

Modelling and Analysis of Interconnects for Deep Submicron Systems-on-Chip

Dinesh Pamunuwa



**ROYAL INSTITUTE
OF TECHNOLOGY**

Stockholm 2003

Laboratory of Electronics and Computer Systems
Department of Microelectronics and Information Technology
Royal Institute of Technology, Stockholm, Sweden

*Thesis submitted to the Royal Institute of Technology in partial
fulfilment of the requirements for the degree of Doctor of Technology*

Pamunuwa, Dinesh
Modelling and Analysis of Interconnects for Deep Submicron Systems-on-Chip

ISBN 91-7283-631-8
ISRN KTH/IMIT/LECS/AVH-03/07--SE
ISSN 1651-4076
TRITA-IMIT-LECS AVH 03:07

© Dinesh Pamunuwa 2003

Royal Institute of Technology
Department of Microelectronics and Information Technology
Laboratory of Electronics and Computer Systems
Isafjordsgatan 39
SE 16440, Stockholm-Kista, Sweden

And near me on the grass lies Glanvil's book-
Come, let me read the oft-read tale again!
The story of the Oxford scholar poor,
Of pregnant parts and quick inventive brain,
Who, tired of knocking at preferment's door,
One summer-morn forsook
His friends, and went to learn the gipsy-lore,
And roam'd the world with that wild brotherhood,
And came, as most men deem'd, to little good,
But came to Oxford and his friends no more.

- from "Scholar Gipsy" by Mathew Arnold

Abstract

The last few decades have been a very exciting period in the development of microelectronics and brought us to the brink of implementing entire systems on a single chip, on a hitherto unimagined scale. However an unforeseen challenge has cropped up in the form of managing wires, which have become the main bottleneck in performance, masking the blinding speed of active devices. A major problem is that increasingly complicated effects need to be modelled, but the computational complexity of any proposed model needs to be low enough to allow many iterations in a design cycle.

This thesis addresses the issue of closed-form modelling of the response of coupled interconnect systems. Following a strict mathematical approach, second order models for the transfer functions of coupled *RC* trees based on the first and second moments of the impulse response are developed. The 2-pole-1-zero transfer function that is the best possible from the available information is obtained for the signal path from each driver to the output in multiple-aggressor systems. This allows the complete response to be estimated accurately by summing up the individual waveforms. The model represents the minimum complexity for a 2-pole-1-zero estimate, for this class of circuits.

Also proposed are new techniques for the optimisation of wires in on-chip buses. Rather than minimising the delay over each individual wire, the configuration that maximises the total bandwidth over a number of parallel wires is investigated. It is shown from simulations that there is a unique optimal solution which does not necessarily translate to the maximum possible number of wires, and in fact deviates considerably from it when the resources available for repeaters are limited. Analytic guidelines dependent only on process parameters are derived for optimal sizing of wires and repeaters.

Finally regular tiled architectures with a common communication backplane are being proposed as being the most efficient way to implement systems-on-chip in the deep submicron regime. This thesis also considers the feasibility of implementing a regular packet-switched network-on-chip in a typical future deep submicron technology. All major physical issues and challenges are discussed for two different architectures and important limitations are identified.

Keywords: delay and noise modelling in VLSI circuits, cross-talk, interconnect modelling, timing analysis, transfer function, on-chip bus, bandwidth maximization, throughput maximization, high-speed interconnect, interconnect delay, repeater insertion, wire optimization, ULSI, high-performance, system-level communication

Acknowledgements

First and foremost, I would like to thank my supervisor, Prof. Hannu Tenhunen for his guidance and support throughout my PhD studies. He has that rare gift of being both profound and lucid, and under his tutelage a project that seemed initially to be shadowy and wraithlike, took form and substance, and led to this thesis. I am also much indebted to him for his many acts of kindness, and for backing me up in numerous situations.

I owe a huge debt of thanks to Shauki Elassaad of Cadence Berkeley Laboratories in Berkeley, CA, for his excellent supervision while I was there, and for having faith in my ability to solve the problems that he laid at my door. Subsequent productive discussions are also gratefully acknowledged. I am also indebted to Dr. Li-Rong Zheng of LECS for his valuable contributions in countless discussions. In spite of being harried to the point of distraction by numerous project deadlines, he always took time to discuss any issue I brought before him, and came up with very good practical suggestions each time.

My colleagues in the lab, Andreas, Abhijit, Steffen, Ingo, Raimo, Wim, Micke, Li Li and others helped in no small way through discussions and many acts of friendship. A special word of thanks is due to Andreas, for his boundless enthusiasm which pushed me to match his own.

I am thankful to Prof. Axel Jantsch and Dr. Johnny Öberg of LECS for giving me the opportunity to work on the NoC project, and the many interesting and stimulating conversations I have had with them.

I am also very grateful to Prof. Eby Friedman, of the University of Rochester -whose erudition and humanism I admire and respect in equal measure- for being kind enough to share his expertise. A special word of thanks as well to Prof. Jari Nurmi of TUT for his hospitality in Tampere, and for organising many productive discussions.

My thanks are due to Cathy Larimer for helping me out with complicated administrative tasks very efficiently in Berkeley and for being instrumental in expediting my internship at Cadence Berkeley Laboratories. I am very grateful to Hans Bergren for his help with IT support at KTH. I also recall with appreciation the initial help I received from Costantino and Gerd in navigating the complicated paths of Cadence software.

This work would not have been possible without the funding support of Sida, and the Government of Sweden. Sida has been responsible for many successful projects in

Sri Lanka, and I would like to put on record my gratitude both as an individual and as a citizen of Sri Lanka. My thanks are also due to all those who initiated this particular project, Dr. Nimal Ratnayake and Dr. Sanath Alahakoon among others in Sri Lanka, and Prof. Roland Eriksson in Sweden. I am very grateful to Cadence Design Systems Inc., USA, for funding my industrial internship and giving me the opportunity to interact with many industrial experts.

My time in Stockholm was made much more pleasant because of many wonderful friends I made, and I will always count myself fortunate for having met Frank and Harendra and others too numerous to mention here.

Finally I would like to acknowledge my biggest debt, to my parents and my brother, for their constant support.

List of Publications

- [1] D. Pamunuwa, S. Elassaad and H. Tenhunen, “Modelling noise and delay in VLSI circuits”, *Electronics Letters*, Vol. 39 Issue 3, pp. 269-271, Feb. 2003.
- [2] D. Pamunuwa, L. R. Zheng and H. Tenhunen, “Maximizing Throughput over Parallel Wire Structures in the Deep Submicrometer Regime”, *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 11, no. 2, pp. 224-243, April, 2003.
- [3] D. Pamunuwa, S. Elassaad and H. Tenhunen, “Modelling delay and noise in arbitrarily-coupled RC trees”, under review, *IEEE Transactions on Computer-Aided Design of Circuits and Systems*, Nov. 2003.
- [4] D. Pamunuwa, J. Öberg, L. R. Zheng, M. Millberg, A Jantsch and H. Tenhunen, “A study on the implementation of 2-D mesh-based networks-on-chip in the nanometre regime,” under review, *Integration - the VLSI Journal, Special Issue on Networks-on-Chip and Reconfigurable Fabrics, Elsevier*, Aug. 2003.
- [5] J. Öberg, D. Pamunuwa, L. R. Zheng, M. Millberg, A Jantsch and H. Tenhunen, “A feasibility study on the performance and power distribution of two possible Network-on-Chip architectures,” under review, *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, Oct. 2002.
- [6] D. Pamunuwa, S. Elassaad and H. Tenhunen, “Analytic Modeling of Interconnects for Deep Sub-Micron Circuits”, in *Proc. International Conference on Computer-Aided Design (ICCAD 2003)*, (in press), Nov. 2003.
- [7] D. Pamunuwa and S. Elassaad, “Closed Form Metrics to Accurately Model the Response in Arbitrarily-coupled RC Trees”, in *Proc. IEEE International Symposium on Circuits and Systems (ISCAS 2003)*, Bangkok, Thailand, May 2003, vol. 4, pp. 892-5.
- [8] D. Pamunuwa, J. Öberg, L. R. Zheng, M. Millberg, A Jantsch and H. Tenhunen, “Layout, performance and power trade-offs for in mesh-based network-on-chip architectures,” in *Proc. IFIP International Conference on VLSI Systems-on-Chip*, Darmstadt, Germany, Dec. 2003 (in press)

-
- [9] D. Pamunuwa, L. R. Zheng and H. Tenhunen, "Optimising Bandwidth Over Deep Sub-micron Interconnect", in *Proc. IEEE International Symposium on Circuits and Systems (ISCAS 2002)*, Scottsdale, Arizona, USA, May 2002, vol. 4, pp. 193-196.
- [10] H. Tenhunen and D. Pamunuwa, "On Dynamic Delay and Repeater Insertion", in *Proc. IEEE International Symposium on Circuits and Systems (ISCAS 2002)*, Scottsdale, Arizona, USA, May 2002, vol. 1, pp. 97-100.
- [11] D. Pamunuwa and H. Tenhunen, "On Dynamic Delay and Repeater Insertion in Distributed Capacitively Coupled Interconnects", in *Proc. IEEE International Symposium on Quality Electronic Design (ISQED 2002)*, San Jose, California, USA, March 2002, pp. 240-245.
- [12] D. Pamunuwa and H. Tenhunen, "Repeater Insertion to Minimise Delay in Parallel Coupled Interconnects", in *Proc. International Conference on VLSI Design (VLSI Design 2001)*, Bangalore, India, January 2001, pp. 513-517.
- [13] L. R. Zheng, D. Pamunuwa and H. Tenhunen, "Accurate a priori signal integrity estimation using a multilevel dynamic interconnect model for deep submicron VLSI design", in *Proc. of the 2000 European Solid State Circuits Conference (ESSCIRC 2000)*, Stockholm, Sweden, Sept. 2000, pp. 324-327.
- [14] D. Pamunuwa, L. R. Zheng and H. Tenhunen, "Combating Digital Noise in High Speed ULSI Circuits Using Binary BCH Encoding", in *Proc. IEEE International Symposium on Circuits and Systems (ISCAS 2000)*, Geneva, Switzerland, May 2000, vol. 4, pp. 13-16.
- [15] D. Pamunuwa, L. R. Zheng and H. Tenhunen, "Error-Control Coding to Combat Digital Noise in Interconnects for ULSI Circuits", in *Proc. Norchip Conference (NORCHIP 1999)*, Oslo, Norway, Nov. 1999, pp. 275-282.
- [16] J. Liu, D. Pamunuwa, L. R. Zheng, and H. Tenhunen, "A global wire planning scheme for network-on-chip," in *Proc. IEEE International Symposium on Circuits and Systems (ISCAS 2003)*, Bangkok, Thailand, May 2003, pp. 892-895.

CONTENTS

Abstract	v
Acknowledgements	vi
List of Publications	viii
1. Introduction	1
1.1 Developments in Microelectronics	1
1.1.1 Technology Trends and Moore's Law	2
1.1.2 Scaling of Wires	4
1.2 Coping with Deep submicron Effects	5
1.2.1 System-on-Chip Communication Woes	5
1.2.2 Off-Chip Communication	7
1.2.3 Dealing with Complexity	7
1.3 Scope of Thesis and Author's Contribution	8
1.3.1 Second Order Modelling of Arbitrarily-Coupled RC Trees	9
1.3.2 Bandwidth Optimisation of On-Chip Buses	10
1.3.3 Physical Issues in Implementation of Networks-On-Chip	11
1.4 Thesis Organisation	12
Chapter 2: Modelling of Interconnects at High Frequencies	13
Chapter 3: Delay and Noise Analysis of Interconnect	13
Chapter 4: Repeater Modelling and Insertion for Coupled Nets ..	13
Chapter 5: Optimal Signalling for On-Chip Buses	13
Chapter 6: Designing System-on-Chip Communication Networks	13
Chapter 7: Summary and Conclusions	13
2. Modelling of Interconnect at High Frequencies	15
2.1 Field Solvers	16
2.1.1 Finite Element Method (FEM)	16
2.1.2 Moment Method	17
2.1.3 Boundary Element Method (BEM)	17
2.1.4 Finite Difference Time Domain (FDTD) Method	17
2.1.5 Finite Difference Frequency Domain (FDFD) Method	18
2.1.6 Transmission Line Matrix (TLM) Method	18
2.1.7 Partial-Element Equivalent Circuit (PEEC) Method	18
2.2 Analytic Formulae	19
2.2.1 Resistance	19
DC Resistance	19
Resistance with Skin Effect	20
2.2.2 Capacitance	21
Single Conductor over Ground Plane	21
Coupled Wires over Ground Plane	23

