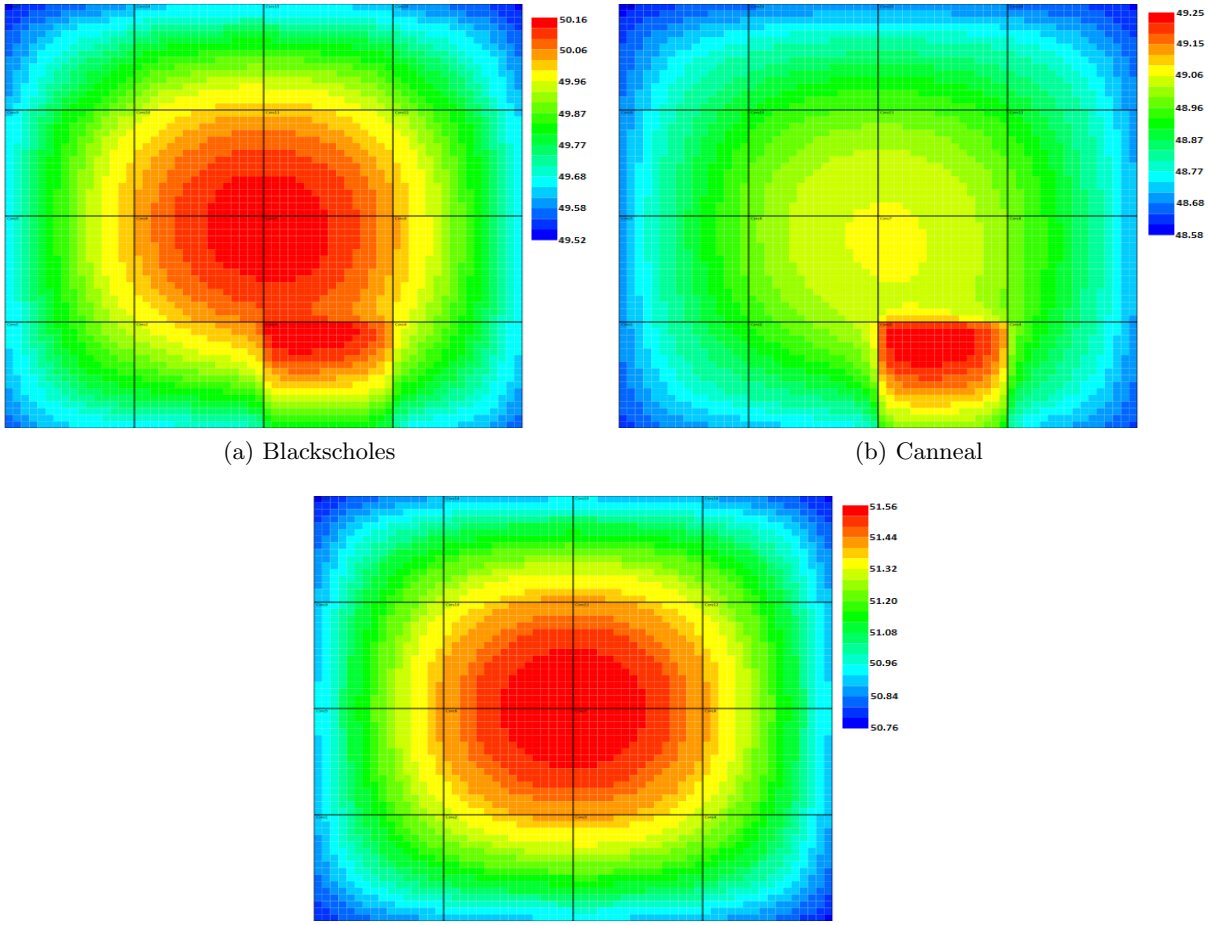

(a) Blackscholes

(b) Canneal

(c) Dedup

Figure 4.1: Thermal Profile Under Normal Operation for PARSEC Benchmarks

4.2 Thermal Profile

In this section, the overall thermal profile of the 16 core system is presented. With increasing energy densities on multi-core chips, the temperatures soar to very high values. Implementing DVFS can improve the thermal profile of the system.

The thermal profile of the system running PARSEC benchmark under normal operating conditions is shown in the Figure 4.1. As can be seen, many cores in the system reach to higher temperatures. Especially the cores at the center are affected by the heat spread from the surrounding cores. The system will not be able to sustain high temperatures for longer duration and its reliability is severely degraded.

There are many phases in the benchmark execution, where multiple cores are active only for a small duration in the total time slot. It is also observed that the high/low active phases generally occur continuously i.e., a core executes some task for some time and then goes into an idle state

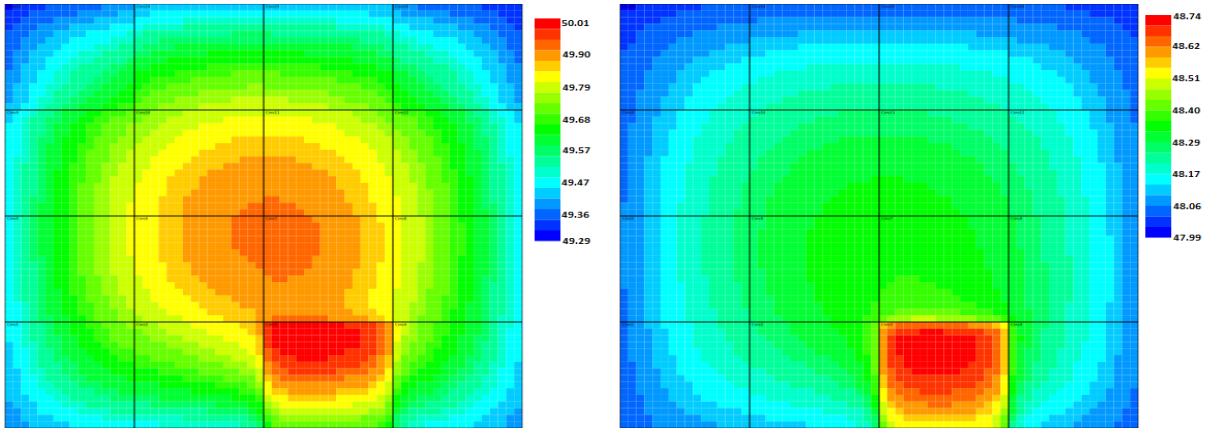

(a) Blackscholes (b) Canneal

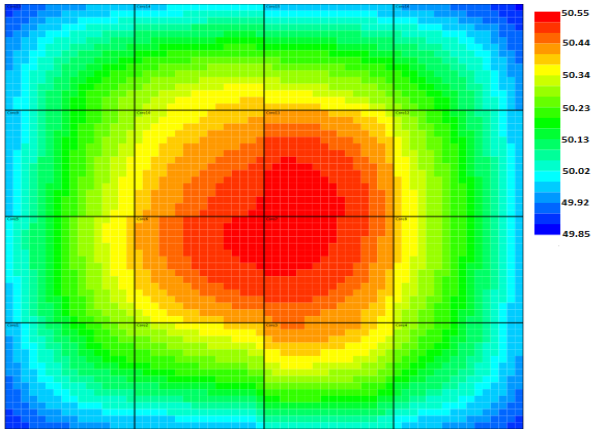

(c) Dedup

Figure 4.2: Thermal Profile Under DVFS Operation for PARSEC Benchmarks

due to stalls, etc. By applying DVFS to exploit these low active states, the temperatures can be reduced to lower sustainable levels for all benchmarks. Figure 4.2 shows the thermal profile of the chip with DVFS for all the benchmarks. Temperature reductions of 0.5%, 1.8% and 1.9% are observed for BLACKSCHOLES, CANNEAL and DEDUP benchmarks respectively. The number of cores attaining higher temperatures is also lowered compared to the normal operation. The spread from the hottest spot to neighboring cores is also significantly lower. The DVFS method seems to balance out the energy dissipation in the cores improving the overall temperature of the chip.

The thermal profile of the chip with all cores operating at low power or high performance states is also observed. Even though the energy consumption in both cases is equal, the cores reach to higher temperatures when operated at latter state. In fact, the coolest spot with high performance operation is hotter than the hottest spot with low power operation.

4.3 Interconnects

4.3.1 Latency Performance

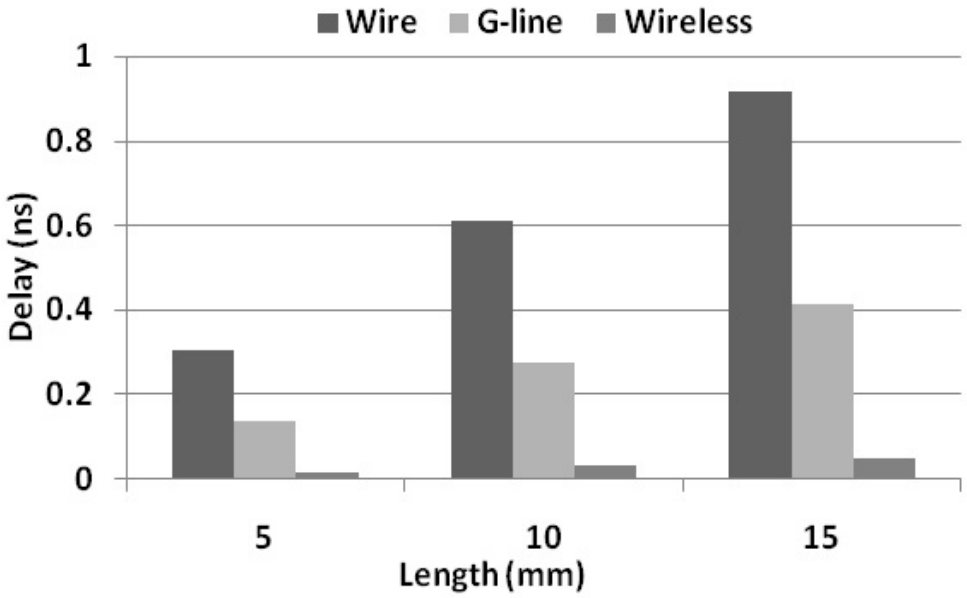

Figure 4.3: Variation of Latency with Interconnect Length for Different Interconnects

Considering a system size of $20mm \times 20mm$, the maximum interconnect length between the controller at center and different cores/clusters in the system can be $15mm$. Using conventional wired interconnects, the latency at this length is $917picosec$, which requires more than once clock cycle. The latency is reduced to $414picosec$ by using G-lines, but the energy per bit consumption increases. Using the wireless interfaces for controller communication, the latency for $15mm$ is $50picosec$. The latency at different interconnect lengths using wired, G-line and wireless interconnects is shown in the Figure 4.3.

4.3.2 Energy Per Bit

Wireless interfaces also provide efficient performance in terms of energy per bit consumption in data transmission. The energy per bit consumption in case of wired and G-line interconnects is $5.025pJ/bit$ and $7.038pJ/bit$ respectively for transmitting data over $15mm$ length. With WIs, it is $0.459pJ/bit$ for the same length. It is less than 10% compared to the wired energy and less than 7% compared to that of G-lines. The variation in energy per bit with length is shown in the Figure 4.4. As can be seen, the G-lines consume almost 50% more energy compared to wired interconnects at all lengths. Also, in both cases, the energy per bit consumption increases by more than 2 times as the length is doubled. The values are significantly lower with WIs and they scale by the same amount as that of interconnect length.

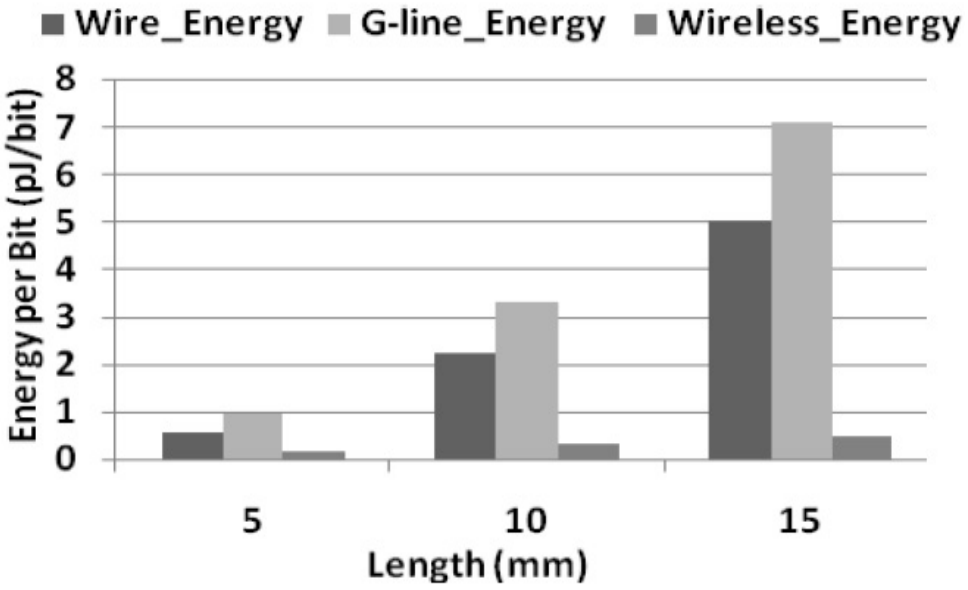

Figure 4.4: Variation of Energy per Bit with Interconnect Length for Different Interconnects

4.4 Overheads

The proposed controller occupies an area of $0.16mm^2$ for a system with 16 cores. As the number of cores in the system increases, the data sizes and controller area increases taking up significant area of the total chip. To keep the area overhead to minimum, time multiplexing of different clusters as described in the 3.3 section of chapter 3. In this case, the system is divided into four clusters and cores in each cluster are updated serially in 4 steps. Hence the controller area overhead remains the same. The controller consumes a total power of $0.5469mW$ of which $0.5384mW$ is the dynamic power and $0.008mW$ is the leakage power.