2.2.3	Inductance	26
2.2.4	Inductance vs Capacitance Extraction	29
2.3	Electrical Level Modelling	30
2.3.1	The General Transmission Line	30
	Modelling the General Transmission Line	32
2.3.2	Simplified Transmission Line Models	32
	Lumped Single Element Models	32
	RC Transmission Line	33
2.4	Choosing a Wire Model.	33
2.5	Summary	34
3.	Delay and Noise Analysis of Interconnect	37
3.1	Introduction	37
3.2	Background	38
3.2.1	Delay Modelling	38
3.2.2	Noise Modelling	41
3.3	Modelling the System Transfer Function	42
3.3.1	Response to Different Switching Events	44
3.4	Calculation of Moments	44
3.4.1	Notation	45
3.4.2	Switching of Victim Driver	46
3.4.3	Switching of Aggressor Driver	49
3.5	Matching Moments to Characteristic Time Constants in Circuit	50
3.5.1	Guaranteeing Stability	51
3.5.2	Switching of Victim Driver	55
3.5.3	Switching of an Aggressor Driver	56
3.6	Physical Basis of the Model	58
3.7	Computational Complexity	59
3.7.1	Background: Incremental Computation of the Elmore Delay	59
3.7.2	Computational Complexity of Proposed Metrics	61
	First order metrics	61
	Second order metrics	61
	Summary	65
3.8	Explicit Noise Models	65
3.9	Results	66
3.10	Summary	76
3.11	Limitations and Future Work.	77
4.	Repeater Modelling	79
4.1	Introduction	79
4.1.1	Background	79
4.2	Signal Delay in Long Uniformly Coupled Nets	82
4.3	Repeater Insertion	86
4.3.1	Minimum Delay	91

4.4	Model Verification	92
4.4.1	Aggressor Alignment	92
4.4.2	Testing with Real Repeaters	94
4.5	Estimating Device Characteristics	96
4.6	Summary	97
5.	Optimal Signalling Over On-chip Buses	99
5.1	Introduction	99
5.2	Interconnect Modelling and Delay Analysis	101
5.2.1	Parasitic Modelling	101
5.2.2	Line Delay and Repeater Insertion	103
5.3	Optimal Signalling Over Parallel Wires	103
5.3.1	Fixed Wire-Width and -Pitch	105
5.3.2	Variable Wire-Width and -Pitch	107
	Simulations	107
	Validity of Analysis	111
	Analytic Guidelines	111
5.4	Error-Control Coding for Lossy Lines	116
5.4.1	Noise Analysis and Modelling	117
5.4.2	Boundary Conditions	117
5.4.3	Genesis of Binary BCH Codes	119
5.4.4	Coding Gain	121
5.5	Summary and Conclusions	122
6.	Designing SoC Communication Networks	125
6.1	Introduction	125
6.1.1	Background	125
6.1.2	Feasibility Study	126
6.2	NoC backbone	127
6.2.1	Architecture	127
6.2.2	Network Protocol	128
6.3	Modelling issues	130
6.3.1	Technology Scaling	130
6.3.2	Switches and Inter-Switch Links	131
	Square Switch	131
	Thin Switch	131
	Network Links	132
6.3.3	Resources	133
6.3.4	Power Estimations	133
6.4	Analysis and results	134
6.4.1	Square-Switch Architecture	134
6.4.2	Thin-switch Architecture	136
6.5	Discussion and Conclusions	138

7. Conclusions	143
7.1 Summary and Conclusions	143
7.2 Limitations and Future Work	144
8. Bibliography	147

1. Introduction

This chapter discusses the motivation for the work, provides a summary of the technical contributions made by the author, and outlines the structure of the thesis.

1.1 Developments in Microelectronics

The digital revolution that started in the 1960s has touched almost every sphere of our lives and it is hard to find any field that has not benefited from digital electronics. The oft-cited example is the field of information technology, and the miraculous information super-highway that is the world-wide-web. Additionally innumerable other advances could be mentioned, in telephony, food and medical technology, transport, construction and production, publishing, the entertainment industry, and the host of improvements that has resulted from computerisation in fields such as accountancy, stock management and record maintenance.

This success story has depended not only on material innovations at the device level, but also on the ability to deal with the complexity in the design process itself. This has been achieved by design automation, built on the two cornerstones underlying all engineering solutions to large and unwieldy problems, namely partitioning and hierarchy. Partitioning a design is basically breaking it up into several smaller tasks, so that each individual problem represents a lesser challenge, and the overall design is achieved by fitting together the completed pieces. Hierarchical design is the practice of building up a module from sub-modules, each of which in turn will compose of their own sub-modules, a process which is continued until the basic building blocks for the design are reached. Key to both these design principles is the ability to abstract many electrical level details at different levels of the design hierarchy.

However with device technology continuing to advance at a breakneck speed, this abstraction is being threatened, and new issues have come up that require careful treatment and analysis, and challenge the entire automation process. Higher signal frequencies mean that wires can no longer be treated as equipotential regions, and often prove the bottleneck to blindingly fast gates; increased integration, falling voltages and smaller rise times make noise issues very important; clock and power distribution have to be carried out over increasingly longer distances with tighter tolerances. Some of these issues are examined in detail in the following sections. It is illuminating to start the discussion with a historical time-line of important developments.

1.1.1 Technology Trends and Moore's Law

The first transistor was invented at Bell Laboratories in December, 1947 by John Bardeen, Walter Brattain and William Shockley¹ [Bardeen48]. This was, with hindsight, perhaps the most important electronics event of the 20th century, as it later made possible the integrated circuit and microprocessor that are the basis of modern electronics. Prior to the transistor (TRANSfer resISTOR) the only alternative to its current regulation and switching functions was the vacuum tube, which could only be miniaturized to a certain extent, and wasted a lot of energy in the form of heat.

The picture in Figure 1.1 shows the first point contact transistor built by Walter Brattain. It consisted of a plastic triangle lightly suspended above a germanium crystal which itself was sitting on a metal plate attached to a voltage source. A strip of gold was wrapped around the point of the triangle with a tiny gap cut into the gold at the precise point it came in contact with the germanium crystal. The germanium acted as a semiconductor so that a small electric current entering on one side of the gold strip came out the other side as a proportionately amplified current.




Figure 1.1: The first point contact transistor and its co-inventors William Shockley (seated), John Bardeen (left) and Walter Brattain (right)²

1. As far back as 1926, Dr. Julius Edgar Lilienfeld from New York filed for a patent on what we would now recognize as an NPN junction transistor being used in the role of an amplifier. However the properties of semiconductors were not well understood at the time [Lee98].
2. Source: Archives at Lucent Technologies website of Bell Labs innovations, protected by copyright law as given in <http://www.lucent.com/copyright.html>

In 1950, Shockley invented a new device called a bipolar junction transistor, which was more reliable, easier and cheaper to build, and gave more consistent results than point-contact devices [Schockley48], [Schockley49]. In 1962, Steven Hofstein and Fredric Heiman at the RCA research laboratory in Princeton, New Jersey, built the first stable practical metal-oxide semiconductor field-effect transistor (MOSFET)¹. In 1959, the Swiss physicist Jean Hoerni invented the planar process, in which optical lithographic techniques were used to diffuse the base into the collector and then diffuse the emitter into the base. One of Hoerni's colleagues, Robert Noyce, invented a technique for growing an insulating layer of silicon dioxide over the transistor, leaving small areas over the base and emitter exposed and diffusing thin layers of aluminum into these areas to create wires. The processes developed by Hoerni and Noyce led directly to modern integrated circuits (ICs). It was Jack Kilby working for Texas Instruments however, who first succeeded in fabricating multiple components on a single piece of semiconductor in the summer of 1958. Kilby's first prototype was a phase shift oscillator, and although manufacturing techniques subsequently took different paths to those used by Kilby, he is credited with the creation of the first true integrated circuit.

In 1965, while working at Fairchild Semiconductor Industries, Gordon Moore (later co-founder of Intel) made the remarkably visionary prediction that the number of transistors that could be integrated on a single die would grow exponentially with time² [Moore65]. This has held true since then, and is expected to hold true at least until 2010, with an annual increase in integration density of roughly 50% [Schaller97]. Now along with the increase in device density, to a first order, the delay of a simple gate has been decreasing linearly with gate length, at approximately 13% a year [Dally98]. This means that the speed of functions has also been increasing exponentially with time. Since the capability of a chip depends on both the number of functions on it and the speed of those functions, the combination of the 13% growth in speed and the 50% growth in device density results in a 70% annual increase in the overall capability of a chip. Unfortunately, the potential of a full-sized chip with gates at the maximum possible speed is unlikely to be realized because of the sheer complexity of managing the inter-connectivity of billions of devices. A key problem is the increasing wiring delay as shown in the next few sections, and provides the motivation for this thesis.

-
1. The physics of the device was worked out by Julius Lilienfeld as early as 1925, and also independently by O. Heill in 1935 [Rabaey96]. Insufficient knowledge of materials and problems with gate stability however defeated all attempts to build a working device until the success of Hofstein and Heiman.
 2. So extraordinarily accurate was this prediction that it became known as *Moore's Law* and has always been the single most verbally cited reference in all conferences on microelectronics that this author has attended.

1.1.2 Scaling of Wires

The importance of wires in terms of delay, power and density increases in comparison with devices as technology scales. Consider Table 1.1 which is reproduced from [Dally98].

It shows three pairs of columns which describe the scaling formula and the change per year for a device sized wire, a 1 μm wire and a global wire (one that traverses the length of the chip). The term x refers to the annual scaling factor for gate length while y refers to that for chip edge.

All the parameters given in the table are important in various ways for delay and power metrics. To a first order, the parasitic capacitance decreases over time for a device sized wire, while it remains constant for a 1 μm wire and increases for a global wire. The resistance increases for all types of wires, and at a significantly higher rate for global wires. The ratio of the IR drop over the wire to the rail voltage also increases for all types of wires, and again, much more significantly for global wires. One of the most important metrics is the RC product for a wire, which increases at a rate of 50% per year for global wires. The ramifications of this parasitic scaling will be examined in the next section.




Figure 1.2: Growth of integration density for microprocessors and memories (reproduced from [Zheng01])

Table 1.1 Scaling of Wire Properties

Parameter	Device		1- μm		Chip	
	C	x	0.87	1	1.00	y
R	$1/x$	1.15	$1/x^2$	1.32	y/x^2	1.40
I	x	0.87	1	1.00	y	1.06
IR	1	1.00	$1/x^2$	1.32	$(y/x)^2$	1.49
IR/V	$1/x$	1.15	$1/x^3$	1.51	y^2/x^3	1.71
RC	1	1.00	$1/x^2$	1.32	$(y/x)^2$	1.49
RC/τ	$1/x$	1.15	$1/x^3$	1.51	y^2/x^3	1.71

1.2 Coping with Deep submicron Effects

1.2.1 System-on-Chip Communication Woes

As feature sizes shrink into the so called *deep* submicron (DSM) regime consisting of lengths of around a hundred nano metres and less, the potential exists to pack together hundreds of millions of transistors on a single die. Since it is easier to make use of the vast number of logic gates available by a modular approach, a chip will in most likelihood consist of an interconnected set of resources with very different functionalities and design and implementation styles (examples being processor cores, DSP cores, FPGA blocks, dedicated hardware blocks, mixed signal blocks, analog and/or RF blocks, and different memory blocks such as RAM, ROM and CAM), comprising an entire system that would typically be implemented over several chips in the submicron regime. The collection of different types of resources on a single chip has led to the coining of the term *System-on-Chip* and the acronym *SoC*.

Now if the interconnections are scaled at the same rate as devices, according to Table 1.1 the RC delay of a constant-length wire increases by 32% per year, and that of a global wire by 49% per year. The most telling tale however, is told by the comparison between wire delay and gate delay; the parameter RC/τ increases by 51% and 71% respectively for constant length and global wires. In 2002, a typical gate with a fan-out of 4 has a delay of approximately 120 ps while a global wire has a delay of approximately 1 ns giving a ratio of 1:8. In 2010 the delays scale to 40 ps and 26 ns respec-




Figure 1.3: Delay for local and global wiring versus feature size (reproduced from [ITRS01])

tively, giving a ratio of 1:650! Figure 1.3 which is reproduced from [ITRS01] shows the relative delay where scaling has been done in a hierarchical manner for local and global wires. When global wires are scaled less aggressively, the wire delay to gate delay ratio is not so high, but is still a considerable factor.