The dual-band wireless transceiver (inclusive of OOK modulator/demodulator, LNA_1 and LNA_2, PA, VCO_1 and VCO_2, Mixer) occupies a total area of $0.6mm^2$. It can sustain a data rate of $16Gbps$ with power consumption $73.4mW$. Two on-chip log-periodic planar antennas are simulated and integrated on the same substrate separated by a distance $20mm$. The antenna is $1.1825mm$ long. The return loss of planner log periodic antenna (S_11 parameter) from [40]shows that the antenna resonates at two frequencies, $44GHz$ and $60GHz$. The gain of this antenna is $38.65dB$ at $60GHz$.

Chapter 5

On-Chip Wireless Channel Modeling

Using Hybrid NoC architectures, as described in previous chapter certainly reduces the delay for long distance communication on chip with very low energy per bit consumption. The results presented in the previous chapter assume an ideal environment in the calculation of delay and energy per bit. In an ideal scenario, we assume that the signal transmitted from a wireless antenna will propagate through free space in the chip unhindered. But this is not the actual case since the signal propagation is affected by different components present in a chip. In a real chip environment, the signal propagation is affected by the silicon substrate and metal interconnects present in the chip. The delay associated with this signal is much different compared to the free space propagation delay. Hence the propagation of wireless signal in a realistic chip environment needs to be modeled and analyzed to obtain the actual improvements in delay over wired interconnects.

5.1 On Chip Model

To study and analyze the properties of on chip wireless channel, we characterize the chip structure using a two dimensional multi-layered model as shown in the Figure 5.1. The entire chip model is mainly comprised of four significant regions

1. Si Substrate Layer
2. Interconnect Layer
3. SiO_2 Passive Layer
4. Free Space

Si Substrate Layer

- It represents the bulk of the silicon wafer into which the transistors and gates of the design are embedded.

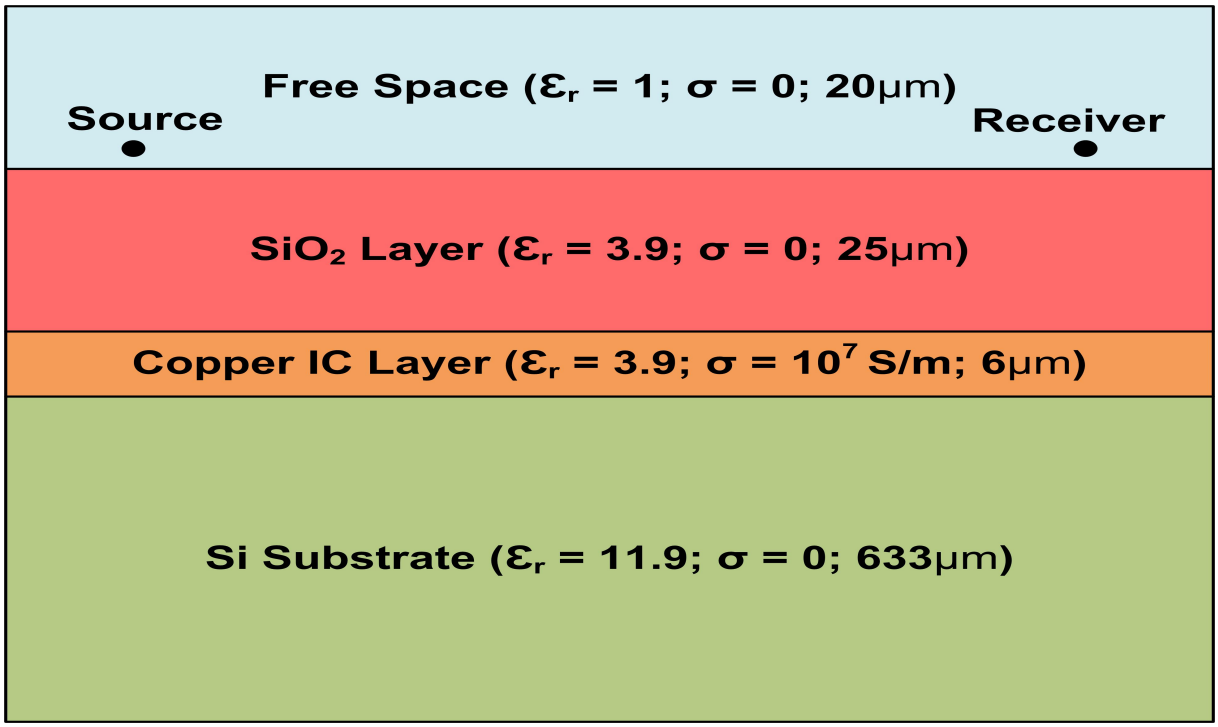

Figure 5.1: 2D Model Representing Different On Chip Components

- It is modeled as a single block of silicon with dielectric constant $\epsilon_r = 11.9$ and conductivity $\sigma = 0$.
- The thickness of the substrate layer is $633\mu m$.

Interconnect Layer

- It represents the stack of metal interconnects in the chip. The interconnects in the chip are stacked into multiple layers separated by a passive material like SiO_2.
- All the metal layers are characterized by a single entity in the model with dielectric constant $\epsilon_r = 3.9$ (ϵ_r of SiO_2) and conductivity $\sigma = 10^7 S/m$.
- The thickness of the interconnect layer is $6\mu m$.

SiO_2 Passive Layer

- It represents a passive layer grown over the below layers to prevent grounding of the antenna due to metal interconnects.
- The dielectric constant of the SiO_2 layer is $\epsilon_r = 3.9$ and its conductivity is $\sigma = 0$.
- The thickness of the free space layer is varied.

Free Space

- It represents the vacuum of air present between the actual die of the chip and its packaging materials.
- It is characterized by free space properties of zero conductivity and unit dielectric constant.
- The thickness of the SiO_2 layer is $20\mu m$.

All the dimensions for the layers are chosen to be same as in [44]. The interconnect dimensions used are obtained from PTM [2] $65nm$ technology models.

Source and Receiver

- A Hertzian dipole radiator with uniform radiation pattern in all directions is used as the source antenna.
- A wide band Gaussian pulse is used as the source signal to analyze the propagation characteristics at different frequencies.
- The source and receiver are placed just above the interface between free space and SiO_2 passive layer.
- They are separated by a distance of $1\mu m$ from the interface to account for the dimensions of the antenna.

The antenna is considered to be vertically polarized and hence only excites Transverse Magnetic (TM) modes in the structure. This is done to keep the propagation model simple so as to identify different components easily. The 2D chip model assumes that the layered structure is infinite in the third dimension unlike a real chip which has finite dimensions in all directions. There will be reflections and diffraction at these edges of the chip that effect the signal propagation. But these effects are not studied in the scope of this work.

5.2 FDTD Simulation Method

To characterize the electromagnetic propagation, Finite-Difference Time-Domain (FDTD) method is used. FDTD method, first proposed by Kane Yee in [46], is a numerical analysis technique used to model and solve problems in electromagnetics using Maxwell's equations. It employs finite differences as approximations to both the spatial and temporal derivatives that appear in the Maxwell's equations.

We use the two dimensional FDTD formulations to solve these equations within the on-chip problem space as described in previous section. One of the compelling features of the FDTD method is that its simplicity from one dimension is maintained in higher dimensions [41]. The computational complexity, although increased with number dimensions, is not as substantial as other numerical techniques. As we move to higher dimensions, multidimensional arrays are required to represent and store the multidimensional grid data. The two dimensional computational complexity and memory requirements for different problem space sizes are discussed in further sections.

5.2.1 FDTD Equations and Algorithm

In a 2D electromagnetic simulation with direction of propagation in z-dimension, the TM mode is composed of the field components $E_z,\ H_x\ and\ H_y$. The Maxwell's equations for this case can be written as

$$\frac{\partial D_z}{\partial t} = \frac{1}{\sqrt{\epsilon_0 \mu_0}}(\frac{\partial H_y}{\partial x} - \frac{\partial H_x}{\partial y}) \tag{5.1}$$

$$D_z(\omega) = \epsilon_0 \epsilon_r^*(\omega) E_z(\omega) \tag{5.2}$$

$$\frac{\partial H_x}{\partial t} = -\frac{1}{\sqrt{\epsilon_0 \mu_0}}(\frac{\partial E_z}{\partial y}) \tag{5.3}$$

$$\frac{\partial H_y}{\partial t} = \frac{1}{\sqrt{\epsilon_0 \mu_0}}(\frac{\partial E_z}{\partial x}) \tag{5.4}$$

Where,

$\epsilon_0,\ \mu_0$: Free space permittivity and permeability
ϵ_r^* : Effective permittivity of the medium

In FDTD algorithm, the space and time are discretized so that the electric and magnetic fields are staggered in both space and time. All the partial derivatives in the equations (5.1)-(5.4) are replaced with finite differences. The resulting difference equations (5.5), (5.6) and (5.7) are then solved to obtain future fields from past fields. The magnetic and electric fields in the next future time step are evaluated alternatively till the end of simulation duration.

$$D_z(i,j,n) = D_z(i,j,n-1) + \frac{1}{2C_0}[H_y(i,j,n) - H_y(i-1,j,n) - H_x(i,j,n) + H_x(i,j-1,n)] \tag{5.5}$$

$$H_x(i,j,n+1) = H_x(i,j,n) + \frac{1}{2\mu C_0}[E_z(i,j,n) - E_z(i,j+1,n)] \tag{5.6}$$

$$H_y(i,j,n+1) = H_y(i,j,n) + \frac{1}{2\mu C_0}[E_z(i+1,j,n) - E_z(i,j,n)] \tag{5.7}$$

Where,

$i,\ j$: x-dimension and y-dimension space step
n : Time Step
C_0 : Speed of Light (3×10^8 m/sec)
μ : Permeability of the medium

The FDTD algorithm in two dimensions can be best illustrated using the leap frog diagram shown in the Figure 5.2 [43].

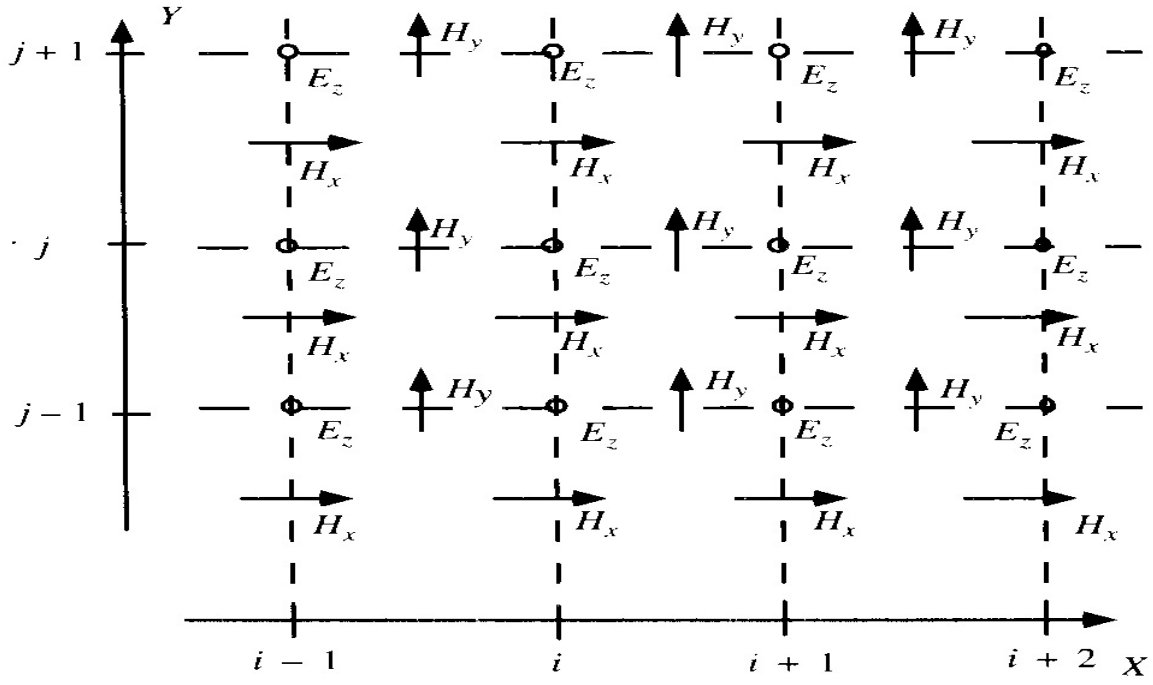

Figure 5.2: Interleaving of E and H Fields for 2D FDTD TM Formulation

After setting up the problems space parameters and initial values of the fields, the magnetic field components are updated at current time instant are computed using field components at n_{th} time step as past fields. Then the electric field components are updated and the time step is incremented to $n+1$. This process is continued till the last time step iteration.

In order to descretize space and time, the terms Δx, Δy and Δt needs to defined which determine the resolution of the problem space. Δx and Δy are the smallest space measurements that can be made in the x- and y-dimensions respectively. Δt is the smallest time that can be measured or observed. The space steps are chosen depending upon the wavelength corresponding to the maximum frequency component present in the source signal. In general, for effective implementation, equation (5.8) gives the ideal space step value.

$$\Delta x = \Delta y = \lambda_{min}/10 \tag{5.8}$$
$$\lambda_{min} = C_0/f_{max}$$

Where,

f_{max} : Maxmimum frequency component in the signal

In case of non-homogeneous grids, Δx and Δy can be chosen to be different but must adhere to the minimum conditions given in equation (5.8). Unless specifically specified, in the remainder of this work, we consider $\Delta x = \Delta y$ and is denoted by Δx. Based on the chosen space step, equation (5.9) gives the condition for the time step.

$$\Delta t \ \leq \ \Delta x/2C_0 \tag{5.9}$$

For non-homogeneous grid,

$$\Delta t \ \leq \ \sqrt{\Delta x^2 + \Delta y^2}/2C_0 \tag{5.10}$$

5.2.2 Perfectly Matched Layer (PML)

One of the long standing issues with the use of FDTD method is boundary conditions. In numerical methods like FDTD, the size of the problem area that can be simulated is determined the computational resources available. In the real life, the problem space is not confined. The region of interest is surrounded by other mediums, like in case of on-chip model the chip is surrounded by air or external components on board. But as wave propagates in FDTD simulation, it eventually reaches the edge of the defined problem space and unpredictable reflections will be generated and propagate inwards. And it is not possible to identify the unwanted components from real signal. To alleviate this issue, Absorbing Boundary Conditions (ABCs) are defined and the Perfectly Matched Layer (PML) [5] is one of the most flexible and efficient ABCs.

The complex coordinate approach of PML is based on analytic continuation of Maxwell's equations into complex spatial coordinates where fields are exponentially decaying. It adds an absorbing anisotropic boundary layer around the region of interest as shown in Figure 5.3, which prevents an inward reflection and decays the signal as it propagates forward. The basic idea of PML is based on the fact that the amount of reflection as a wave propagates from one medium to another is depended on the intrinsic impedances of the two media.

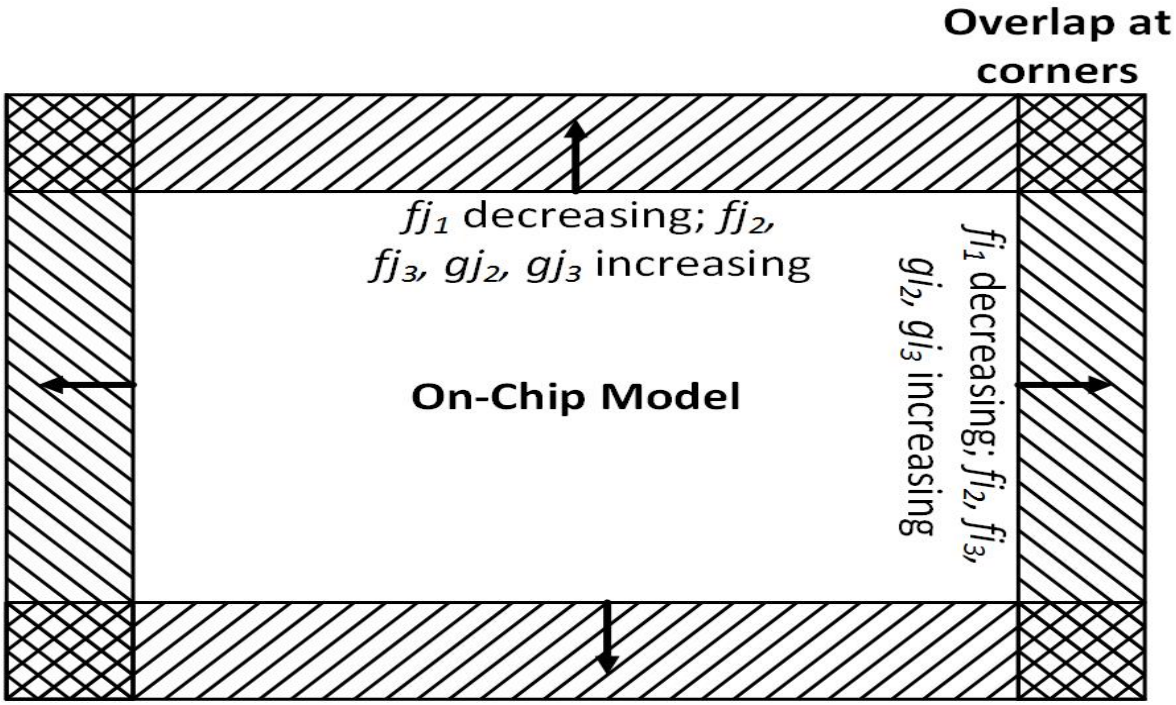

Figure 5.3: On-Chip Model Problem Space with PML

$$\Gamma = \frac{\eta_A - \eta_B}{\eta_A + \eta_B} \tag{5.11}$$
$$\eta = \sqrt{\frac{\mu}{\epsilon}}$$

For this μ is changed with ϵ so that η remains constant and there is no reflection. And we want the wave to decay completely before it hits the boundary. To accomplish this, ϵ and μ are made complex [43] since imaginary part causes decay of the wave. The terms f and g shown in the figure are the terms introduced into FDTD equations to accommodate the conductivity terms in the imaginary part of ϵ and μ. These terms are varied as represented to vary conductivities accordingly in the PML region.

Although PML offers an efficient solution for boundary conditions and works well in most cases, it suffers from some unavoidable reflections. Once the wave is discretized for simulations, some small numerical reflections appear. These reflections distort the actual signal slightly and the effects are analyzed and discussed in the results section.

5.2.3 Memory Requirements

Based on the maximum frequency of operation, the upper bounds for the space and time resolutions needed in effective implementation of FDTD are decided. FDTD method works effectively in the domains where the characteristic dimensions of the problem space are of the order of wavelength in size. Within these constraints, the space resolution, Δx is chosen such that the smallest dimension in the problem space can be modeled by minimum required number of grid steps. The time step is then fixed according to the equality condition in equation (5.9).

So, the smallest dimension in the problem space decides the total number of grid steps required to represent it and the number of time steps in the total simulation period. Since the electric field, magnetic field and related variables need to be defined at all points in the problem space, two dimensional arrays are required to store the field data in 2D FDTD simulations. And the total memory requirement is decided by the resolution chosen and total number of grid steps. Any variation in Δx reflects in a corresponding quadratic change in memory requirements and cubic change in execution time.

For example, if a new resolution is chosen such that $\Delta x_{new} = \Delta x/2$, the number of grid steps to represent the same horizontal and vertical dimensions is twice the previous value. The number of time steps required is also twice the previous value. The array sizes to represent the field variables is four times the previous array size and so is the memory requirement. The total execution period is eight times the previous time. Hence any change in the space-time resolution results in a significant change in memory requirements and execution time.

In our 2D on-chip model case, the smallest dimension is the interconnect layer that represents the metallic wires in the chip. Each metal wire is of the order of $1\mu m$ and the total layer thickness is less than $10\mu m$, whereas the SiO_2 passive layer, free space layer are of order of tens of μm, Si substrate is order of hundreds of μm and transmitter-receiver distance is of order thousands of μm. Due to this disproportionate lengths, the total number of grid steps is significantly high for our problem space; of the order of millions of steps. The actual values chosen and memory requirements are discussed in the results section.

Chapter 6

Propagation Results

6.1 Experimental Setup

6.1.1 Problem Space

The FDTD method is implemented using the MATLAB [1] tool. We used the version R2011a and R2013b versions of MATLAB. To keep the memory requirements and execution time within tolerable limits in the initial stages, we have scaled down the dimensions in the original problem space described in the section 5.1 of chapter 5. We present the results for these two problem spaces. The dimensions used in each case are presented in the Table 6.1

Table 6.1: Dimensions of the Problem Space

Problem Space	**Width**	**Height**				**PML**
		Si	**Interconnect**	**SiO_2**	**Free Space**	**Thickness**
$\rho = 100\mu m$	$250\mu m$	$25\mu m$	$6\mu m$	$25\mu m$	$20\mu m$	$8\mu m$
$\rho = 1mm$	$1150\mu m$	$633\mu m$	$6\mu m$	$25\mu m$	$20\mu m$	$8\mu m$

6.1.2 PML

To keep the unwanted reflections from PML boundary to a minimum, we have tried different PML thickness sizes. In both cases, thickness of $8\mu m$ has given us the best results. And to reduce the effect of PML reflections on source, a separation of $75\mu m$ is maintained between source and PML boundary layer in both the cases. The same is also maintained at the receiver.

6.1.3 Source and Receiver

The source used is a wide band Gaussian mono pulse. We have chosen a wide band signal to analyze the propagation characteristics at different frequencies. The field is also analyzed at different receiver distances to observe the effect of distance on received signal. The source signal properties, step sizes and receiver distances used for both cases are presented in the Table 6.2.

Table 6.2: Source Signal Properties

Problem Space	**Mean** ($picosec$)	**Variance** ($picosec$)	**Maximum Receiver Distance**	**Space Step** (μm)	**Time Step** ($femtosec$)
$\rho = 100\mu m$	0.03	0.005	$100\mu m$	0.25	0.4167
$\rho = 1mm$	1	0.1	$1000\mu m$	0.5	0.8333

6.1.4 Step Size

The space step and time step used in both cases are shown in the Table 6.2. The values are chosen to satisfy the conditions in equation (5.8) and to keep the array sizes within limits. With the values chosen, all the fields should be calculated at 1064×368 steps for the problem space, $\rho = 100\mu m$ and at 2332×1400 steps for the problem space, $\rho = 1mm$.

6.1.5 Simulation Time

The execution time is chosen such that all the components of the field reach the receiver point and to meet the required frequency resolution. In case of scaled problem space, the execution time is 2 *picosec* and for original problem space, it is 33.333 *picosec*.

6.2 $\rho = 100\mu m$ Problem Space

To identify different components from each layer in the model, we have first run the simulations using three variants of the model,

- Free Space Model, in which all the layers in the problem space are modeled as free space. This give us an understanding of the direct wave traveling through free space.
- Air-SiO_2-Si Model, in which all layers except the interconnect layer are modeled. This gives us the effect of reflections from SiO_2 passive layer on the received signal.
- Air-SiO_2-Cu-Si Model, which contains all the layers and gives us the effect of interconnects on received signal.

6.2.1 Free Space Model

Figure 6.1 shows the received electric field at a distance of $100\mu m$ from the source and in direct line of sight. The Direct Wave, shown in the figure, is the wave traveling through free space at the speed of light ($C_0 = 3 \times 10^8 m/sec$) from source to receiver and is the desired significant component of the received signal. The delay in received signal is the time taken by the wave to travel $100\mu m$ in free space,

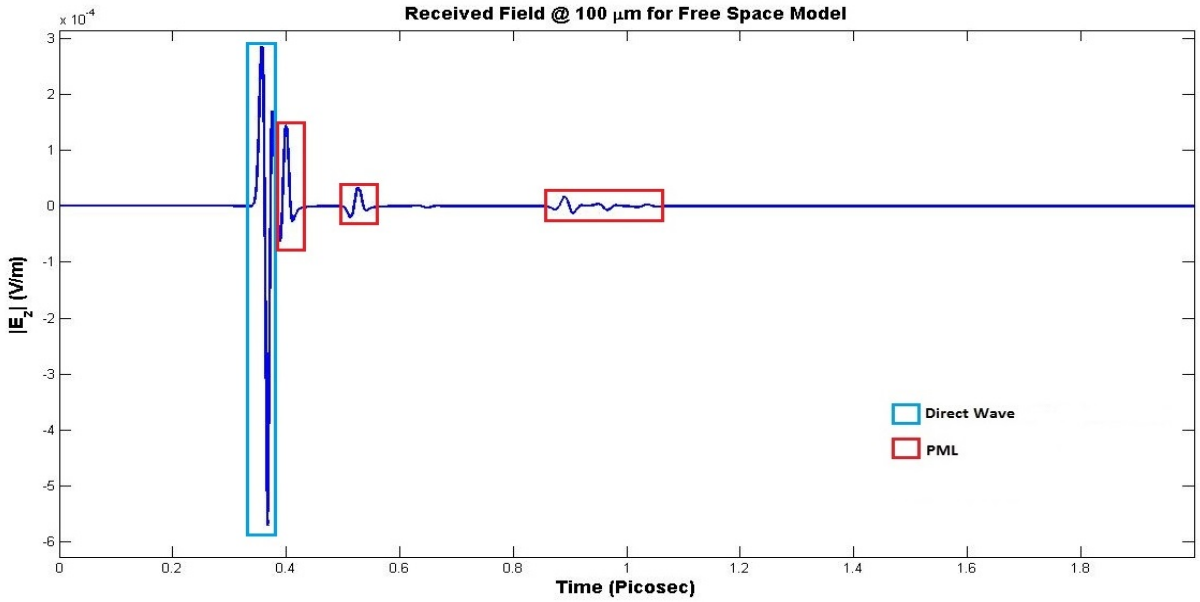

Figure 6.1: Received Field vs Time for Free Space Model

$$t_{pd_{direct}} = (100 \times 10^{-6})/(3 \times 10^{8}) = 0.333 picosec \tag{6.1}$$

The source peak is at 0.03 *picosec* and hence received signal should have its peak at 0.363 *picosec* and the simulation results show approximately the same. Since the PML is not ideal, the reflections from top, left and bottom boundaries eventually reach the receiver and are as marked in the Figure 6.1.