These analyses were based on a first-order approximation of the wire capacitances. In the DSM regime, the fringing capacitance of device sized wires can be an order of magnitude higher than the parallel-plate capacitance. Most of this fringing capacitance is to an adjacent wire, resulting in capacitive cross-talk. Cross-talk couples a noise pulse onto the line, causing spurious switching across thresholds, and affecting the propagation delay. The RC nature of the wires means that the delay is diffusive, and a linear mapping of the rise time, which increases uncertainty and precludes pipelining. Hence wire delay has the potential to be a severe performance bottleneck.

To overcome this, innovative signalling techniques that are tailored to RC lines to get the best out of them are required. The most commonly practised techniques involve repeaters. Low swing techniques are also adopted, but these are much more susceptible to noise. Row 5 of Table 1.1 shows that the IR drop to rail voltage ratio grows at 15% for local wires, 51% for standard length wires and 71% for global wires, which means that the noise margin available is shrinking all the time. With low swing techniques, there is a trade-off between delay and noise immunity. With repeaters it is between delay and area or power. At the system level, architectures that exploit locality and relax the need for global communication are necessary [Sylvester98].

1.2.2 Off-Chip Communication

While the number and speed of gates on a chip increase exponentially with time, the number of pins have been increasing at best linearly with the chip dimension for peripheral bonding techniques. This is because noise considerations restrict the length of the bonding wire from the package contact to the die, and hence the peripheral depth to which pins can be packed on the package. Therefore off-chip communication is severely constrained, requiring more advanced transceiver designs to increase the bandwidth per pin [Dally97]. The dearth of pins also places more restrictions on the design of the power supply network, as more gates per pin per means longer current paths and increased current in each path, requiring more on-chip bypass capacitance [Dally98]. With the availability of area array bonding techniques where the pins are placed over the entire surface of the package, the number of pins grows with the square of the chip dimension. Also the inductance of the bond wire is eliminated, and the resistive drop over the on-chip power supply grid is much less, as the current paths are shorter. This eases the requirements on on-chip bypass capacitance, but there still exists a need for innovative off-chip signalling schemes. Multi-chip packaging techniques, where several chips exist in one package in a vertical stack (*System-In-Package* or *SiP*), and the inter-chip links are implemented locally, are another option to *SoC* [Zheng01].

1.2.3 Dealing with Complexity

As narrated, advances in IC fabrication technology have led to the Moore law scaling of device density in VLSI chips and resulted in equally dramatic increases in device speed over the past twenty five years. This led to the digital revolution that has seen a proliferation of cheaper products with greater functionality in virtually all spheres of electronic applications. A key factor of this success story has been the ability, at each stage of advancement in device density, to cope with the increasing complexity of design through automation. A crucial concept in automation of the digital design process

is abstraction. Abstraction allows the details of an implementation to be hidden and replaced with a black box view or model characterized by fewer components, allowing both partitioning and hierarchical design to be easily adopted in a top-down approach.

In the DSM regime, the digital abstraction that allows design simplicity is challenged by electrical level issues that affect signalling, timing, power and noise. Additionally, the sheer complexity of the systems means that computer-aided design (CAD) tools such as timing and verification programs take unacceptably high run times to analyse entire chips. Hence at the same time, more complicated effects need to be modelled and the models themselves should be simpler! This is one reason that architectures which exploit locality ([Sylvester98], [Sylvester99a]) and standardise on-chip communication via common protocols are being touted as the way forward. Such architectures are referred to as *Network-on-Chip (NoC)* architectures and have been proposed by several research groups ([Dally01], [Sgroi01], [Benini01], [Hemani99]). The limited size of each block, although big enough to place a microprocessor resource, is not so big that wire delay is pathologically high, and is also small enough to utilise CAD tools. The overhead associated with the network is acceptable when considering the overall size of the chip, and the benefits of modularity it confers.

1.3 Scope of Thesis and Author's Contribution

The scope of this thesis is to examine interconnect modelling and analysis techniques that are suitable for DSM SoC. This includes parasitic extraction techniques, electrical level modelling of different kinds of interconnect, analysis of these circuits, and delay minimisation techniques for on-chip signalling.

The technical contributions of this thesis are threefold. Firstly analytic second order models for the transfer functions of general arbitrarily-coupled *RC* trees are derived, which represent the minimum complexity associated with second-order models for this class of circuits. They are suitable for delay and noise estimations in complex, coupled interconnect systems, early in the design flow, when computational speed is of paramount importance.

Secondly, new techniques are proposed for the optimisation of on-chip buses. A novel method of simultaneous repeater and wire optimisation is described, that derives the unique optimal configuration when the available area resources are limited. Analytic guidelines dependent only on process parameters are derived for optimal sizing of wires and repeaters.

Finally, an investigative study is reported of the physical issues related to implementation of packet-switched networks on chip, and cost and performance metrics are ex-

tracted for a typical 65 nm technology expected to be available in 2007. All major issues related to the implementation are discussed.

Since this thesis is written in the form of a monograph, these main topics and the papers on which they are based are listed below.

1.3.1 Second Order Modelling of Arbitrarily-Coupled RC Trees

- [A] D. Pamunuwa, S. Elassaad, and H. Tenhunen, "Modelling noise and delay in VLSI circuits," *Electronics Letters*, Vol. 39 Issue 3, pp. 269-271, Feb. 2003.

Technical Contribution in Paper- New models for estimating delay and noise in VLSI circuits, based on closed form expressions for the first and second moment of the impulse response in coupled RC trees are reported.

Author's Contribution to Paper - The first author came up with the concept, theoretical formulations and simulations, and wrote the manuscript.

- [B] D. Pamunuwa, S. Elassaad, and H. Tenhunen, "Modeling delay and noise in arbitrarily-coupled RC trees," under review at *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, Nov. 2003.

Technical Contribution in Paper- The models reported in [A] are derived from first principles, and their physical and mathematical bases are investigated. Test cases involving different corner cases are considered.

Author's Contribution to Paper - The first author came up with the concept, theoretical formulations and simulations, and wrote the manuscript.

- [C] D. Pamunuwa and S. Elassaad, "Analytic modeling of interconnects for deep submicron circuits," in *Proc. International Conference on Computer-Aided Design (ICCAD 03)*, San Jose, CA, Nov. 2003 (in press).

Technical Contribution in Paper- The computational complexity of the models reported in [A] are discussed in detail. It is shown that they lend themselves to *incremental computation*, and that a single order n traversal of the tree (where n is the number of nodes) is sufficient to incorporate changes after the initial traversals of the tree.

Author's Contribution to Paper - The first author came up with the concept, theoretical formulations and simulations, and wrote the manuscript.

1.3.2 Bandwidth Optimisation of On-Chip Buses

- [D] D. Pamunuwa and H. Tenhunen "On dynamic delay and repeater insertion in distributed capacitively coupled interconnects," in *Proc. International Symposium on Quality Electronic Design (ISQED)*, San Jose, USA, March 2002, pp. 240-245.

Technical Contribution in Paper- Using switch-factor based delay models for important switching patterns in uniformly coupled lines, models that optimise repeaters to compensate for dynamic switching effects are proposed. Optimal repeater insertion for minimising delay with worst-case cross-talk and area constrained optimisation are considered.

Author's Contribution to Paper - The first author came up with the concept, theoretical formulations and simulations, and wrote the manuscript.

- [E] D. Pamunuwa, L. R. Zheng, and H. Tenhunen, "Maximising throughput over parallel wire structures in the deep submicrometer regime," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 11, no. 2, pp. 224-243, April, 2003.

Technical Contribution in Paper- Using closed-form equations that match the wire geometry to the wire parasitics, and the repeater models derived in [D], a novel analysis is conducted for optimising throughput over a given metal area. This analysis is used to show that there is a clear optimal configuration for the wires which maximizes the total bandwidth. Additionally closed form equations are derived, the roots of which give close to optimal solutions. It is shown that for wide buses, the optimal wire width and spacing are independent of the total width of the bus, allowing easy optimization of on-chip buses.

Author's Contribution to Paper - The first author came up with the concept, theoretical formulations and simulations, and wrote the manuscript.

- [F] D. Pamunuwa, L. R. Zheng, and H. Tenhunen, "Combating digital noise in high speed ULSI circuits using binary BCH encoding," in *Proc. IEEE International*

Symposium on Circuits and Systems (ISCAS), Geneva, Switzerland, May 2000, vol. 4, pp. 13-16.

Technical Contribution in Paper- This paper examines the issue of high speed signalling in DSM technologies and proposes the use of particular BCH codes to improve the bit-error-rate in the face of noise. Simulations are conducted to investigate their performance, and plots showing coding gain are derived.

Author's Contribution to Paper - The first author carried out all the analysis work including simulations, and wrote the manuscript.

1.3.3 Physical Issues in Implementation of Networks-On-Chip

- [G] D. Pamunuwa, J. Öberg, L. R. Zheng, M. Millberg, A Jantsch and H. Tenhunen, "Layout, performance and power trade-offs in mesh-based network-on-chip architectures," in *Proc. IFIP International Conference on VLSI Systems-on-Chip*, Darmstadt, Germany, Dec. 2003 (in press).

Technical Contribution in Paper- Under the assumption that a packet-switched network of resources is a viable option for future systems-on-chip, some physical issues in the design of the switches and inter-switch connections are investigated. A study is conducted for a CMOS technology expected in about 5 years, and the overhead cost and bandwidth performance of the network are investigated. Parameters of interest are extracted, and trade-offs in the layout of the network are discussed.

Author's Contribution to Paper - The first author carried out all the analysis work including derivation of models and simulations, and wrote the manuscript.

- [H] D. Pamunuwa, J. Öberg, L. R. Zheng, M. Millberg, A Jantsch and H. Tenhunen, "A study on the implementation of 2-D mesh-based networks-on-chip in the nanometre regime," under review, *Integration - the VLSI Journal, Special Issue on Networks-on-Chip and Reconfigurable Fabrics*, Elsevier, Aug. 2003.

Technical Contribution in Paper- For the mesh-based architectures considered in [G], a study is carried out for a future technology with parameters as predicted by the International Technology Roadmap for Semiconductors to yield a quantitative comparison of the performance and power trade-off.

Author's Contribution to Paper - The first author carried out all the analysis work including derivation of models and simulations, and wrote the manuscript.

- [I] J. Öberg, D. Pamunuwa, L. R. Zheng, M. Millberg, A Jantsch and H. Tenhunen, "A feasibility study on the performance and power distribution of two possible Network-on-Chip architectures," under review, *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, Oct. 2002.

Technical Contribution in Paper- Additional issues in the implementation of NoC structures and the resource blocks, including regulation of the power-supply network, and the trade-off between device integration and increasing area requirement for on-chip smoothing capacitance, are discussed.

Author's Contribution to Paper - The second author carried out a significant portion of the analysis work, including derivation of certain models and carrying out simulations, and also wrote a portion of the manuscript.

1.4 Thesis Organisation

This thesis is structured in the following manner. This first chapter constitutes an introduction to the thesis, and the second a general introduction to the field and the material covered in the thesis, and do not contain any original contributions by the author. Each chapter that follows is based upon one or more papers mentioned in section 1.3. Hence these chapters are structured like papers, starting with a short introduction that puts the specific work reported into context, and going on to describe the work carried out by the author in detail.

The chapters are organised so that there is a logical progression from modelling of different elements to employing those models in different applications. The interconnect modelling is divided into two parts, with chapter 2 covering parasitic extraction techniques and different electrical level models. The next chapter describes the analysis of these different kinds of circuit models. Chapter 4 discusses the optimisation of repeaters for signalling over nets in the face of crosstalk. In chapter 5, work related to bandwidth optimisation in on-chip buses is reported. The next chapter describes the feasibility study carried out on the physical issues related to Networks-On-Chip. Chapter 7 is the summary, and chapter 8 the list of references cited in this thesis. A brief description of each chapter follows.

Chapter 2: Modelling of Interconnects at High Frequencies

Interconnect modelling techniques that have been reported in the literature are discussed. Parasitic extraction techniques are reviewed extensively, and models suitable for use in modern VLSI systems are identified.

Chapter 3: Delay and Noise Analysis of Interconnect

Existing models are revised and new second order models for estimating the transfer functions of coupled trees derived from a rigorous approach. Their performance is tested extensively on different testbeds including all possible corner cases. The mathematical and physical basis of the models are also examined. Details of implementation are discussed which show that these models represent the minimum computational complexity for this class of circuits. Comparisons are made against other moment-based models.

Chapter 4: Repeater Modelling and Insertion for Coupled Nets

Existing methods of modelling and analysing repeater insertion are reviewed, and some modifications which allow the optimisation of repeaters for coupled nets are presented. Area-constrained optimisation is considered.

Chapter 5: Optimal Signalling for On-Chip Buses

Optimisation of wide on-chip buses for DSM circuits is investigated in this chapter. Using the repeater models derived in the previous chapter and closed form models that map the wire geometry to the parasitics, solutions that allow simultaneous wire and repeater optimisation are derived. It is shown that there is a unique optimal solution for a given set of area and repeater resources. Analytic guidelines are derived.