6.2.2 Air-SiO_2-Si Model

The electric field at a distance of $100\mu m$ and $1\mu m$ above the Air-SiO_2 interface is shown in the Figure 6.2. The surface wave component is the field propagating along the Air-SiO_2 interface. There are multiple reflections from the SiO_2-Si interface within the SiO_2 layer. These waves reach the receiver much later because they travel at approximately half the speed of light in the SiO_2 medium. Finally, the direct wave component as can be seen from the figure is much weaker as compared to the free space component, which is due to the reflection from Air-SiO_2 interface. This can be explained using the perpendicular polarized [4] wave incident at the interface of two mediums as shown in Figure 6.3.

The reflection coefficient in the above case is given by equation (6.2),

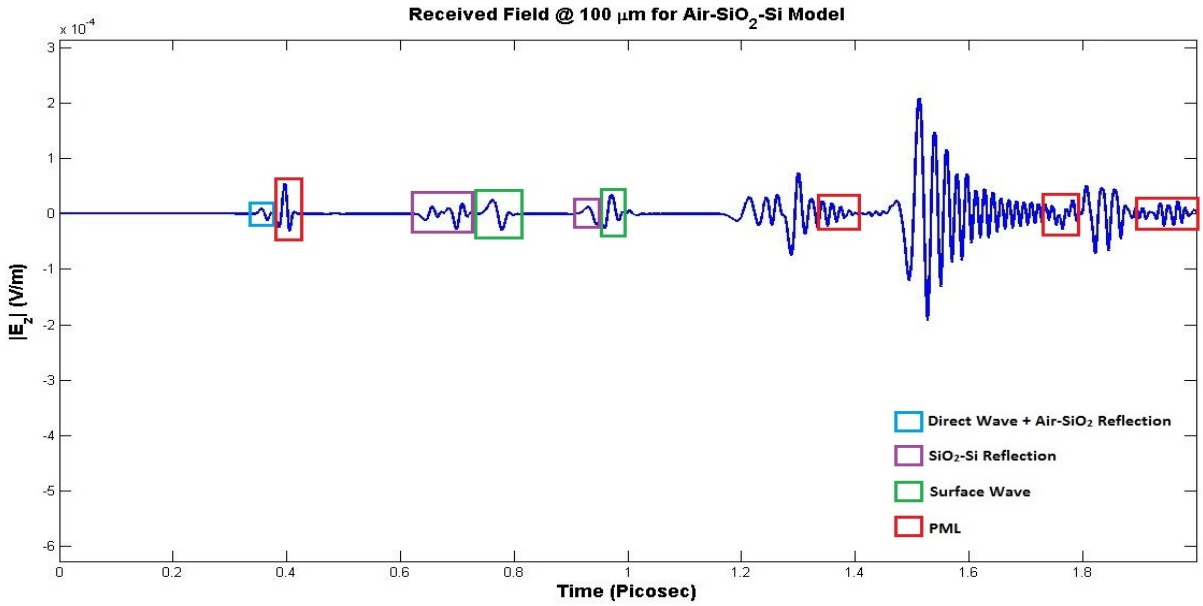

Figure 6.2: Received Field vs Time for Air-SiO_2-Si Model

$$\Gamma = \frac{\eta_2 \cos\theta_i - \eta_1 \cos\theta_t}{\eta_2 \cos\theta_i + \eta_1 \cos\theta_t} = \frac{\sqrt{\mu_2/\epsilon_2}\cos\theta_i - \sqrt{\mu_1/\epsilon_1}\cos\theta_t}{\sqrt{\mu_2/\epsilon_2}\cos\theta_i + \sqrt{\mu_1/\epsilon_1}\cos\theta_t} \tag{6.2}$$

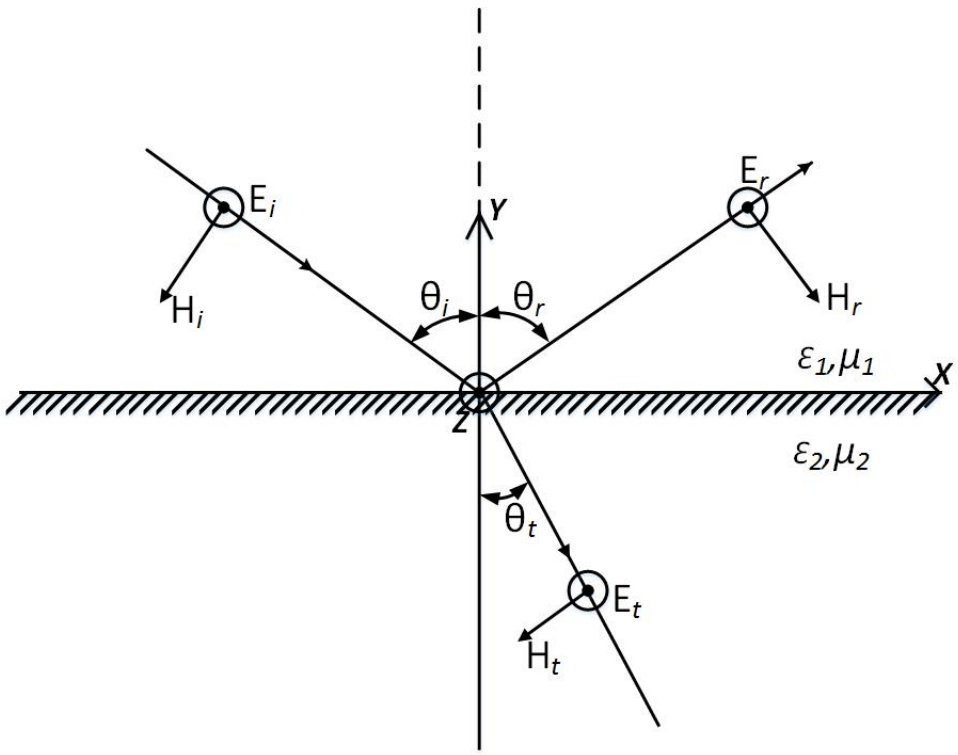

Figure 6.3: Perpendicular Polarized Wave Incident at Oblique Angle

Where,

$\eta_1,\ \eta_2$: Characteristics Impedances of the Mediums
θ_i : Angle of Incidence
θ_t : Angle of Transmission

Since for most dielectric media, $\mu_1 \approx \mu_2$, the equation in (6.2) is reduced to equation (6.3).

$$\Gamma|_{\mu_1=\mu_2} = \frac{\cos\theta_i - \sqrt{\epsilon_2/\epsilon_1}\sqrt{1-(\epsilon_1/\epsilon_2)\sin^2\theta_i}}{\cos\theta_i + \sqrt{\epsilon_2/\epsilon_1}\sqrt{1-(\epsilon_1/\epsilon_2)\sin^2\theta_i}} \tag{6.3}$$

Since the source separation from the Air-SiO_2 interface is very small ($1\mu m$) compared to the distance between source and receiver ($100\mu m$), the angle of incidence at the interface is $\theta_i \approx 90°$. Substituting in equation (6.3), the reflection coefficient is $\Gamma \approx -1$ irrespective of the mediums. If E_i and E_r are the incident and reflected fields, then $E_r = \Gamma E_i$. And the distance traveled by the reflected wave ($100.02\mu m$) is approximately same as the direct wave and both components reach the receiver at the same time. Hence, the total field in the incident medium, given by $E_1 = E_i + E_r$, is almost zero near the interface of two mediums and the desired direct wave component is completely canceled out by the equal and out-of-phase reflected wave component.

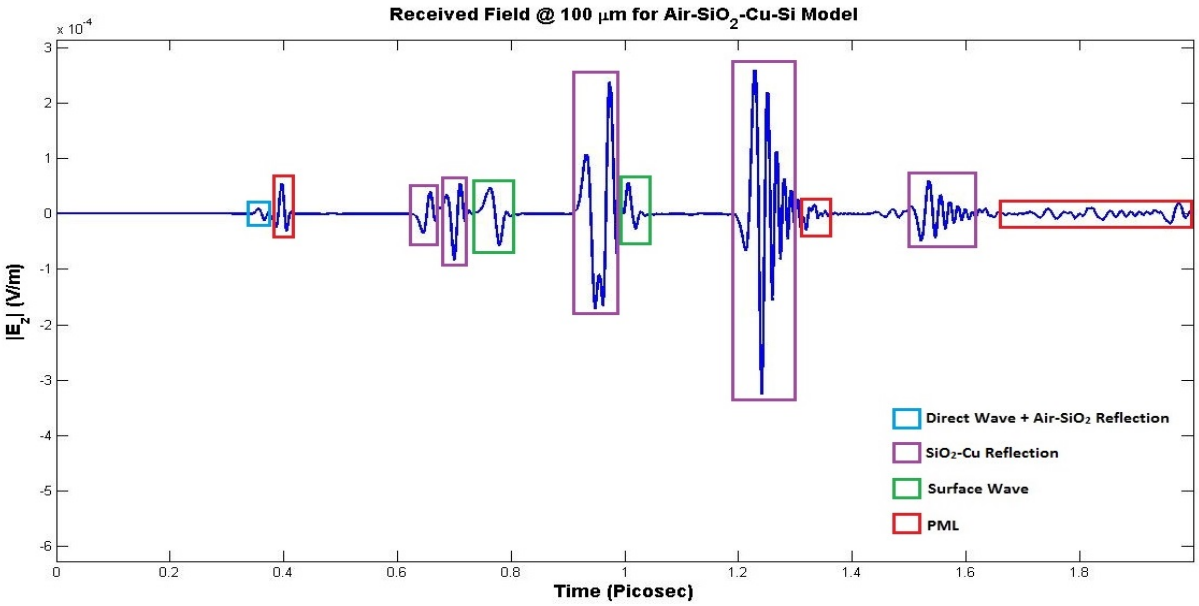

Figure 6.4: Received Field vs Time for Air-SiO_2-Cu-Si Model

6.2.3 Air-SiO_2-Cu-Si Model

Finally in the scaled model with all layers modeled, the received electric field and its components are as shown in the Figure 6.4. The direct wave component and the Air-SiO_2 reflection remain the same and cancel out each other as in the Air-SiO_2-Si case. Since the interconnect layer is a conducting medium, the field incident on the SiO_2-Cu interface is completely reflected and then transmitted to the the receiver through Air-SiO_2 interface. Multiple reflections exist within SiO_2 layer that eventually reach the receiver the same way. The residual surface wave components from the field transmitted/reflected into the SiO_2 layer propagates along the interface to reach the receiver.

6.2.4 Wave Components

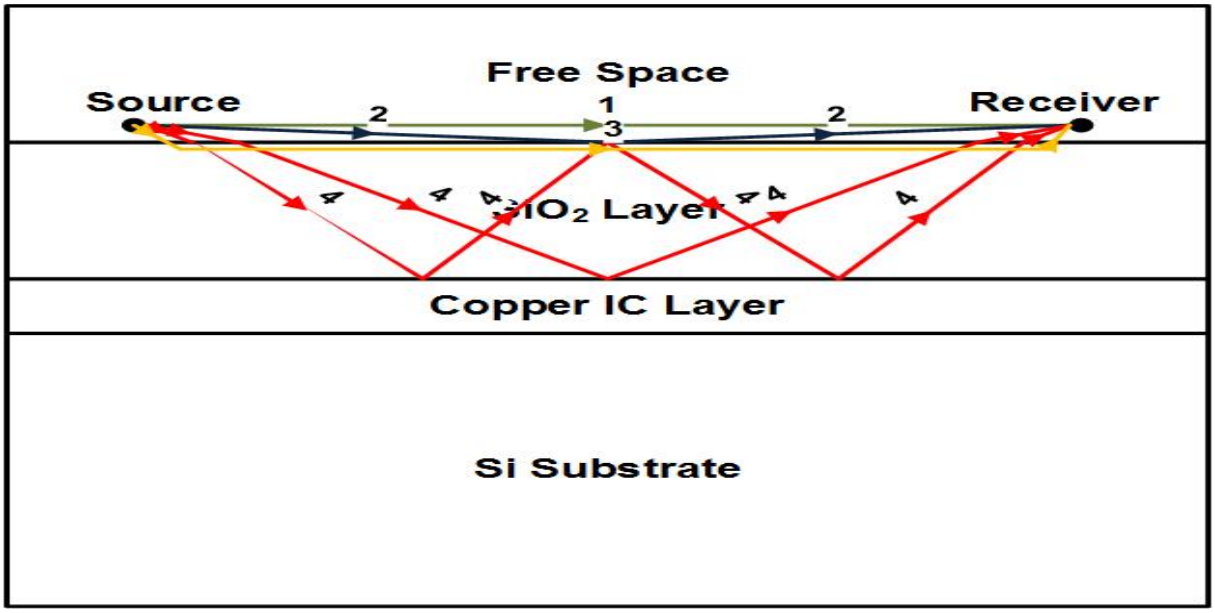

Figure 6.5: Different Wave Components in the Received Electric Field

The different wave components in the received field can be represented as shown in the Figure 6.5. The first component is the direct source-to-receiver contribution and the second component is the reflection from the Air-SiO_2 interface. The third component is the sum of all residual surface wave components along the interface. The fourth component represents the multiple reflections from the interconnect layer and within the SiO_2 layer. The propagation delay of each of these components can be calculated using the reflection and transmission theory for oblique incidence at the interface of two mediums.

The magnitude of the received electric field as a function of frequency is shown in the Figures 6.6, 6.7. The phase is shown in the Figure 6.8. Using range gating, we have separated different

wave components from the total field. As can be seen the surface wave and reflection wave components dominate over the direct wave component. The trend remains the same over the entire frequency range. At very high frequencies, beats exist in the reflection and surface wave components, which can be due to the multiple reflections within the SiO_2 layer that can interfere destructively to give rise to nulls in the signal. The phase of the different wave components vary linearly with frequency.

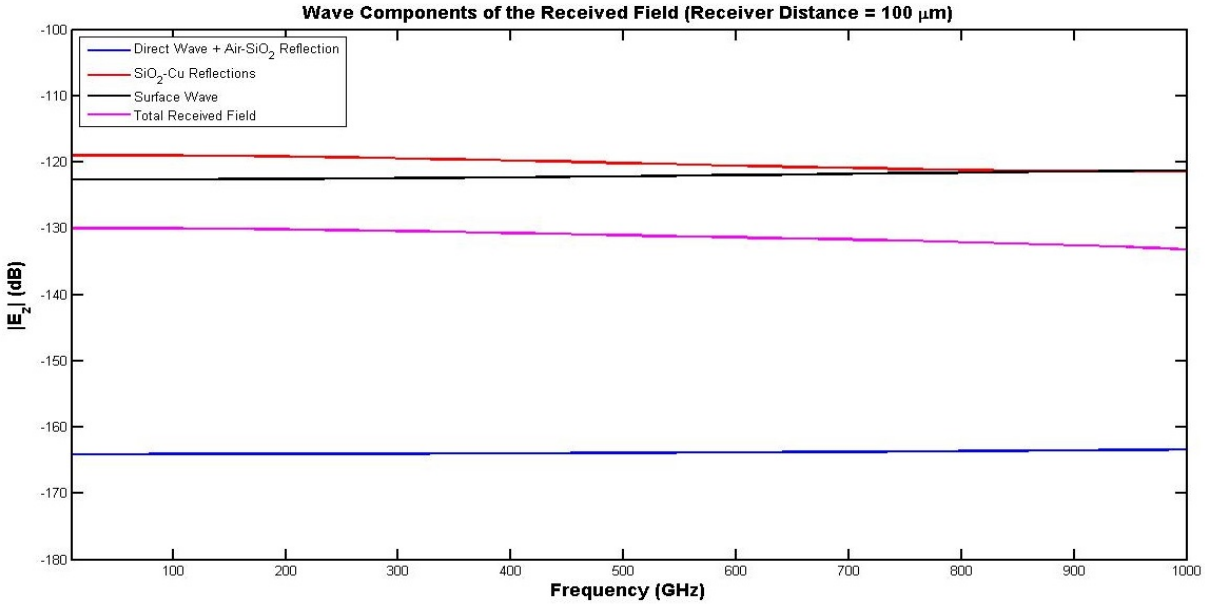

Figure 6.6: Magnitude of the Electric Field at the Receiver

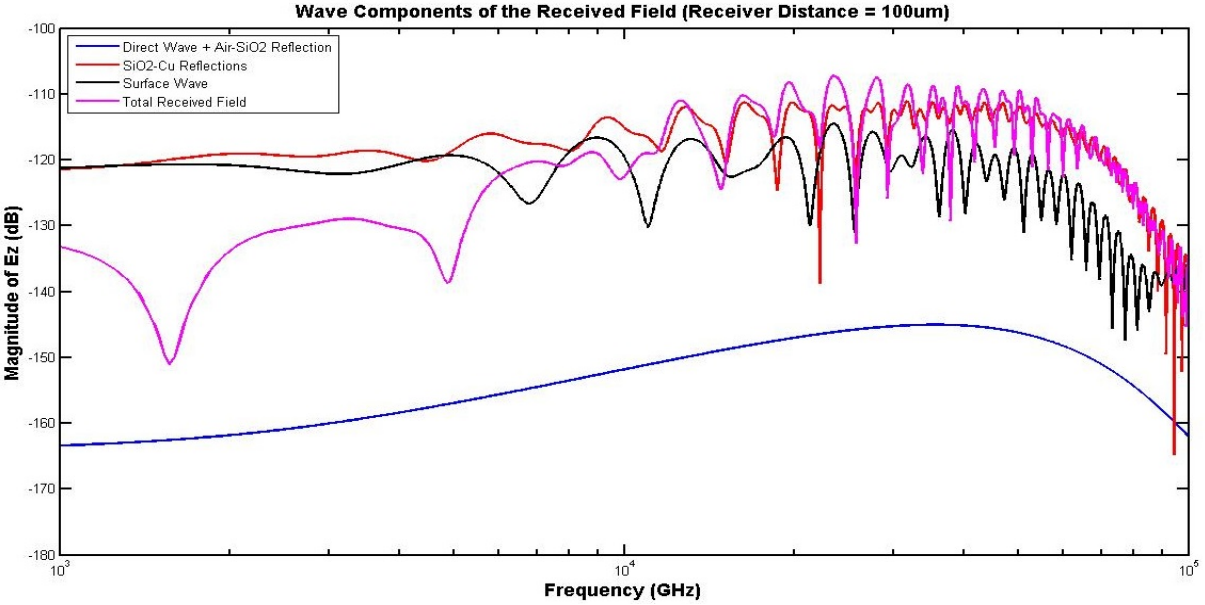

Figure 6.7: Magnitude of the Electric Field at the Receiver

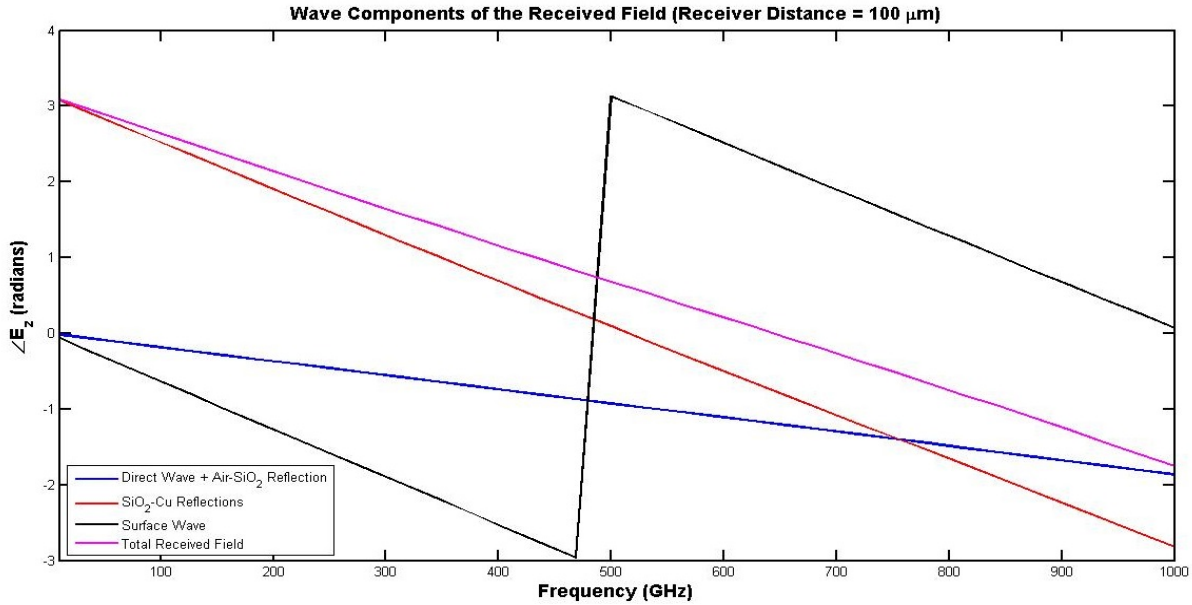

Figure 6.8: Phase of the Electric Field at the Receiver

6.3 $\rho = 1mm$ Problem Space

The received electric field at a distance of $1mm$ from the source for SiO_2 thickness of $30\mu m$ is shown in the Figure 6.9. The direct wave component phenomenon remains the same and is much more weaker (almost zero) as the field decays as $1/\rho$ with respect to distance ρ from the source. The reflections from the SiO_2-Cu interface completely dominate the received electric field.

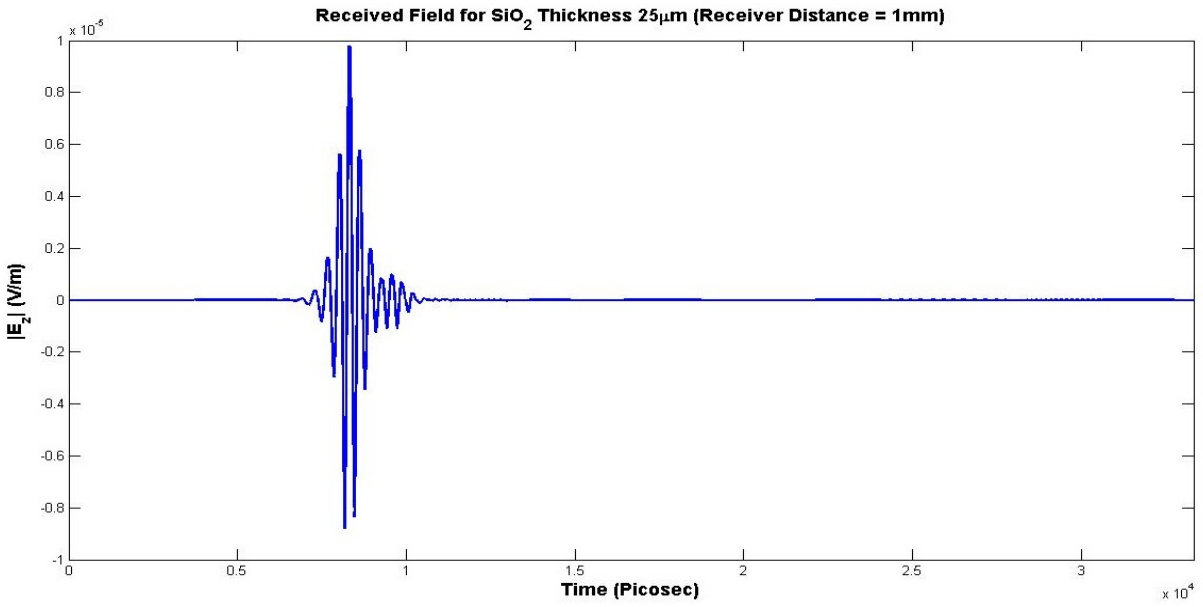

Figure 6.9: Received Electric Field at a Distance of 1mm from Source

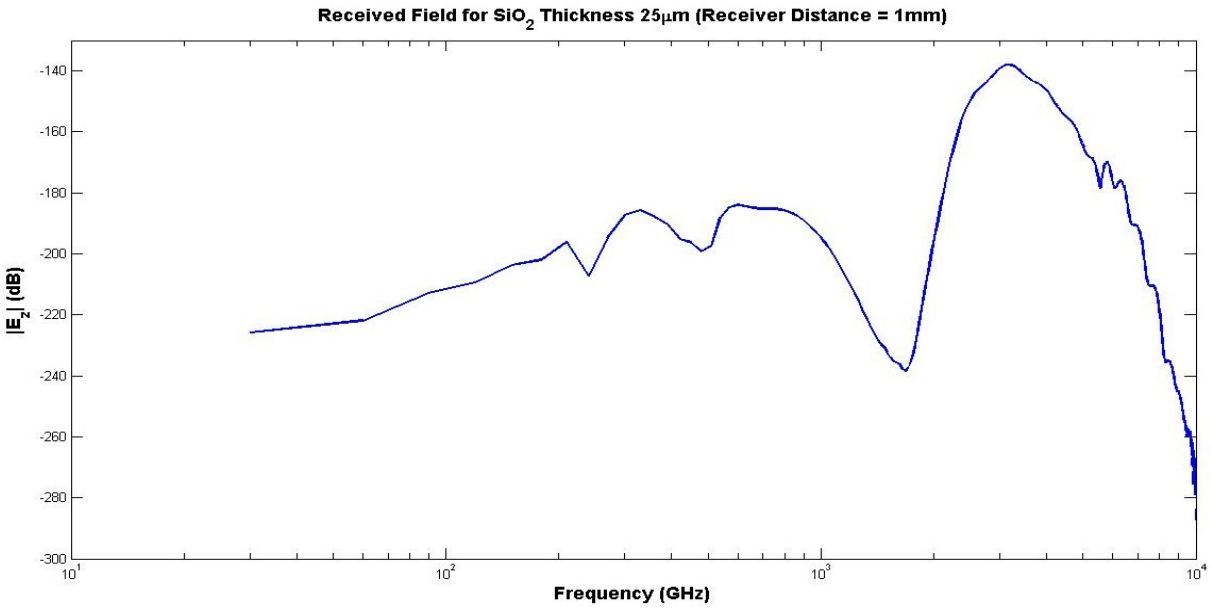

Figure 6.10: Received Electric Field at a Distance of 1mm from Source

The variation of the received field with respect to frequency is shown in the Figure 6.10. The field variation with frequency can be divided into mainly three frequency regions. A significant portion of the field lies in the very high frequency range above 1500 GHz. At lower frequency ranges i.e., below 120 GHz, the field strength is low but varies almost linearly with frequency. In the range between, the field is distributed into two small bands within which the field is mostly linear. The field in the middle range is stronger than that at lower ranges, but still is much lower than that at higher frequencies.