Chapter 6: Designing System-on-Chip Communication Networks

Chapter 6 carries out a feasibility study for implementation of mesh-based network-on-chip architectures in an example 60nm technology that will become available in the time period 2008. Cost and performance metrics are extracted and all relevant issues are discussed.

Chapter 7: Summary and Conclusions

The work is summarised and areas for future work are identified.

2. Modelling of Interconnect at High Frequencies

This chapter considers the non-ideality of wires at high frequencies, and reviews established techniques of modelling interconnects for different applications. The concept of parasitic extraction is introduced, and various reported models are examined. Different electrical level models are discussed, along with the issue of choosing the appropriate model for a given application. This chapter serves as a general introduction to the technical issues covered in this thesis.

2.1 Introduction

A digital system comprises collections of chips on circuit boards, collections of different boards housed in chassis, and chassis placed in a rack or racks. At each level of hierarchy, signals are transported on different kinds of interconnections. On-chip wires constitute the lowest level in a hierarchy that spans chip- to package-level connections (such as bond wires, package vias and solder balls, and package traces), circuit-board-level connections (thick film wires), backplane-level wires (thick film metal layers or cables), chassis-level connections (more cables) and finally rack-level connections (such as bus bars made of solid metal straps or rods for power connections) [Zheng01], [Dally98]. This thesis is concerned primarily with on-chip wires.

The ideal approximation of a wire assumes that it can be treated as an equipotential region without any loss. A real wire however presents a load to the signal driver, requires non-zero time for the signal to propagate across it (sometimes requiring multiple reflections), and consumes power. The non-ideality of a wire introduces *parasitic* capacitance, resistance and inductance. Depending on the dimensions of the wire and the rise (fall) times of the signals with which they are gated, the electrical circuit model of the wire differs. A clear distinction can usually be made between the characteristics of off-chip and on-chip wires by dint of the fact that the former kind are mostly lossless, while the latter kind are very lossy with usually negligible inductance. Hence off-chip wires are mostly modelled by *LC* lines, and on-chip wires by *RC* lines, both of which are simplifications of the general *LRC* transmission line. In *LC* lines phenomena such as reflections and ringing can be observed, while in *RC* lines the signal propagation is more diffusive in nature, characterized by slow rise times at the far end.

At frequencies where the circuit dimensions become comparable to signal wavelengths, it is not always possible to identify discrete parasitic elements. The voltages and currents are true field quantities, and correctly calculating any parameter requires solving Maxwell's equations in three dimensions. This is accomplished by tools known as field solvers. Field solvers however are very expensive in terms of computation time, and are usually used only for small subsections that require very careful analysis. For the most part, it is possible to use the circuit approach successfully; i.e. modelling the interconnects as connected lumped elements and using circuit techniques to solve for the required parameter.

The rest of this chapter gives a brief introduction to field solvers, and then considers analytic methods of obtaining circuit parasitics from the geometry of the structures in detail. Different circuit models are considered, and metrics for choosing a model for a given application are discussed. The analysis of the electrical circuit for the parameter of interest is covered in Chapter 3.

2.2 Field Solvers

The accurate extraction of parasitics requires solving Maxwell's equations in 3 dimensions, which is accomplished by tools known as field solvers. Using a 3D field solver is very expensive computationally, and impossible over an entire chip, or even large sub-circuits. A second class of tools that consider strips or slots with uniform cross-sections and sacrifice accuracy for run time reduction, are known as 2D field solvers. A third class that falls in between these two are named (predictably) 2.5D field solvers. They allow arbitrary metal patterns in one or more planes, and discretize only the metal in each plane. Usually using even a 2D field solver is justified only for the most critical portions of the chip.

A field solver uses numerical techniques to solve for the fields in the regions of interest, from which the required frequency dependent parameters such as capacitance or inductance are extracted. These numerical techniques have been classified under various names depending on their approach. The following synopses follow the classification of [Hubing91].

2.2.1 Finite Element Method (FEM)

A finite-element analysis (FEA) discretizes a continuous domain into a number of small homogeneous elements so that the field variation within the element can be approximated by simple models. These elements are connected to each other through nodal points. Most FEMs use variational techniques to obtain the field solution at each

node. That is, some expression known to be stationary about the true solution is minimised or maximised at each node, that together with the boundary conditions result in a set of algebraic equations, the solutions to which result in the parameter of interest at each node.

2.2.2 Moment Method

The method of moments (MoM) refers to an FEA that uses the method of weighted residuals (MWR) to solve a set of differential equations at each node. The MWR reaches a solution in a leap-frog manner by substituting an approximate solution into the equations, and then summing the weighted residual iteratively until the solution converges. Hence it is an integral solution method.

2.2.3 Boundary Element Method (BEM)

BEM is essentially a subset of the method of moments. It is a moment-method technique whose expansion and weighting functions are defined only on a boundary surface. It is derived through the discretization of an integral equation that is mathematically equivalent to the partial differential equation that governs the solution in a domain. Most moment methods utilise a BEM technique in a general purpose solver.

2.2.4 Finite Difference Time Domain (FDTD) Method

In the FDTD method, Maxwell's (differential form) equations are simply modified to central-difference equations, discretized, and solved iteratively. Maxwell's curl equations are:

$$\nabla \times H = J + \epsilon \frac{\partial E}{\partial t} \quad (2-1)$$

$$\nabla \times E = -\mu \frac{\partial H}{\partial t} \quad (2-2)$$

From an examination of (2-1) and (2-2) it can be seen that the time derivative of the E field is dependent on the curl of the H field. This can be simplified to state that the temporal change in the E field (the time derivative) is dependent on the spatial change in the H field (the curl). The result is the basic FDTD equation that the new value of the

E field is dependent on the old value of the E field (hence the difference in time) and the difference in the old value of the H field on either side of the point in consideration in space. The H field is found in the same manner.

2.2.5 Finite Difference Frequency Domain (FDFD) Method

Similar to FDTD techniques, FDFD results from a finite difference approximation of Maxwell's curl equations. However now the time harmonic versions of the equations are employed:

$$\nabla \times H = (\sigma + j\omega\epsilon)E \quad (2-3)$$

$$\nabla \times E = -j\omega\mu H \quad (2-4)$$

Since there is no time stepping, uniformity of mesh spacing is not required. Hence optimal FDFD meshes resemble optimal finite-element meshes.

2.2.6 Transmission Line Matrix (TLM) Method

In the TLM method analysis is performed in the time domain and the entire region of the analysis is gridded similar to FDTD. Instead of interleaving the E-field and H-field grids however, a single grid is established and the *nodes* of this grid are interconnected by virtual transmission lines. Excitations at the source nodes propagate to adjacent nodes through these transmission lines at each time step.

2.2.7 Partial-Element Equivalent Circuit (PEEC) Method

Methods involving PEEC models are very popular because they allow a quasi-static analysis with equivalent elements of resistors, capacitors and inductors. From Maxwell's differential equations, a single integro-differential equation that describes the E field can be constructed:

$$E(r,t) = \frac{J(r,t)}{\sigma} + \frac{\partial}{\partial t}A(r,t) + \nabla\Phi(r,t) \quad (2-5)$$

Each term of (2-5) can be used to define an equivalent circuit element for each metal segment in the grid, with the first corresponding to the resistive term, the second to the inductive term and the third to the capacitive term. Subsequently KVL and KCL are used to solve the resulting circuits using a circuit simulator such as Spice.

Several techniques other than the ones described above have been developed, and a good overview can be found in [Hubing91]. For the most part, the use of field solvers is restricted to critical portions of the chip due to the complexity of the numerical techniques and the consequent high run time. In the following sections, alternative approaches (comprising for the main part analytic formulae) to model resistive, capacitive and inductive parasitics that are cheaper, though necessarily less accurate are examined. These are important in their own right, as they provide an intuitive understanding of the variation in performance with geometry, and are perfectly adequate for a large number of cases.

2.3 Analytic Formulae

2.3.1 Resistance

DC Resistance

A rectangular cross-section is a fairly good approximation for an on-chip wire, as shown in Figure 2.1 (for more complex structures including curves, corners and combinations of square and round cross-sections, the resistance is calculated with the help of field solvers). The DC resistance R_{DC} of such a uniform strip of material is given by (2-6) where ρ is the resistivity of the material.

$$R_{DC} = \rho \frac{l}{tw} \quad (2-6)$$

Since the thickness t is usually a constant for a given technology, it is customary to incorporate it and the resistivity into a single constant called the *sheet resistance* of the material, given in (2-7).

$$R_q = \frac{\rho}{t} \quad (2-7)$$

Then the resistance is given by (2-8)

$$R_{DC} = R_q \frac{l}{w} \quad (2-8)$$




Figure 2.1: Cross-sectional dimensions of a conductor

The sheet resistance gives the resistance of a square of that material for a constant height. Given in Table 2.1 are the resistivities of several materials commonly used for conductors. Although silver is the best in terms of conductivity, its high cost means that it is used only for special applications. The most commonly used are aluminium, which is the most economical, and copper which is more expensive but has much better conductivity.

Table 2.1 Resistivity of materials used for conductors in VLSI circuit fabrication

Material	Resistivity (ρ) $\Omega\text{-m}$
Tungsten (W)	5.5×10^{-8}
Aluminium (Al)	2.7×10^{-8}
Gold (Au)	2.2×10^{-8}
Copper (Cu)	1.7×10^{-8}
Silver (Ag)	1.6×10^{-8}

Resistance with Skin Effect

At high frequencies, the current density inside a conductor is not uniform, but drops away exponentially with depth into the conductor. A cut-off frequency where this phenomenon begins can be identified, and an empirical approximation for this value is given in (2-9) where μ refers to the permeability, and δ_c to the skin depth as given in (2-10). This last is an approximation given in [Zheng01].

$$f_c = \frac{\rho}{\pi\mu\delta_c^2} \quad (2-9)$$

$$\delta_c = 1.5tw(h+w) \quad (2-10)$$

A good explanation of the working is given in [Paul92] and [Dally98].

Below f_c the current is assumed to be spread uniformly across the entire cross-sectional area of the conductor, resulting in the DC resistance given in (2-8), while above it the resistance increases with the square of the frequency. Hence the frequency dependent resistance R_{HF} can be conveniently expressed as given in (2-11).

$$R_{HF} = R_{DC} \left(\frac{f}{f_c} \right)^{1/2} \quad (2-11)$$

2.3.2 Capacitance

The definition of capacitance between two conductors is given in (2-12)

$$C = \frac{Q}{V} = \frac{\iint_s \epsilon \vec{E} \cdot d\vec{S}}{\int_c \vec{E} \cdot d\vec{l}} \quad (2-12)$$

The capacitance of a wire depends strongly on the geometry of the adjoining structures as well. However the electric field lines permeate only a short distance from the conductor in question, and it is possible to estimate the capacitance fairly accurately by simple means. Various empirical equations have been formulated with this end in mind, which are very useful in a priori timing and signal integrity analyses. An overview of closed form modelling of capacitance is given below.

Single Conductor over Ground Plane

The per unit length parallel plate capacitance (the capacitance assuming the E field is entirely contained within the two plates) of the micro-strip line structure given in Figure 2.2.a is given by (2-13):

$$C_p = \frac{w\epsilon}{h} \quad (2-13)$$




Figure 2.2.a: rectangular conductor

Figure 2.2.b: circular conductor

Figure 2.2.c: Shape for estimating fringe capacitance

Figure 2.2: Cross-sectional dimensions of a conductor

The simple parallel plate approximation underestimates the capacitance of a wire by as much as an order of magnitude if applied to the very high aspect ratio (height / width) wires in DSM technologies. It is essential that the contribution of the *fringe* components of the E field to the capacitance is taken into account. The basis of most such estimations is the decomposition of the capacitance into two components, one proportional to the parallel plate capacitance, and another proportional to the capacitance of a circular conductor of diameter t [Lee98]. Hence it is useful to look at the capacitance of a wire over a ground plane as shown in Figure 2.2.b, which is given in (2-14):

$$\frac{2\pi\epsilon}{\ln(2h/t)} \quad (2-14)$$

One of the early approaches¹ detailed in [Yuan82] gives the empirical formula (2-15) which has a straightforward physical motivation.

$$C = \epsilon \left[\frac{w}{h} + \frac{2\pi}{\ln\{1 + (2h/t)(1 + \sqrt{1 + t/h})\}} - \frac{t}{2h} \right] \quad (2-15)$$

The accuracy of this equation however drops rapidly when the ratio w/h falls below values of about 2-3. The trend in modern technologies is to have increasing numbers of metal layers, thus increasing h , and shrinking wire sizes, decreasing w , making the regime below this ratio the most interesting, and hence rendering (2-15) unusable.