6.3.1 Variation with Passive Layer Thickness

To investigate the effect of SiO_2 passive layer on the received field, we varied the thickness of SiO_2 layer and observed the received electric field at a distance of $1mm$ from the source. The thickness is varied from $2\mu m$, $5\mu m$, $10\mu m$ till $60\mu m$. The variation of the field with respect to frequency for different values of thickness is shown in the Figure 6.11.

It can be seen that as the thickness of the SiO_2 layer increases, the strength of the received field increases. The variation of the field with respect to frequency at all thickness is similar, distributed into three separate frequency bands. But this is not the case for SiO_2 thickness of $2\mu m$. At frequencies below 1 THz, the variation remains the same, but the field falls off instead of increasing. The same effect, if not so pronounced can be observed for thickness of $10\mu m$ and $15\mu m$.

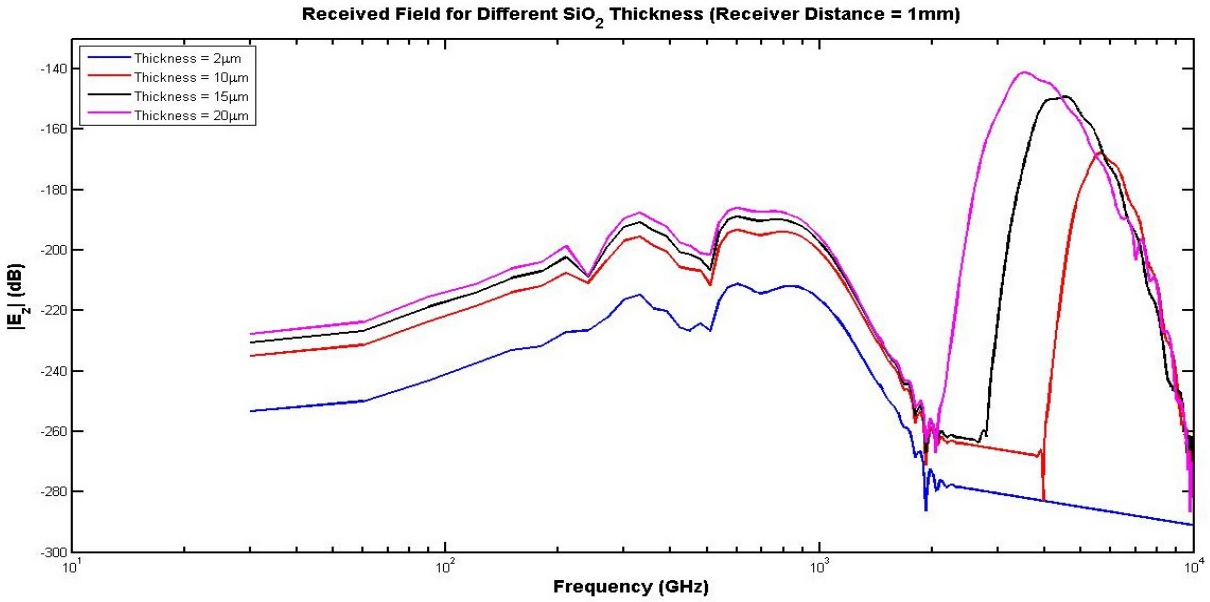

(a)

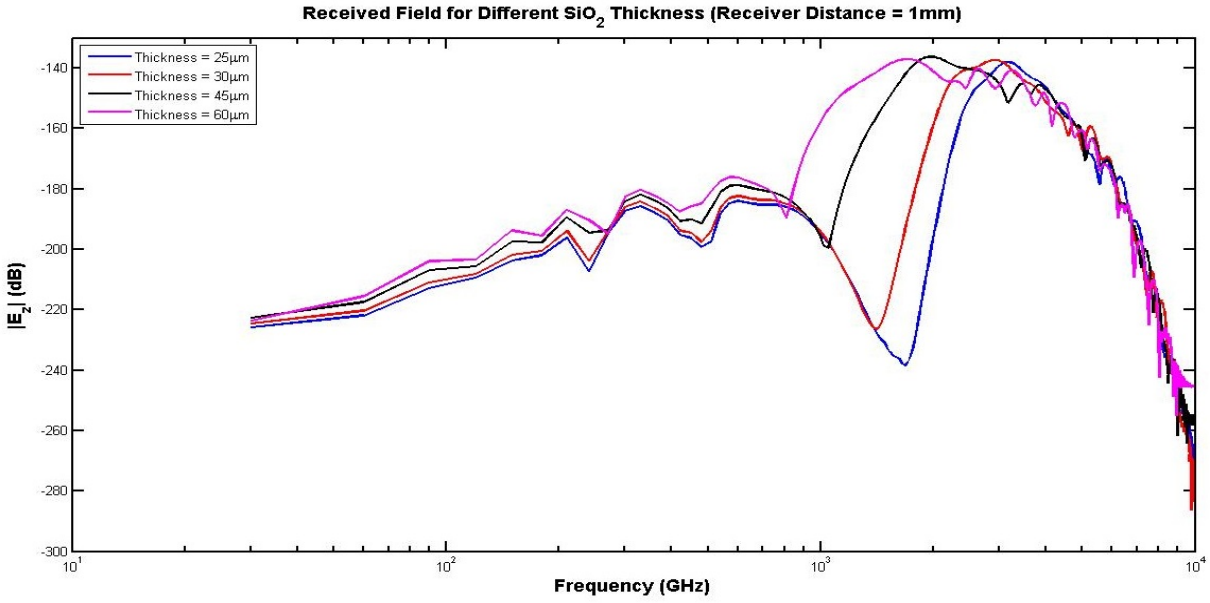

(b)

Figure 6.11: Variation of Received Electric Field with Frequency for Different SiO_2 Layer Thickness

As the thickness of SiO_2 layer increases, the null present at higher frequencies becomes less significant i.e., the strength of the field at the null increases with thickness. At $60\mu m$ and $45\mu m$, the variation of field in middle and higher frequency bands is not so clearly demarcated.

The variation of the normalized electric field with frequency is shown in the Figure 6.12. As can be seen, the field is much stronger at high frequencies. As the SiO_2 layer thickness increases, the frequency band in which the significant portion of the energy is concentrated shifts towards the lower frequencies. Also the bandwidth of the field increases with increasing thickness. But

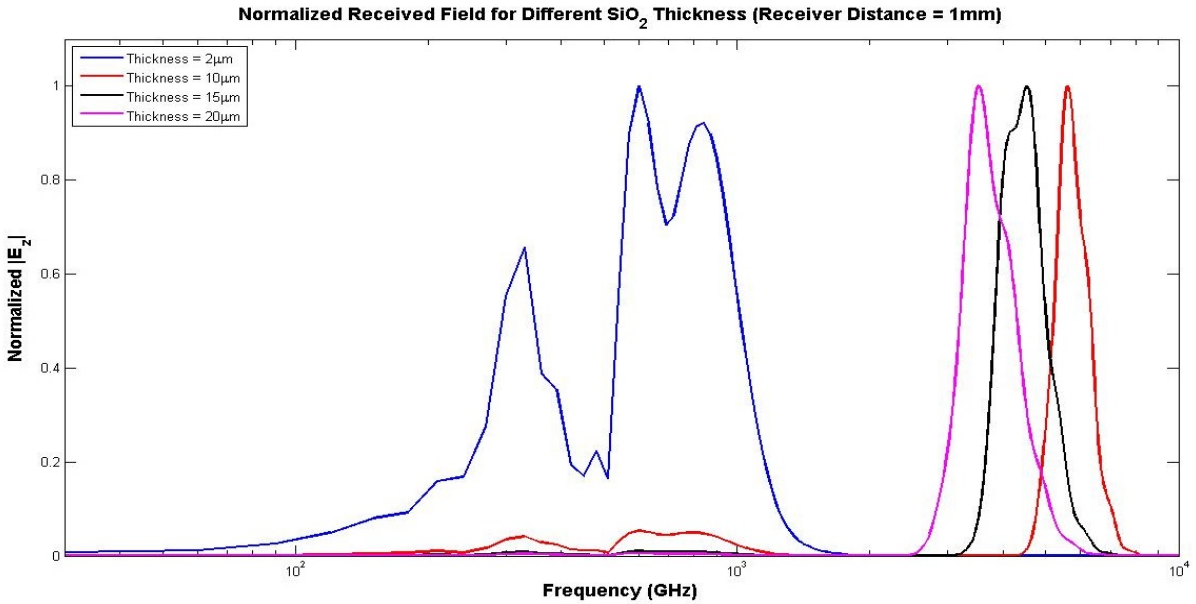

(a)

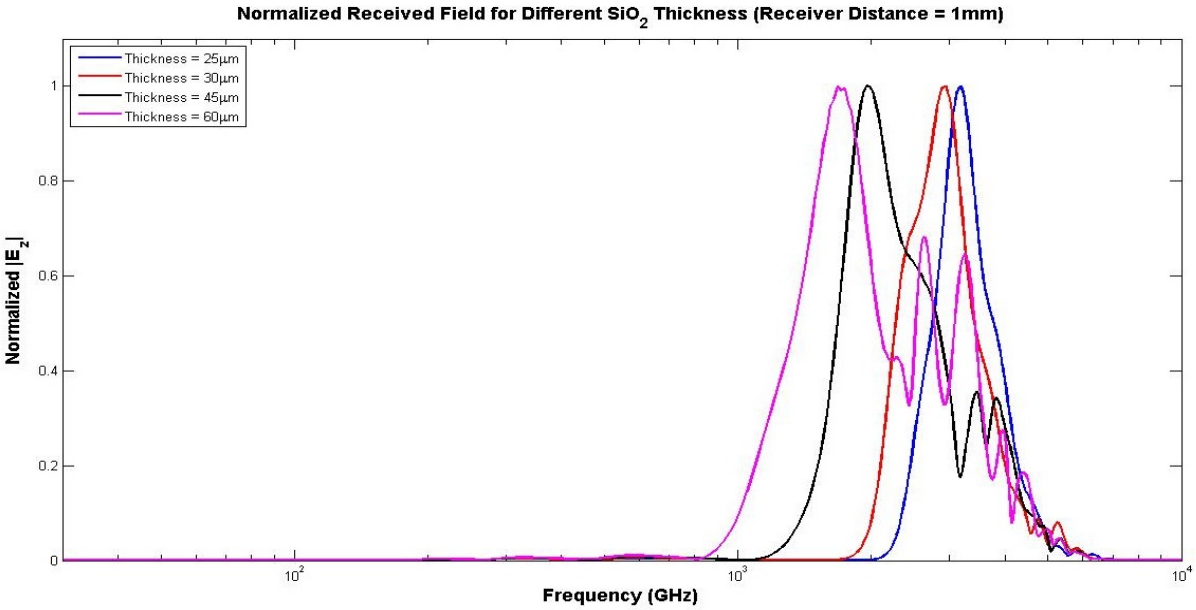

(b)

Figure 6.12: Variation of Normalized Electric Field with Frequency for Different SiO_2 Layer Thickness

for $2\mu m$ thickness case, as already discussed since the field falls at higher frequencies, the field is concentrated in the middle frequencies.

The variation of the received field with respect to SiO_2 layer thickness at different frequencies is shown in the Figure 6.13. The field strength increases initially with thickness but saturates beyond $25\mu m$ at all frequencies. At high frequencies, the field strength even starts decreasing beyond $40\mu m$ thickness. From the figures, any value between $20\mu m$ and $30\mu m$ for SiO_2 layer thickness would give similar results at most frequencies.