1. A comparison of capacitance extraction techniques for a single conductor over a ground plane is carried out in [Barke88].

In [Sakurai83a], the physically motivated approach was abandoned, and one based completely on curve fitting adopted, resulting in (2-16):

$$C = \epsilon \left[\frac{w}{h} + \frac{0.15w}{h} + 2.8 \left(\frac{t}{h} \right)^{0.222} \right] \quad (2-16)$$

This has better accuracy than (2-15) when the w/h ratio drops below 2-3, but is still increasingly inaccurate as the ratio continues to decrease.

Another formula which is slightly more complex and reported in [Meijs84] is given in (2-17):

$$C = \epsilon \left[\frac{w}{h} + 0.77 + 1.06 \left(\frac{w}{h} \right)^{0.25} + 1.06 \left(\frac{t}{h} \right)^{0.5} \right] \quad (2-17)$$

This is reported to be the most accurate in [Barke88] for the values of dielectric thickness ($h=0.75\mu\text{m}$) and conductor thickness ($t=1.3\mu\text{m}$) that were used in the study. Since then several other empirical models have been reported, most of them connected to multi-net structures, which will be considered in the next section.

Coupled Wires over Ground Plane

Shown in Figure 2.3 is a multi-net structure with different capacitance terms. In the same article that was cited above for a single line capacitance, [Sakurai83a], an equation (2-18) for a mutual capacitance is reported.

$$C_c = \epsilon \left[0.03 \frac{w}{h} + 0.83 \frac{t}{h} - 0.07 \left(\frac{t}{h} \right)^{0.222} \right] \left[\frac{s}{h} \right]^{-1.34} \quad (2-18)$$

The total capacitance of the middle wire, C_t is then given by (2-19) where C_s refers to the single line capacitance defined in (2-16):

$$C_t = C_s + 2C_c \quad (2-19)$$

The total capacitance given by this equation is in very good agreement with that predicted by a field solver for the total capacitance of the middle wire, but the individual components are not intended to provide decomposition of the total into ground and coupled components. Since the presence of the adjacent conductors significantly affects the electric field around the central conductor, accurate decomposition requires

that the proximity of the neighbouring conductors, or in a mathematical sense the quantity s , has to be modelled in the expression for the self capacitance. It then follows that the expressions for mutual capacitance are also unusable on their own. Hence although these equations are quite useful for certain applications, they not suitable for any analysis which requires that the distribution of the capacitance into self and mutual components be accurate. An excellent discussion including independent verification of this can be found in [Lee98].

Since then, formulae which better partitioned the components into ground and coupling components have been proposed, in [Chern92] and [Lee97] among others, which are not reproduced here. A range of values for t , h and w are given where the proposed models are valid. Simple expressions for estimating the cross-over capacitance between vertically and horizontally placed interconnects are also proposed in those articles.

A complete set of equations were also proposed in [Zheng00]. The capacitance C_f and C_f' shown in Figure 2.3 refer to the fringing components with and without an adjacent conductor. They are defined in (2-20) and (2-21) respectively.

$$C_f = \epsilon_k \left[0.075 \left(\frac{w}{h} \right) + 1.4 \left(\frac{t}{h} \right)^{0.222} \right] \quad (2-20)$$

$$C_f' = C_f \left[1 + \left(\frac{h}{s} \right)^\beta \right]^{-1} \quad (2-21)$$

The quantity ϵ_k refers to the relative permittivity of the medium, while β is a curve fitting constant. The mutual or coupling capacitance C_c is defined in (2-22):




Figure 2.3: Multi-net Configuration

$$C_c = C_f - C'_f + \epsilon_k \left[0.03 \left(\frac{w}{h} \right) + 0.83 \frac{t}{h} - 0.07 \left(\frac{t}{h} \right)^{0.222} \right] \left(\frac{h}{s} \right)^{1.34} \quad (2-22)$$

This particular partitioning allows a self capacitance to be defined both for a conductor sandwiched between two other conductors, and also for one which has just one adjacent conductor. The capacitance in the first instance is given in (2-23) and in (2-24) for the second.

$$C_{s, mid} = C_p + 2C'_f \quad (2-23)$$

$$C_{s, corner} = C_p + C_f + C'_f \quad (2-24)$$

These equations are deemed to be valid when the wire geometries are in the range defined by the set of inequalities given in (2-25), when the error was contained to within 10%.

$$0.3 < (w/h) < 30 \quad 0.3 < (t/h) < 10 \quad 0.3 < (s/h) < 10 \quad (2-25)$$

The empirical constant β is calculated by generating a database with a field solver, and then using curve fitting techniques. The closest physical interpretation of it is that it is related to the cross-sectional dimensions of the wire. Its value varies between 1 and 2, and 1.75 is typical for DSM wires.

When the wire geometry is such that it goes out of the range defined in (2-25), it is possible to treat the rectangular conductors as equivalent round wires if condition (2-26) is satisfied where H is as defined in (2-27).

$$w \leq 2H \quad (2-26)$$

$$H = h + t/2 \quad (2-27)$$

The radius of the *equivalent* round conductor is then approximated to be:

$$R = 0.25w + 0.335t \quad (2-28)$$

Now the self capacitance changes to the sum of (2-29) and the parallel plate capacitance, while the mutual capacitance term changes to the expression given in (2-30)

$$C_f = \frac{\pi \epsilon_k}{\ln\left(2H\sqrt{(2H^2) + (d+w)^2} / [R(d+w)]\right)} - \frac{w\epsilon_k}{2H} \quad (2-29)$$

$$C_c = 2\pi\epsilon_k \ln\left(\frac{\sqrt{(2H^2) + (d+w)^2}}{(d+w)}\right) / \left(\ln\left(2H\sqrt{(2H^2) + (d+w)^2} / [R(d+w)]\right)\right) \ln\left(2H(d+w) \left[R\sqrt{(2H^2) + (d+w)^2}\right]\right) \quad (2-30)$$

In general when the wires are relatively less tightly coupled, intuition derived from analysis of simple round conductors can be employed. It can be seen that in these two equations, terms proportional to the capacitance of a round conductor over a ground plane are present. These equations in turn are more accurate when the geometry does not exceed the range specified in (2-31) and (2-32). In comparison with values calculated from a 2D field solver, the error was contained to within 12% for geometries within this range.

$$t/w < 2 \quad (2-31)$$

$$d/w > 1 \quad (2-32)$$

2.3.3 Inductance

The definition of the inductance of a wire loop is given in (2-33):

$$L = \frac{\iint_s \vec{B} \cdot d\vec{S}}{i} = \frac{\oint_c \vec{A} \cdot d\vec{l}}{i} \quad (2-33)$$

The parasitic inductance, unfortunately, is much more dependent on the global environment than the parasitic capacitance. The inductance of a wire, as can be seen in the

definition, depends on the loop which comprises the signal path, *and* the return path. Initial work assumed that the return path was contained within the substrate ([Priore93], [Jarvis63] and [Eo93]). However subsequent work established that the current return path is primarily in the power distribution network, and other adjacent wires ([Deutsch90], [Deutsch95a], [Deutsch97], [Shoji96], [Deutsch95b], [Massoud98], [Krauter98] and [Deutsch96]). The reason is that in digital designs, there is a profusion of metal wires, and only a very sparse metal environment will cause the return current to choose the substrate. The loop formed by the signal wire and the return path can potentially extend to several hundred micrometers away from the wire under consideration. This vastly complicates the extraction of parasitic inductance of a given wire, as it depends not only on the characteristics of that particular wire, but also potentially on the characteristics of several thousand other wires. Hence it would appear that expensive 3D numerical techniques are necessary for good accuracy.

However, two characteristics of on-chip inductance can be exploited to use simpler techniques. The first is that the signal waveform on a wire is relatively insensitive to errors in the value of the inductance used. This is particularly true for two important characteristics of the waveform: rise time and propagation delay. The second is that the inductance is a slow varying function of the wire width and geometry of the surrounding conductors. Hence the need for accurate modelling of the geometry of the conductors is not as paramount as with resistance and capacitance. A good discussion can be found in [Ismail01], where the authors run simulations for an *LRC* tree with the correct inductance as calculated by a field solver and a constant inductance. They show that for nets where inductance is significant, errors of up to 30% in the extracted inductance leads to errors between 4% and 10% in the propagation delay, depending on the damping factor. Neglecting the inductance altogether and using an *RC* model (for those nets where inductance needs to be modelled) results in a much bigger error. (This immediately raises the question of when inductance should be modelled in a line, and that will be covered in later, after electrical models of wires are introduced.) Hence it is possible with reasonable accuracy, to use even a uniform inductance per line length, or use table look-up or curve fitting techniques to obtain a value that is dependent on characteristics such as the width and pitch of the signal wire and surrounding power distribution network.

Now in calculating the parasitic inductance, one way of approaching the signal return problem is to use the concept of partial inductance developed in 1908 [Rosa08]. This technique assigns portions of the loop inductance to segments along the loop, which can then be summed up to yield the total inductance for a wire, consisting usually of self and mutual inductance terms. The partial inductance terms may be either negative or positive depending on the relative orientation of the currents.

One simple method to calculate the inductive parasitics is to use the identity give in (2-34) where C and L refer to the inductance and capacitance.

$$CL = \epsilon\mu \quad (2-34)$$

This is valid only when the conductors are surrounded by a uniform dielectric, but in the presence of a dielectric boundary (such as for a micro-strip line which has the Si substrate below and air or possibly SiO₂ above), it is possible to define an “average” dielectric that takes this effect into account. In the procedure outlined in [Zheng01], the capacitances are calculated for a three net structure using (2-22), (2-23) and (2-24) with air as the surrounding dielectric. These values are used with (2-34) to yield (2-35) and (2-36) where either C_{s_mid} or C_{s_corner} is substituted for C_s .

$$L_s = \frac{\epsilon_0\mu_0}{2} \left(\frac{1}{C_s^0} + \frac{1}{C_s^0 + 2C_c^0} \right) \quad (2-35)$$

$$L_m = \frac{\epsilon_0\mu_0}{2} \left(\frac{1}{C_s^0} - \frac{1}{C_s^0 + 2C_c^0} \right) \quad (2-36)$$

This assumes that the signal return is contained within the adjacent conductors. For more accurate extraction, partial inductance techniques as given in [Krauter98] and [Shepard00] may be used. Briefly, in these techniques the return is limited to some reasonable area, usually bounded by the closest power and ground lines. Then frequency independent resistance and inductance values are computed for each cross-section assuming uniform current distribution. If the skin effect is noticeable at the frequencies under consideration, the cross-sections are subdivided into sections smaller than the skin depth at the maximum frequency of interest. To calculate the partial inductances, equations formulated in the early 1900s [Rosa08] are used. A comprehensive tabulation of expressions for self and mutual inductances for different geometries is provided in [Grover62]. Just the formulae for self and mutual inductances of a rectangular wire are shown here, in (2-37) and (2-38).

$$L_s = 2l \left(\ln \left(\frac{2l}{w+t} \right) + 0.5 + 0.22 \frac{(w+t)}{l} \right) \quad (2-37)$$

$$L_s = 2l \left(\ln \left(\frac{2l}{d} \right) - 1 + \frac{d}{l} \right) \quad (2-38)$$

Finally, the best accuracy can be achieved by field solvers. Techniques that use PEEC models are very popular, and is the basis of FASTHENRY, a 3D inductance extraction tool [Kamon94] that is popular and freely available.

2.3.4 Inductance vs Capacitance Extraction

There is a certain duality in the problems inherent in capacitance and inductance extraction [Shepard00]. Capacitance is very localised in that the field lines from a given conductor tend to terminate on the nearest neighbour conductors. This makes the capacitance matrix sparse (since only the terms related to the coupling between close wires need to be included, the others being insignificant), and hence analytic formulae need only model the geometry of the wire in question and the adjacent wires. However the non-zero interaction terms have a very strong geometry dependence. This makes the accuracy of analytic formulae somewhat limited, and an error contained to within roughly 10% is about the best that can be hoped for in different capacitive components of complex structures.

By contrast, strong geometry dependence does not exist for inductance and local calculation is rather easy, rendering analytic formulae for partial inductances quite accurate. But again, contrary to the situation with capacitance, the locality problem is much harder. Current loops defining flux linkages can, and often do, extend far beyond the conductor in question, making the inductance matrix very dense. Hence sparsifying the inductance matrix is a difficult problem. Because of the relative insensitivity of signal waveforms to variations in the parasitic inductance though, expensive extraction techniques can be avoided to a fair extent for most circuits, with some approaches even adopting a constant precharacterized inductance. For greater accuracy, inductance can be extracted by methods that seem to fall into three categories. The first is to invert the capacitance matrix, which is very simple, but yields only approximate values as it assumes transverse electromagnetic (TEM) propagation, and uses averaging techniques to account for heterogeneous media. The second is to use empirical formulae for partial inductances, and sum up the values for a window which is estimated to contain the current return paths. Finally the most accurate and also most expensive method is to use a field solver.

2.4 Electrical Level Modelling

In most instances, the behaviour of the field quantities can be captured by simplified models such as isolated transmission lines, coupled transmission lines, cascaded *LRC* and *RC* networks, lumped *LRC* and *RC* circuits, and simple one element capacitive or inductive loads. In this section different models will be considered. Assuming TEM propagation, all wires can be generalized to transmission lines which have series resistance and inductance, and parallel capacitance and conductance. All other models are simplifications of the general transmission line. Hence equations governing the general line will first be described with some attention being given to results that are well known and derived from basic transmission line theory. Subsequently, simplified models and their suitability for use in different cases are considered.

2.4.1 The General Transmission Line

In heterogeneous media such as layered dielectrics, the propagation mode is not restricted to the TEM mode. However if the separation of the conductors is small compared to the wavelengths, which it generally is in a chip, approximate TEM propagation exists. Shown in Figure 2.4 are the lumped model consisting of cascaded sections, and the infinitesimal model of a single section in the limit where the distance Δx tends to zero. The voltage $V(x,t)$ and current $I(x,t)$ on this line are functions of position x , and time t . The partial differential equations that describe the behaviour of the line can be obtained by observing that the gradient of the voltage is the drop across the series elements, and the gradient of the current is the current through the parallel elements, leading to (2-39) and (2-40).

$$\frac{\partial V}{\partial x} = rI + L \frac{\partial I}{\partial t} \quad (2-39)$$




Figure 2.4.a: Lumped model




Figure 2.4.b: Infinitesimal model

Figure 2.4: Lumped and Infinitesimal models of a transmission line

$$\frac{\partial I}{\partial x} = gV + c \frac{\partial V}{\partial t} \quad (2-40)$$

Combining these two equations results in the general wave equation given by (2-41):

$$\frac{\partial^2 V}{\partial x^2} = (rc + lg) \frac{\partial V}{\partial t} + lc \frac{\partial^2 V}{\partial t^2} \quad (2-41)$$

The driving point impedance of an infinitely long line can be shown to take the value given in (2-42).

$$Z_0 = \sqrt{\frac{r + j\omega l}{g + j\omega c}} \quad (2-42)$$

A line which has a finite length and is terminated by its characteristic impedance will appear as an infinitely long line to the driving source, and Z_0 defines the ratio of voltage to current at any point along the line. However if the load Z_L is different from Z_0 , it will impose its own ratio of voltage to current at the termination point. The only way to reconcile this conflict is for some portion of the incident signal to reflect back towards the source. Adopting the convention that $x=0$ at the load end, the solution to (2-41) can be expressed as the sum of two voltage waves, one corresponding to the incident wave and the second to the reflected wave.

$$V = V_i e^{-\gamma x} + V_r e^{\gamma x} \quad (2-43)$$

The symbol γ is known as the propagation constant, and is defined in (2-44).

$$\gamma = \sqrt{(r + j\omega l)(g + j\omega c)} \quad (2-44)$$

The ratio of the reflected voltage (current) to the incident voltage (current) at any point along the line is given by the reflection coefficient Γ as defined in (2-45) where Γ_L is the ratio of the reflected to incident quantities at the load end.

$$\Gamma = \Gamma_L e^{2\gamma x} \quad (2-45)$$

In the next few sections various simplified solutions to (2-41) will be considered, with ways of reconciling the zero impedance return path shown in the models of Figure 3.1.1 with the non-zero values associated with real returns being investigated.

Modelling the General Transmission Line

A transmission line is often approximated by cascaded lumped sections as shown in Figure 2.4 when using circuit simulators. To maintain acceptable accuracy, the sections must be chosen so that the resonant frequency in each LC section is small compared to the highest frequency of interest in the circuit. Additionally, the time step Δt used in the simulator must be small compared to the period of the LC circuit in each section [Dally98]. This leads to the inequality given in (2-46).

$$\Delta t \ll 2\pi\Delta x\sqrt{lc} \ll 2t_r \quad (2-46)$$

This implies that the number of sections needed for reasonable accuracy may be arbitrarily high, and is in contrast to the situation for RC lines, where the response converges very fast with increasing sections regardless of the relative magnitudes of the signal rise time and the time constants of each section.

2.4.2 Simplified Transmission Line Models

Lumped Single Element Models

Lumped capacitive load

If the length of a wire is much shorter than the shortest wavelength of interest (corresponding to the rise time of the signal over the wire), and there is no DC current over it, the resistive and inductive parasitics may be safely ignored and the capacitance treated as a lumped element [Dally98]. Such a model is used for very short on-chip wires that drive static gates, and short off-chip wires. If the metal environment is dense, most of the capacitance may be to other wires, when capacitive cross-talk occurs, and the model has to be appropriately modified.

Lumped resistive model

An on-chip wire that distributes substantial amounts of DC current (such as in the power distribution network) can usually be modelled as a lumped resistive load. The

capacitance and inductance can be ignored if the line holds a relatively steady voltage. The main issue with such wires is the IR drop across the wire.

Lumped inductive model

A line that carries substantial amounts of AC current (such as in the power distribution network) can on occasion be modelled as a lumped inductor. If the Ldi/dt drop associated with the inductance usually dominates, the resistance and capacitance may be neglected.

RC Transmission Line

This is one of the most important simplifications of the general transmission line, as RC lines are typical of a majority of on-chip wires. When the inductance is negligible, (2-41) simplifies to (2-47).

$$\frac{\partial^2 V}{\partial x^2} = rc \frac{\partial V}{\partial t} \quad (2-47)$$

This is known as the diffusion equation, and is widely encountered in heat conduction problems. For the initial conditions that are of interest, the solution is an infinite sum of exponential terms. However, RC lines usually exhibit a response that is governed by a few low frequency poles, and quite often even by a single pole. This allows successful reduced-order modelling.

Depending on the length of the wire, the number of sections with which the wire is modelled has to be chosen. The response of a cascaded section model converges to the true solution very quickly with the number of sections. Usually a five section π model is 99% accurate to the response of a true distributed line. The difference between a five-section π model and a one-section π model can be as high as 41%. Hence the number of sections with which a wire is modelled should be chosen after a careful analysis of its physical length.

2.5 Choosing a Wire Model

The most important issue when choosing a wire model is the question of whether inductance should be included. There are a growing number of works in the literature that address this point [Deutsch97], [Ismail99], [Krauter99], [Lin00] and [Banerjee01]. These expressions, though formulated in different ways, are for the most part equivalent. Reproduced here are the expressions from [Ismail99], because they neatly quan-

tify a window where inductance is important, and have a straightforward physical motivation.

$$\frac{t_r}{2\sqrt{lc}} < length < \frac{2}{r}\sqrt{l} \quad (2-48)$$

A lossy transmission line has series resistive and inductive segments and parallel capacitive segments (the conductive loss to ground can be safely ignored for the vast majority of VLSI circuit applications). The symbols r , l and c in (2-48) refer to the per unit length quantities while $length$ refers to the length of the wire. Now in a qualitative sense, if the combined capacitive and inductive reactance at the highest frequency of operation (defined by the rise time at the output of the driver) is comparable with the series resistance, inductance cannot be ignored. This condition defines the second inequality of (2-48); if the line is longer, the loss is high enough to mask out the inductive effect. However the line also has to be long enough for the delay at the speed of light in the medium to be comparable to the rise time; if not, the gating signal is too slow for the reactance to compete with the resistance. This defines the first inequality. Additionally, this window may never exist, if the combination of the rise time and loss is such that short lines have a time of flight delay that is much less than the rise time, and long lines have far too much loss for inductance to be important. This condition is defined in (2-49).

$$t_r > 4\frac{l}{r} \quad (2-49)$$

The inequalities (2-48) and (2-49) can be used to show that for the majority of signal nets in VLSI circuits, inductance can be safely ignored.

2.6 Summary

On-chip wires at high frequencies exhibit non-idealities, which can exactly be captured only by computing the characteristics of the electromagnetic field generated by excitation of the complete distributed interconnect structure. This is carried out by a tool set known as field solvers, and requires expensive simulations. One way of reducing the complexity is to partition the problem into calculating a set of geometry dependent parasitics, and solving a discrete electrical network made up of parasitic elements. This allows a delineation of the two tasks, and the requirement in this approach is to do each of the two tasks as efficiently and as accurately as possible.

Parasitic extraction basically consists of calculating equivalent resistive, capacitive and inductive elements to build a discrete electrical network. Computing resistive elements is more or less straightforward, as the DC resistance can be used for the majority of cases. Extraction techniques for capacitive and inductive parasitics are coloured by their characteristics, which dictate high geometry dependence and a sparse matrix for the former, and a loose dependence on geometry and a dense matrix for the latter. Although based on many assumptions, closed-form formulae work quite well for most cases.

Different electrical level models lead to different effects being modelled and accuracy being traded off for computational complexity. A key question is when inductance is important, and analysing a distributed *RLC* transmission line under the assumption of TEM propagation leads to some simple metrics which can be applied as a test. Most of the time, a coupled *RC* circuit is adequate for modelling interconnected structures of signal nets. This is the subject of the next chapter.

3. Delay and Noise Analysis of Interconnect

The analysis of interconnected networks of linear(ised) elements is covered in this chapter. Existing techniques are revised, and new second order models for estimating the transfer functions of arbitrarily-coupled RC trees are proposed. Their computational complexity is examined and the accuracy compared to other moment-based models.

3.1 Introduction

The previous chapter discussed electrical level modelling of interconnect, from the extraction of *parasitic* information, to circuit models, to choosing between different kinds of circuit models. This chapter discusses analysis methods for obtaining delay and noise information from such circuits. With decreasing gate delays and increasing wiring density, noise modelling and its impact on performance and functionality has become very important. The majority of signal wires are typically very lossy, and higher aspect ratios to control the resistance result in increased capacitive coupling. This, together with smaller signal rise times, results in heavy cross-talk, which couples a noise voltage onto the victim net. A distinction is usually made between the coupled noise amplitude, and the effect of noise on delay. The former can cause functional failures by causing the voltage to swing above or below the logic threshold, while the latter has an impact on the cycle time. The circuit that captures this behaviour when no restrictions are made on the topology, is an *arbitrarily-coupled RC tree*.

Now the ability to put billions of transistors on a single die has also imposed severe restrictions on the computational complexity of noise and delay models used in an iterative design flow. While more accurate modelling is necessary, the sheer size of the systems prohibits expensive dynamic simulation. Consequently the subject of delay and noise modelling for VLSI circuits has received a vast amount of attention in the literature. The three attributes of accuracy, computational simplicity and generality, are however difficult to encompass in a single integrated model. Most reported models that consider the effect of cross-talk on noise and delay either use heuristics that are tailored for specific topologies, or use multiple moments that make them expensive.

In this chapter new models are described for generating second order transfer functions from any driver to the receiver in general arbitrarily-coupled trees with guaranteed stability. The summation of all waveforms results in the complete response to all switching events at the node of interest, with no restriction on arrival times, and allows

both delay and noise estimations. The models proposed exhibit good accuracy, and represent the minimum possible complexity for second order estimations for this class of circuits without compromising on generality.

3.2 Background

An accurate analysis of interconnects requires solving Maxwell's equations in three dimensions, which is prohibitively expensive in terms of computation time. However it is possible to use simplified models in most cases to capture the important effects in the regime of interest [Chiprout98]. A particular concern with falling rise times is how important inductive effects are, and when and how they should be modelled (see previous chapter). A growing body of literature now exists that address this issue. They propose metrics that relate the physical dimensions of the wire to the signal rise time by assuming TEM propagation, to determine when neglecting inductive effects results in significant errors. The general consensus now is that modelling inductance is necessary for special nets such as clock and power lines, and that the majority of signal lines can be accurately modelled by networks of resistors and capacitors, even with very small rise times. It is important however to consider the effect of capacitive cross-talk, which is exacerbated by sharper slew rates. Hence the circuit model of an arbitrarily-coupled RC tree is very important in current and future technologies.