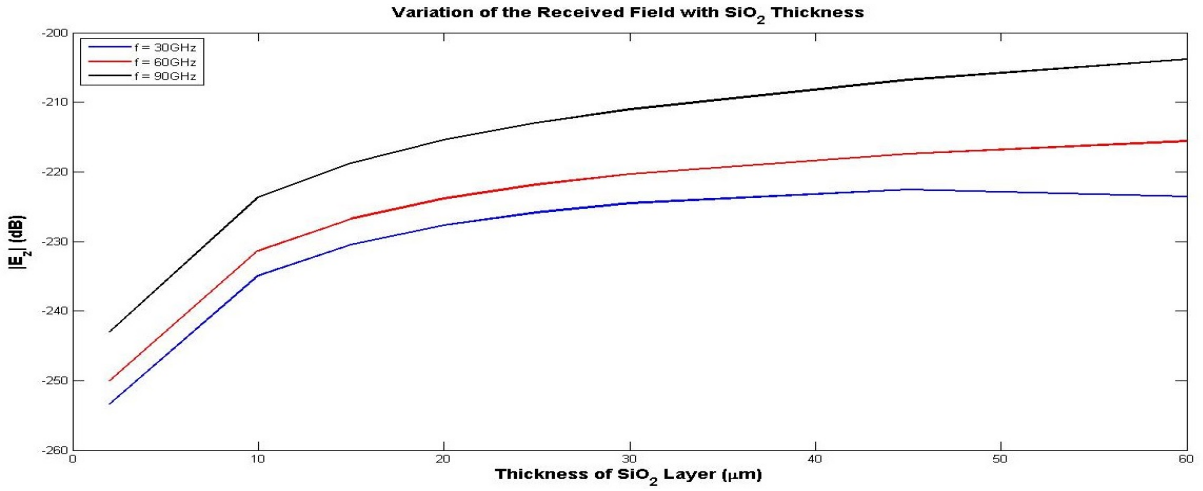

(a)

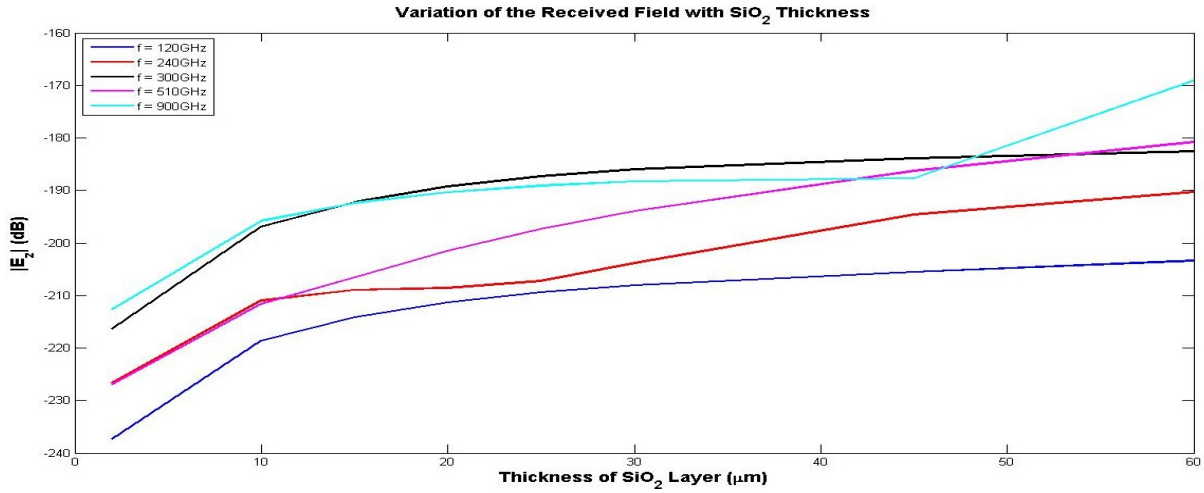

(b)

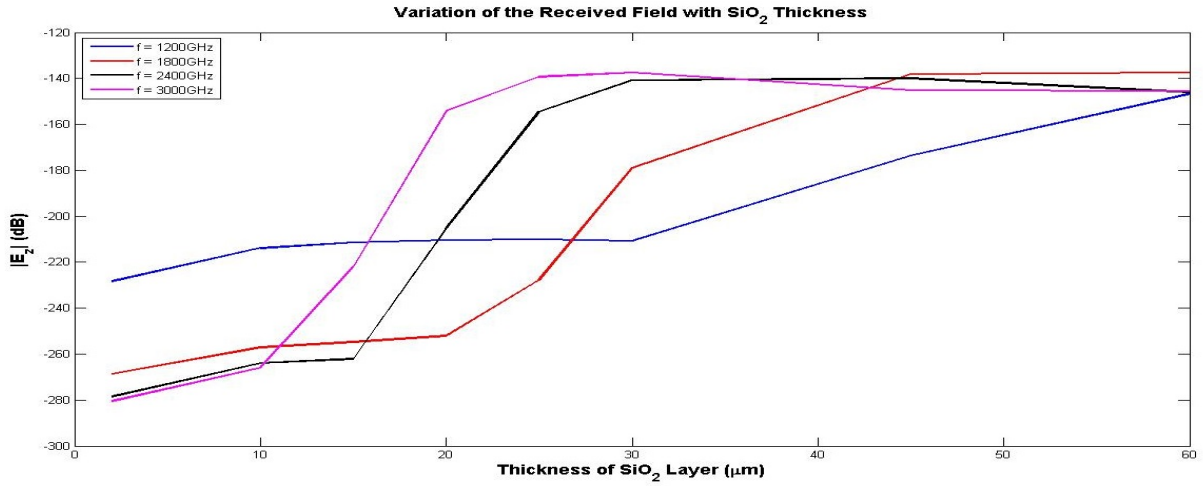

(c)

Figure 6.13: Variation of Received Electric Field with SiO_2 Layer Thickness

6.3.2 Variation with Distance

The variation of the received electric field with receiver distance for SiO_2 thickness of $30\mu m$ at different frequencies is shown in the Figure 6.14. The field strength falls with respect to the distance from the source.

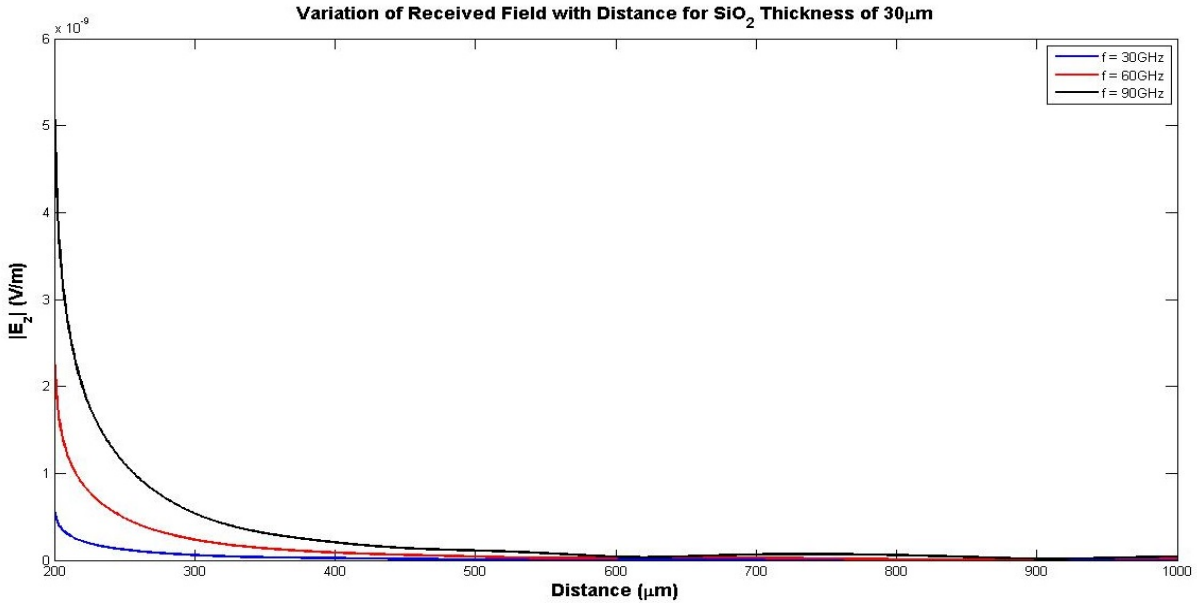

Figure 6.14: Variation of Received Electric Field with Receiver Distance

6.4 Delay

The primary advantage desired for using the wireless interfaces is their latency performance. Using the wireless interfaces, single hop delays can be achieved for data transmission in long distance communications on chip unlike wired interconnects which take multiple hops. But as discussed in the previous sections, the direct wave which provides the shortest delay in signal transmission is not the significant component of the received field. Hence, in this section, we investigate if the delay of the significant component is within the limits of single hop delay and still outperforms the wired interconnects.

Figure 6.15 shows the delay variation with respect to receiver distance. The signal delay increases almost linearly with distance and for a distance of 1 *mm*, the delay is 8.2567 *picosec*. Several previous papers provide a source-receiver separation of 5 *mm* [30] [48] [49]. Since the variation is linear, the delay for 5 *mm* distance can be 41.2835 *picosec*, which is more than double the delay assuming ideal wireless interfaces is around 17 *picosec*. But it still remains considerably less than the single hop delay (even assuming a 4 GHz operating frequency). The delay for 10 *mm* and 15 *mm* is 82.5670 *picosec* and 123.8505 *picosec* respectively. The delay

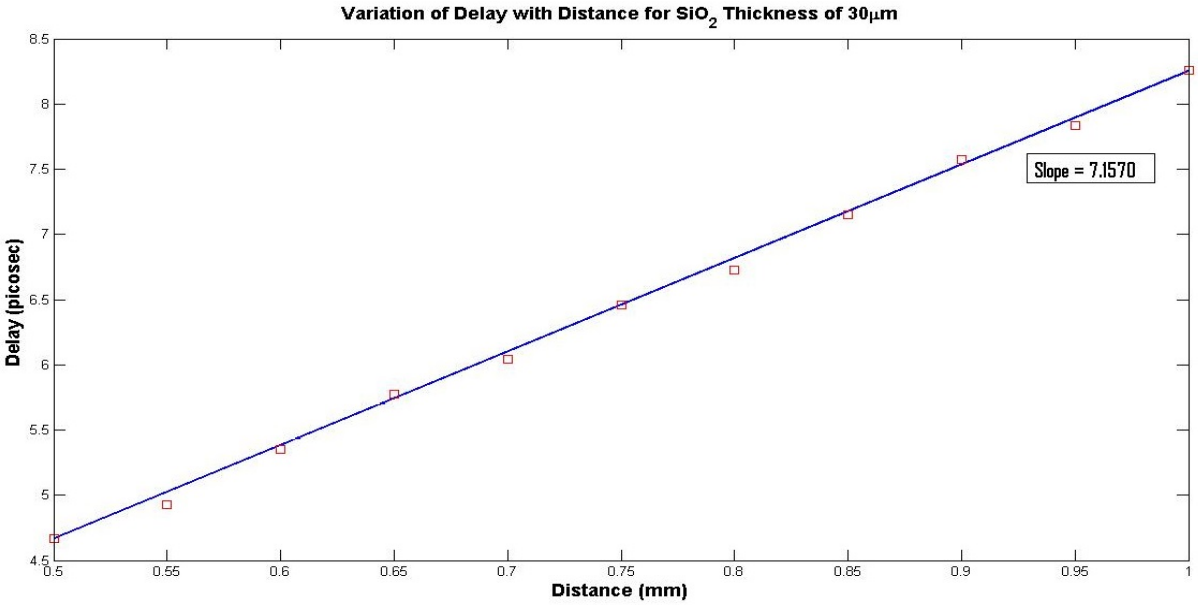

Figure 6.15: Variation of Signal Delay with Receiver Distance

for 15 *mm* case is more than the ideal delay of 50 *picosec* [34], but still provides better results compared to wired interconnects (0.9 *nanosec* [34]).

6.5 Computational Statistics

6.5.1 Memory Requirements

The FDTD method requires properties of the problem space and field data at all points in the problem space to calculate the field values at next time instant. MATLAB stores this data in two dimensional arrays for each required field variable, dielectric constant and relative permeability for the entire region. The memory required for storing these values is dependent on the problem size and the space resolution chosen. For a given problem space, the array sizes and subsequently the memory required varies as $O(\Delta x^{-2})$ with space resolution Δx. Hence as Δx decreases, the memory requirement increases quadratically. For our problem space, the memory required goes upto tens of gigabytes for scaled problem and hundreds of gigabytes for original problem space to save just the electric field data for the entire simulation period.

6.5.2 Execution Time

The time required to complete the execution of an FDTD simulation is depended on simulation duration, time resolution and space resolution. Since the field at each time step needs to be calculated at all points, the execution time varies quadratically with space resolution and linearly with time resolution. But the time resolution is directly proportional to space resolution as shown

in equation (5.9). Hence the execution time varies as $O(\Delta x^{-3})$ with Δx. Any change in space resolution results in corresponding cubic change in execution time. For the original problem space, with MATLAB running on a server with two Intel Xeon E5 chips and $64\,GB$ of RAM, the execution time for each each simulation is around 48 hours.

The space resolution chosen greatly affects both the memory requirements and execution time of a FDTD simulation. Therefore it is of paramount importance to choose the value properly.

Chapter 7

Conclusion and Future Work

In this work, two important problems pertaining to ULSI or SoC architectures, power consumption and propagation characteristics of wireless interfaces are tackled. Firstly to reduce the power consumption in the chip, a centralized controller that operates at the chip level is implemented. The controller applies a DVFS method to tackle the power in cores and power gating technique to the wireless interfaces in the system. The DVFS method sets the voltage and frequency for next time slice based on the core state and utilization for current and all past time slices. The power gating is applied to all the wireless interfaces on chip which are not in any active data transmission. The performance of the controller is evaluated using GEM5 and McPAT tools. A complete thermal profile of the system under different operating conditions is analyzed using HotSpot tool. The controller is then synthesized using Synopsys tools and to alleviate the issue of long delay in transmitting control signals, wireless interfaces are used for this purpose.