3.2.1 Delay Modelling

Timing analysis in VLSI circuits has long been carried out using the simplified model of an RC tree where all capacitors are connected to ground, which circuit model shall be called a *simple tree* (Fig. 3.1). There is a large body of literature that deals with delay modelling in simple trees. One of the most important and widely used metrics, the first moment of the impulse response, was proposed back in 1948 as an upper bound for the delay in valve circuits [Elmore48], and is known as the Elmore delay. Its attraction is that it uses minimum information and has unmatched algorithmic simplicity and ele-




Figure 3.1: Example of *simple RC tree*

gance, explicitly matches the circuit elements to the delay, and yet exhibits good *fidelity*, giving results as good as more expensive models when used in interconnect optimization algorithms. The error of the Elmore delay however, can be as high as several hundred percent, especially for near-end nodes. Bounds and metrics that gave an indication of when the Elmore model was poor were developed in [Rubinstein83].

A stable approximation to the second order transfer function for simple trees based on the first and second moment of the impulse response, and the sum of the open circuit time constants was proposed in [Horowitz84] and expanded in [Chu87] to encompass charge sharing networks. Later, generic moment-based techniques applicable to any circuit comprising linear elements that allowed the calculation of an arbitrary number of poles, were developed in [Pillage90]. This technique of waveform estimation which was termed Asymptotic Waveform Evaluation or AWE by the authors, calculates as many poles as are necessary, by means of matching the response to multiple moments of the impulse response. A vast body of research now exists which identified and corrected some stability problems with the initial approach and improved on the computational complexity of the algorithms [Feldmann95], [Chiprout92], [Yu99].

An implementation that is optimised for the tree like structures of interconnects was proposed in [Ratzlaff94]. These techniques depend on the Padé approximation, which typically requires $2q$ moments for a q^{th} order approximation. Hence, even though the computation is very efficient, obtaining a second-order model requires the calculation of four moments. Other estimators based on the Arnoldi algorithm [Silveira95] match q moments to a q^{th} order approximation. An example is [Odabasioglu98], which gives reduced order models for linear systems. However the nodal matrices of the system need to be formed, and at least one LU decomposition of the admittance matrix (which has a cubic complexity) is necessary. For initial analysis of complex systems which involves many iterations, such techniques are best avoided when possible.

Hence numerous models have been proposed, that occupy some position in the spectrum defined by the accurate though expensive solution offered by generic moment matching techniques such as AWE (and similar methods) at one end, and the simplicity offered by the Elmore delay at the other. For simple trees, the models of [Horowitz84] represent the minimum computational complexity for a second-order model. Alternate second-order models for the transfer function include those reported in [Kahng95] and [Kahng97], which involve generating equivalent circuits and are more suited for highly inductive lines; and that reported in [Tutuianu96] which yields a stable model from the first three moments.

Now a two (or higher) pole model cannot be solved explicitly for the delay at a given threshold. Hence there are quite a few works that attempt to garner more information

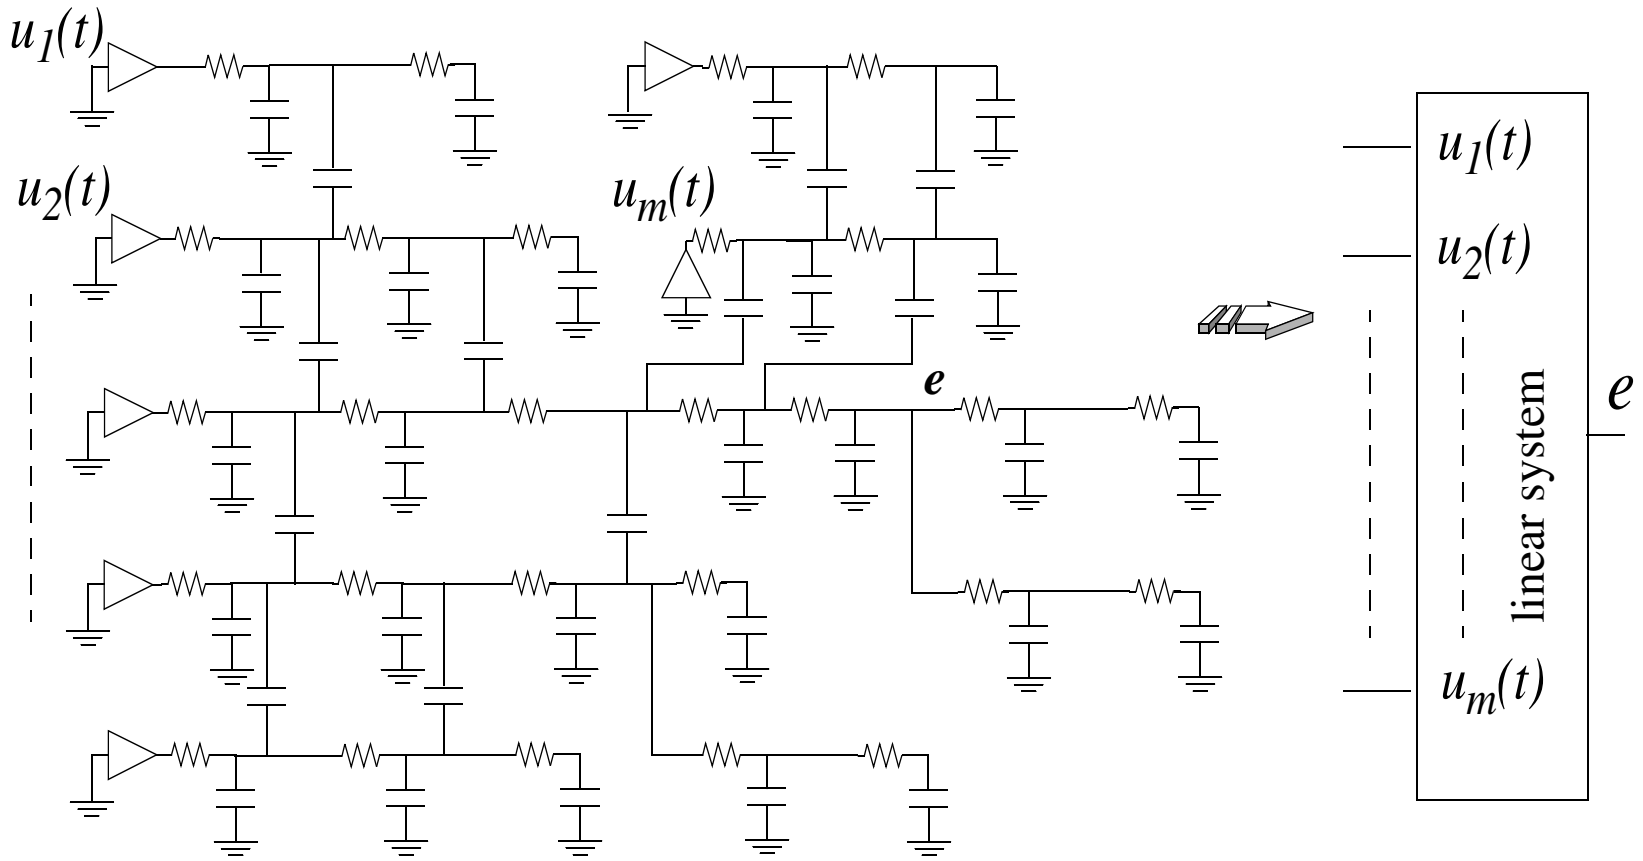


Figure 3.2: Example of *coupled RC tree*

than the first moment (Elmore delay) from the circuit, and match it explicitly to the delay via some heuristic, such as in [Kahng95], [Kahng97], [Tutuianu96] and [Acar99]. The authors of [Alpert01] present two delay metrics, one based on the first two moments, and another based on an effective capacitance model which seeks to overcome the effect of resistive shielding that makes the Elmore delay inaccurate at near-end nodes. Explicit delay models for inductive lines were proposed in [Ismail00]. Different approaches were suggested in [Kay98] and [Lin98], where the moments of the circuit are matched to parameters of probability density functions to yield the delay.

In today's circuits, as mentioned, considering the effect of noise is important. Finding the response of such systems involves solving circuits with multiple drivers and coupling capacitors, consisting of simple trees coupled to each other via series capacitors, which circuit model we shall call a *coupled* tree (Fig. 3.2). General moment matching techniques can of course be applied to solve coupled trees, but again simplified techniques are necessary for use early in the design flow. Timing analysers often use the concept of worst-case, average and best-case delay, using a switch factor that takes the value of 2, 1 or 0 to modify the Elmore delay. The capacitance for a line is modelled as the sum of two components, one of which represents the capacitance to ground, and the other the capacitance to adjacent nets. This second component is multiplied by a factor depending on whether the coupled net is expected to be quiet or not, and if not, on the direction of switching. This method of modelling is not accurate except in certain very simple situations, such as uniform structures or simultaneously switching nets, and indeed was recently shown to not even represent an upper bound on the delay [Kahng00]. A lot of research has focused on certain simplified configurations of interest. In [Kawaguchi98] the authors use the first moment of the impulse response to generate single-pole responses for uniformly coupled RC lines, while [Kahng99] presents a two-pole response for a one-section coupled π circuit with arbitrary ramp inputs. They extend it to accommodate multiple segmented aggressors in [Kahng01], but the allowed topology is still very limited.

3.2.2 Noise Modelling

Now as mentioned, it is often necessary to know the coupled noise amplitude explicitly, to check for spurious errors caused by switching nets disturbing the logic state of a quiescent net. A single-pole noise metric for coupled trees was proposed in [Devgan97]. Although computationally efficient, some simplifying assumptions in the formulation of the metric cause the results to be mostly very pessimistic. Some of the works mentioned above which present models for estimating the effect of noise on delay also report noise metrics [Acar99], [Kawaguchi98] and [Kahng01]. In

[Takahashi01] the authors use circuit transformations to simplify a general tree to a 2- π model when analytic formulae can be used, but intermediate steps require the calculation of admittances at each branch point and the estimation of equivalent capacitances which increase run time and impact on the accuracy respectively. In [Vittal99], a two pole waveform is analysed to derive expressions for cross-talk noise amplitude and pulse width.

When dealing with multiple driver systems such as depicted in Fig. 3.2, the concept of superposition is very useful, as the coupled RC network is a linear system. The effect of multiple aggressors switching at different times can be estimated by considering one input at a time with all other inputs grounded, and then adding up the individual waveforms. The authors of [Tong00] and [Chen02] adopt such a methodology, where an attempt is made to generate transfer functions from each driver to the receiver. However the only concession to different switching events (and hence different charging paths) is calculating a unique zero; the poles of the transfer function for all switching events are the same, and are the two lowest frequency poles of the system. These poles are estimated from the methodology proposed in [Cochrun73], which gives closed form expressions for the poles of systems with storage elements, and is a technique that has long been used in analog design to estimate the bandwidth of amplifiers. However using the same two lowest frequency poles in all of the transfer functions can result in large errors, as the *significant* poles which determine the response for different switching events can be far apart on the frequency axis. The reason is that though these poles are part of the natural response of the system, and hence do appear in the transfer function from each driver to the receiver, there will *always* be *partial* pole-zero cancellation in systems which have signal paths with widely differing time constants. Since the transfer function is limited to two poles, it is important that for each path the *two-pole-one-zero model that best fits that particular charging path* is calculated.

3.3 Modelling the System Transfer Function

Consider Fig. 3.2 which shows an arbitrary network comprising a victim net and several aggressors coupled to the victim net through banks of series capacitances. Such a network can be represented by an m -input-single-output system as shown. The waveform at the node of interest e , on the victim is defined as $V_e(t)$. The total waveform at e can always be represented by the n^{th} order linear differential equation (3-1) where n is the order of the system.

Setting the right hand side to zero results in the homogeneous equation, the solution to which gives the natural or transient response of the circuit. For a second order approximation, the complementary equation is as shown in (3-2).

$$\begin{aligned}
& a_n \frac{d^n}{dt^n} V_e(t) + a_{n-1} \frac{d^{n-1}}{dt^{n-1}} V_e(t) + \dots + a_1 \frac{d}{dt} V_e(t) + V_e(t) \\
& = b_1 u_1(t) + b_2 u_2(t) + \dots + b_m u_m(t)
\end{aligned} \tag{3-1}$$

$$a_2 s^2 + a_1 s + 1 = 0 \tag{3-2}$$

Assuming that the roots (which are always real and negative for an *RC* tree) are $s = -1/\tau_1$ and $s = -1/\tau_2$, the complete response of the circuit is given by the following two time constant model where $f(t)$ is the particular solution corresponding to the forcing functions.

$$V_e(t) = A e^{-t/\tau_1} + B e^{-t/\tau_2} + f(t) \tag{3-3}$$

Since the inputs are always aperiodic¹ for the case under consideration, $f(t)$ or the steady state value, is always zero or one (for normalized supply rails). The coefficients A and B will depend on the inputs. In the methodology proposed, linear superposition is used where the response for each input is considered with all other inputs grounded, and all those responses are summed up to generate the complete solution (as in all moment-based approaches). Now in general, all the natural poles of the system contribute to the step response for any switching event where the other inputs are grounded, but their relative contribution varies greatly according to the zeros for a particular switching event. For example, with reference to Fig. 3.2, if the victim tree, and say tree A have parasitics that are much higher than those of the other trees, the lowest frequency poles that *significantly* contribute to the response for the events of the victim driver and the driver of A switching are situated much closer to the origin than those for the events of the other drivers switching. That is because there will always be *partial* pole-zero cancellation of those lowest frequency poles which reflects the smaller parasitics in the other trees, for the switching of those drivers. If the same poles are used, the results will be skewed by the highest parasitics in the coupled tree, regardless of their influence on the particular switching event. Hence in general, for each switching event, unique poles have to be considered to achieve acceptable accuracy.