Secondly, the propagation characteristics of the wireless interfaces are studied and analyzed to see the impact of different interference structures on chip. Primarily, there are three major components in the received signal, direct wave, surface wave and reflections from metallic interconnects. It is observed that the direct wave through free space is completely canceled out by reflection from Air-SiO_2 interface and is not the dominant component of wave propagation. The reflections from underlying metallic interconnects contribute to the significant component to the received signal. are received much later than the free space component. The impact of receiver distance, passive layer thickness and material on signal strength and delay are studied. It is verified that even with the increased delay, the values still are considerably less and outperform the wired interconnect.

7.1 Future Work

The work presented in this thesis can be extended further in many possible directions. The application of DVFS can be extended to wireless interfaces to vary their operating voltage/frequency according to network traffic and utilization. The design and implementation of voltage regulators (VR) and frequency generators (FG) for centralized controllers can be explored. The pros and cons between using a single VR and FG or one VR and FG for each cluster can be analyzed and proper design choices can be made. One other possible direction of work is to look for other DVFS realizations that can be used in this scenario.

For channel propagation mechanisms, one most important direction of work can be to develop proper mathematical model for different modes in the signal based on problem space and materials used for chip structures. This helps to develop appropriate methods to maximize the gain of received signal.

As observed from results, the reflection from Air-SiO_2 interface cancels out the free space direct wave. So, the idea to use directional antenna to improve the strength of the direct wave can be explored. Theoretically, an antenna with a directional pattern over the upper half of antenna should improve the strength of direct wave. The STI guard bands used to separate devices on chip can be used for placing the antennas since these regions are filled with higher resistivity materials. We have used 2D model for the on-chip structures. To better model the structures accurately, a 3D on-chip model can be developed.

Bibliography

[1] Matlab the language of technical computing.

[2] Predictive technology model (ptm), 2012.

[3] Advanced configuration and power interface specification, 2013.

[4] Balanis, C. A. *Advanced Engineering Electromagnetics*. 1989.

[5] Berenger, J.-P. A perfectly matched layer for the absorption of electromagnetic waves. *Journal of Computational Physics 114*, 2 (1994), 185 – 200.

[6] Bienia, C. *Benchmarking Modern Multiprocessors*. PhD thesis, Princeton University, January 2011.

[7] Binkert, N., Beckmann, B., Black, G., Reinhardt, S. K., Saidi, A., Basu, A., Hestness, J., Hower, D. R., Krishna, T., Sardashti, S., Sen, R., Sewell, K., Shoaib, M., Vaish, N., Hill, M. D., and Wood, D. A. The gem5 simulator. *SIGARCH Comput. Archit. News 39*, 2 (Aug. 2011), 1–7.

[8] Bircher, W., and John, L. Core-level activity prediction for multicore power management. *Emerging and Selected Topics in Circuits and Systems, IEEE Journal on 1*, 3 (Sept 2011), 218–227.

[9] Bircher, W. L., and John, L. Predictive power management for multi-core processors. In *Proceedings of the 2010 International Conference on Computer Architecture* (Berlin, Heidelberg, 2012), ISCA'10, Springer-Verlag, pp. 243–255.

[10] Bohr, M. Interconnect scaling-the real limiter to high performance ulsi. In *Electron Devices Meeting, 1995. IEDM '95., International* (Dec 1995), pp. 241–244.

[11] Cheema, H., and Shamim, A. The last barrier: on-chip antennas. *Microwave Magazine, IEEE 14*, 1 (Jan 2013), 79–91.

[12] Chen, Z. M., and Zhang, Y.-P. Inter-chip wireless communication channel: Measurement, characterization, and modeling. *Antennas and Propagation, IEEE Transactions on 55*, 3 (March 2007), 978–986.

[13] Cho, T., Kang, D., Dow, S., Heng, C.-H., and Song, B. S. A 2.4ghz dual-mode 0.18/spl mu/m cmos transceiver for bluetooth and 802.11b. In *Solid-State Circuits Conference, 2003. Digest of Technical Papers. ISSCC. 2003 IEEE International* (Feb 2003), pp. 88–480 vol.1.

[14] Deb, S., Chang, K., Ganguly, A., Yu, X., Teuscher, C., Pande, P., Heo, D., and Belzer, B. Design of an efficient noc architecture using millimeter-wave wireless links. In *Quality Electronic Design (ISQED), 2012 13th International Symposium on* (March 2012), pp. 165–172.

[15] Deb, S., Chang, K., Yu, X., Sah, S., Cosic, M., Ganguly, A., Pande, P., Belzer, B., and Heo, D. Design of an energy-efficient cmos-compatible noc architecture with millimeter-wave wireless interconnects. *Computers, IEEE Transactions on 62*, 12 (Dec 2013), 2382–2396.

[16] Deb, S., Ganguly, A., Chang, K., Pande, P., Beizer, B., and Heo, D. Enhancing performance of network-on-chip architectures with millimeter-wave wireless interconnects. In *Application-specific Systems Architectures and Processors (ASAP), 2010 21st IEEE International Conference on* (July 2010), pp. 73–80.

[17] Diao, Q., and Song, J. Prediction of cpu idle-busy activity pattern. In *High Performance Computer Architecture, 2008. HPCA 2008. IEEE 14th International Symposium on* (Feb 2008), pp. 27–36.

[18] Faust, G., Zhang, R., Skadron, K., Stan, M., and Meyer, B. Archfp: Rapid prototyping of pre-rtl floorplans. In *VLSI and System-on-Chip (VLSI-SoC), 2012 IEEE/IFIP 20th International Conference on* (Oct 2012), pp. 183–188.

[19] Ganguly, A., Chang, K., Deb, S., Pande, P., Belzer, B., and Teuscher, C. Scalable hybrid wireless network-on-chip architectures for multicore systems. *Computers, IEEE Transactions on 60*, 10 (Oct 2011), 1485–1502.

[20] Garg, S., Marculescu, D., Marculescu, R., and Ogras, U. Technology-driven limits on dvfs controllability of multiple voltage-frequency island designs: A system-level perspective. In *Design Automation Conference, 2009. DAC '09. 46th ACM/IEEE* (July 2009), pp. 818–821.

[21] Gebhart, M., Hestness, J., Fatehi, E., Gratz, P., and Keckler, S. W. Running parsec 2.1 on m5. Tech. rep., The University of Texas at Austin, Department of Computer Science, October 2009.

[22] Guo, X., Li, R., and O, K. Design guidelines for reducing the impact of metal interference structures on the performance on-chip antennas. In *Antennas and Propagation Society International Symposium, 2003. IEEE* (June 2003), vol. 1, pp. 606–609 vol.1.

[23] Herbert, S., and Marculescu, D. Analysis of dynamic voltage/frequency scaling in chip-multiprocessors. In *Low Power Electronics and Design (ISLPED), 2007 ACM/IEEE International Symposium on* (Aug 2007), pp. 38–43.

[24] Ho, R., Ono, I., Liu, F., Hopkins, R., Chow, A., Schauer, J., and Drost, R. High-speed and low-energy capacitively-driven on-chip wires. In *Solid-State Circuits Conference, 2007. ISSCC 2007. Digest of Technical Papers. IEEE International* (Feb 2007), pp. 412–612.

[25] Hoskote, Y., Vangal, S., Singh, A., Borkar, N., and Borkar, S. A 5-ghz mesh interconnect for a teraflops processor. *Micro, IEEE 27*, 5 (Sept 2007), 51–61.

[26] Huang, W., Sankaranarayanan, K., Skadron, K., Ribando, R., and Stan, M. Accurate, pre-rtl temperature-aware design using a parameterized, geometric thermal model. *Computers, IEEE Transactions on 57*, 9 (Sept 2008), 1277–1288.

[27] Jung, Y.-J., Jeong, H., Song, E., Lee, J., Lee, S.-W., Seo, D., Song, I., Jung, S., Park, J., Jeong, D.-K., Chae, S.-I., and Kim, W. A 2.4-ghz 0.25- mu;m cmos dual-mode direct-conversion transceiver for bluetooth and 802.11b. *Solid-State Circuits, IEEE Journal of 39*, 7 (July 2004), 1185–1190.

[28] Kikkawa, T., Rashid, A. B. M. H., and Watanabe, S. Effect of silicon substrate on the transmission characteristics of integrated antenna. In *Wireless Communication Technology, 2003. IEEE Topical Conference on* (Oct 2003), pp. 144–145.

[29] Kim, K., Bomstad, W., and Kenneth, K. A plane wave model approach to understanding propagation in an intra-chip communication system. In *Antennas and Propagation Society International Symposium, 2001. IEEE* (July 2001), vol. 2, pp. 166–169 vol.2.

[30] Kim, K., and O, K. Characteristics of integrated dipole antennas on bulk, soi, and sos substrates for wireless communication. In *Interconnect Technology Conference, 1998. Proceedings of the IEEE 1998 International* (Jun 1998), pp. 21–23.

[31] Kolpe, T., Zhai, A., and Sapatnekar, S. Enabling improved power management in multicore processors through clustered dvfs. In *Design, Automation Test in Europe Conference Exhibition (DATE), 2011* (March 2011), pp. 1–6.

[32] Krishna, T., Kumar, A., Chiang, P., Erez, M., and Peh, L.-S. Noc with near-ideal express virtual channels using global-line communication. In *High Performance Interconnects, 2008. HOTI '08. 16th IEEE Symposium on* (Aug 2008), pp. 11–20.

[33] Li, S., Ahn, J.-H., Strong, R., Brockman, J., Tullsen, D., and Jouppi, N. Mcpat: An integrated power, area, and timing modeling framework for multicore and manycore architectures. In *Microarchitecture, 2009. MICRO-42. 42nd Annual IEEE/ACM International Symposium on* (Dec 2009), pp. 469–480.

[34] Mondal, H., Harsha, G., and Deb, S. An efficient hardware implementation of dvfs in multi-core system with wireless network-on-chip. In *VLSI (ISVLSI), 2014 IEEE Computer Society Annual Symposium on* (July 2014).

[35] Murray, J., Hegde, R., Lu, T., Pande, P., and Shirazi, B. Sustainable dual-level dvfs-enabled noc with on-chip wireless links. In *Quality Electronic Design (ISQED), 2013 14th International Symposium on* (March 2013), pp. 135–142.

[36] Murray, J., Lu, T., Pande, P., and Shirazi, B. Chapter 3 - sustainable dvfs-enabled multi-core architectures with on-chip wireless links. In *Green and Sustainable Computing: Part II*, A. Hurson, Ed., vol. 88 of *Advances in Computers*. Elsevier, 2013, pp. 125 – 158.

[37] Nowak, E. Maintaining the benefits of cmos scaling when scaling bogs down. *IBM Journal of Research and Development 46*, 2.3 (March 2002), 169–180.

[38] Park, S., Park, J., Shin, D., Wang, Y., Xie, Q., Pedram, M., and Chang, N. Accurate modeling of the delay and energy overhead of dynamic voltage and frequency scaling in modern microprocessors. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on 32*, 5 (May 2013), 695–708.

[39] Razavi, B. A 900-mhz/1.8-ghz cmos transmitter for dual-band applications. *Solid-State Circuits, IEEE Journal of 34*, 5 (May 1999), 573–579.

[40] Samaiyar, A., Deb, S., and Ram, S. Millimeter-wave planar log periodic antenna for on-chip wireless interconnects. In *Antennas and Propagation (EuCAP), 2014 8th European Conference on* (April 2014).

[41] Schneider, J. B. Understanding the finite-difference time-domain method, 2010.

[42] Seok, E., and Kenneth, K. Design rules for improving predictability of on-chip antenna characteristics in the presence of other metal structures. In *Interconnect Technology Conference, 2005. Proceedings of the IEEE 2005 International* (June 2005), pp. 120–122.

[43] Sullivan, D. M. *Electromagnetic Simulation Using the FDTD Method*. 2000.

[44] Yan, L., and Hanson, G. Wave propagation mechanisms for intra-chip communications. *Antennas and Propagation, IEEE Transactions on 57*, 9 (Sept 2009), 2715–2724.

[45] Ye, R., and Xu, Q. Learning-based power management for multi-core processors via idle period manipulation. In *Design Automation Conference (ASP-DAC), 2012 17th Asia and South Pacific* (Jan 2012), pp. 115–120.

[46] Yee, K., and Chen, J. The finite-difference time-domain (fdtd) and the finite-volume time-domain (fvtd) methods in solving maxwell's equations. *Antennas and Propagation, IEEE Transactions on 45*, 3 (Mar 1997), 354–363.

[47] Zhang, Y., Hu, X., and Chen, D. Task scheduling and voltage selection for energy minimization. In *Design Automation Conference, 2002. Proceedings. 39th* (2002), pp. 183–188.

[48] Zhang, Y. P., Chen, Z. M., and Sun, M. Propagation mechanisms of radio waves over intra-chip channels with integrated antennas: Frequency-domain measurements and time-domain analysis. *Antennas and Propagation, IEEE Transactions on 55*, 10 (Oct 2007), 2900–2906.

[49] Zhang, Y. P., Sun, M., and Fan, W. Performance of integrated antennas on silicon substrates of high and low resistivities up to 110 ghz for wireless interconnects. *Microwave and Optical Technology Letters 48*, 2 (2006), 302–305.