1. Aperiodic in the sense that the clock period is (much) greater than the settling times.

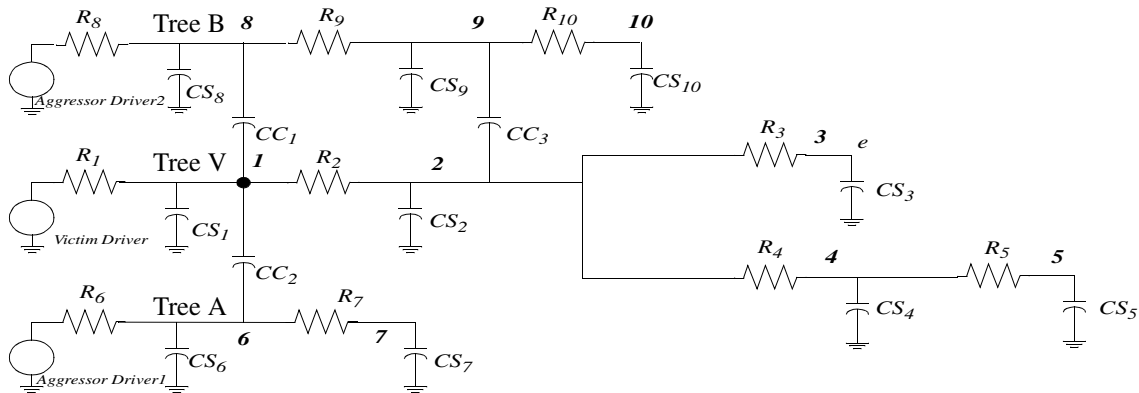


Figure 3.3: Example coupled RC tree for explanation of notation and derivation of expressions

3.3.1 Response to Different Switching Events

A coupled RC tree is characterized by a resistive path from the output node e to the forcing (victim) driver, and series capacitive elements to other (aggressor) drivers. Hence the output for the victim driver switching will always change rails, while it will start and end at the same rail for an aggressor switching. Because of this fundamental difference, the transfer functions characterizing the response to the victim switching and any of the aggressors switching are different. The former will have a zero on the negative part of the real axis:

$$H_v(s) = \frac{1 + s\tau_{z,v}}{(1 + s\tau_{1,v})(1 + s\tau_{2,v})} \quad (3-4)$$

while the latter will have a zero at the origin.

$$H_{a_i}(s) = \frac{s\tau_{z,a_i}}{(1 + s\tau_{1,a_i})(1 + s\tau_{2,a_i})} \quad (3-5)$$

3.4 Calculation of Moments

In all of the following descriptions the tree of Fig. 3.3 can be referred to as an example tree.

3.4.1 Notation

CS_k^p = capacitance to ground at node k in pth tree

$CC_{kj}^{pq} = C_{kj}^{pq}$ = capacitance between node k on pth tree and node j on qth tree where first sub(super) script refers to reference tree

C_k^p = total capacitance at node k on pth tree

V_k^p = voltage at kth node on pth tree

V_{kj}^{pq} = voltage between node k on pth tree and node j on qth tree

V_k^{pq} = voltage between node k on pth tree and corresponding coupled node on qth tree; i.e. 2nd subscript is omitted as it is a more convenient notation when permissible

R_{ke}^p = shared resistance from source to nodes e and k on tree p

Υ_k^n = nth moment of the impulse response at the kth node

It should be noted that superscripts always refer to trees while subscripts always refer to nodes, except in the definition for moments, where the superscript refers to the order of the moment. Additionally, rail voltages are normalized to 0 and 1, and the expressions always derived for a positive step without loss of generality. For negative transitions, the waveforms are simply mirrored. The usage of the notation is illustrated in Fig. 3.4.

Where it is possible to do so without introducing ambiguity, the second subscript will be dropped for convenience. For example, if tree I is the reference tree in Fig. 3.4, cdv_k^{I2}/dt refers to $i(t)$, and node j is implicit in the expression.

Finally, the following quantity is defined:

$$\tau_{D_e}^{t_r t_i} = \sum_{k \in t_r} R_{ke}^{t_r} C_k^{t_r t_i} \quad (3-6)$$

This is the summation over the reference tree t_r , of resistance capacitance products at each node k , where R_{ke} is the shared resistance between node k and sink e , on the path from source to sink. The capacitance term $C_k^{t_r t_i}$ is the capacitance between trees t_r and t_i at node k on t_r . For example with reference to Fig. 3.3, $C_1^{v_b}$ is CC_1 . If the second tree t_i is omitted, the capacitance refers to the total capacitance at node k ; for example, C_1^v is $(CS_1 + CC_1 + CC_2)$. In that case, the second tree would also be omitted in the name, i.e. $C_k^{t_r}$ would be with respect to $\tau_{D_e}^{t_r}$. This notation is used because it makes for a compact description, and also to make it consistent with that adopted in [Horowitz84], which describes second-order models for simple trees. The lower case subscript in τ_{D_e} , which is e in this case, always refers to the output. If the output node is omitted, the only quantity which is with respect to the output, R_{ke} , becomes R_{kk} .

3.4.2 Switching of Victim Driver

When the victim driver switches while all other inputs are grounded, the first moment of the impulse response at the output node e is defined in (3-7):

$$Y_{e,v}^1 = \int_0^{\infty} t h_e^v(t) dt \quad (3-7)$$




Figure 3.4: Illustration of usage of notation

Now the following expression describes the voltage drop from the source to e where $V_e^v(t)$ is the step response at e , and $a_1, a_2 \dots$ are the aggressors. This is obtained by summing up the capacitor currents and adding the drops across each resistor, or in other words, using Kirchoff's voltage and current laws.

$$1 - V_e^v(t) = \sum_{k \in v} R_{ke}^v \left[CS_k^v \frac{dV_k^v}{dt} + CC_k^{va_1} \frac{dV_k^{va_1}}{dt} + CC_k^{va_2} \frac{dV_k^{va_2}}{dt} + \dots \right] \quad (3-8)$$

The impulse response $h_e^v(t)$ is the first time derivative of the step response. Hence (3-7) can be integrated by parts, and (3-8) substituted in it to yield the following expression:

$$\Upsilon_{e,v}^1 = \sum_{k \in v} R_{ke}^v \left[CS_k^v + CC_k^{va_1} + CC_k^{va_2} + \dots \right] = \tau_{D_e}^v \quad (3-9)$$

The second moment of the impulse response at the output node e is given by:

$$\Upsilon_{e,v}^2 = \int_0^{\infty} t^2 h_e^v(t) dt \quad (3-10)$$

Integrating by parts, it can be shown that this is equivalent to:

$$\Upsilon_{e,v}^2 = 2 \int_0^{\infty} t(1 - V_e^v(t)) dt \quad (3-11)$$

Using expression (3-8) for the step response and again integrating by parts, this can be shown to be as given in (3-12):

$$\Upsilon_{e,v}^2 = 2 \sum_{k \in v} R_{ke}^v \left[CS_k^v \int_0^1 t dV_k^v + CC_k^{va_1} \int_0^1 t dV_k^{va_1} + CC_k^{va_2} \int_0^1 t dV_k^{va_2} + \dots \right] \quad (3-12)$$

The constituent integrals can be evaluated by integrating by parts, and using Kirchoff's laws to obtain expressions for the voltages. The first integral is basically the first moment at node k , for which an expression can be obtained by simply substituting k for e in (3-9).

The other integrals are of the form:

$$I_{a_i} = \int_0^1 t dV_k^{va_i} \quad (3-13)$$

Integrating by parts, this simplifies to:

$$I_{a_i} = \int_0^\infty [1 - V_k^{va_i}(t)] dt \quad (3-14)$$

The voltage can be decomposed into two components thus:

$$V_k^{va_i} = V_k^v - V_j^{a_i} \quad (3-15)$$

Now circuit laws can be used to obtain expressions for the individual voltages. This first is:

$$1 - V_k^v(t) = \sum_{K \in v} R_{Kk}^v \left[CS_K^v \frac{dV_K^v}{dt} + CC_K^{va_1} \frac{dV_K^{va_1}}{dt} + CC_K^{va_2} \frac{dV_K^{va_2}}{dt} + \dots \right] \quad (3-16)$$

and the second as given in (3-17) where the superscripts $a_i b_1, a_i b_2 \dots$ in the CC terms indicate the coupling capacitances to tree a_i 's own aggressors.

$$V_j^{a_i}(t) = \sum_{K \in a_i} R_{Kj}^{a_i} \left[CS_K^{a_i} \frac{dV_K^{a_i}}{dt} + CC_K^{a_i v} \frac{dV_K^{a_i v}}{dt} + CC_K^{a_i b_1} \frac{dV_K^{a_i b_1}}{dt} + CC_K^{a_i b_2} \frac{dV_K^{a_i b_2}}{dt} + \dots \right] \quad (3-17)$$

Considering the fact that all nodes not on the victim tree start and end at the same voltage, this simplifies to:

$$\begin{aligned}
Y_{e,v}^2 = & 2 \sum_{k \in v} R_{ke}^v \left\{ CS_k^v \sum_{K \in vic} R_{Kk}^v (CS_K^v + CC_K^{va_1} + CC_K^{va_2} + \dots) \right. \\
& + CC_k^{va_1} \left[\sum_{K \in a_1} R_{Kj}^{a_1} CC_K^{a_1v} + \sum_{K \in v} R_{Kk}^v (CS_K^v + CC_K^{va_1} + CC_K^{va_2} + \dots) \right] \\
& \left. + CC_k^{va_2} \left[\sum_{K \in a_2} R_{Kj}^{a_2} CC_K^{a_2v} + \sum_{K \in v} R_{Kk}^v (CS_K^v + CC_K^{va_1} + CC_K^{va_2} + \dots) \right] + \dots \right\}
\end{aligned} \tag{3-18}$$

This can be expressed in more succinct form by using (3-6):

$$Y_{e,v}^2 = 2 \sum_{k \in v} R_{ke}^v \left\{ C_k^v \tau_{D_k}^v + CC_k^{va_1} \tau_{D_j}^{a_1v} + CC_k^{va_2} \tau_{D_j'}^{a_2v} + \dots \right\} = 2 \left(\tau_{G_e}^v \right)^2 \text{ say} \tag{3-19}$$

3.4.3 Switching of Aggressor Driver

Following an approach identical to that in the former case, the first moment of the impulse response at node e on the victim tree for aggressor a_i switching can be shown to be:

$$Y_{e,a_i}^1 = - \sum_{k \in v} R_{ke}^v CC_k^{va_i} = -\tau_{D_e}^{va_i} \tag{3-20}$$

The second moment can also be calculated from an approach similar to the former case, resulting in:

$$Y_{e,a_i}^2 = -2 \sum_{k \in v} R_{ke}^v \left\{ C_k^v \tau_{D_k}^{va_i} + CC_k^{va_i} \tau_{D_j}^{a_i} \right\} = -2 \left(\tau_{G_e}^{a_i} \right)^2 \text{ say} \tag{3-21}$$

The expressions in (3-9), (3-19), (3-20) and (3-21) along with another first order metric, form the basis of the proposed models. The additional metric, the sum of the open-circuit time constants with respect to the victim driver will be introduced later.

3.5 Matching Moments to Characteristic Time Constants in Circuit

Now the interest is in generating the best two-pole-one-zero transfer function for the response at the output node for any given switching event. The moments can be matched to the characteristic time constants in the circuit by considering the power series expansion of e^x in the definition of the Laplace transform. The Laplace transform of the impulse response is:

$$\begin{aligned}
 H(s) &= \int_0^{\infty} h(t) e^{-st} dt \\
 &= \int_0^{\infty} h(t) \left[1 - st + \frac{s^2}{2!} t^2 - \dots \right] dt \\
 &= \int_0^{\infty} h(t) dt - s \int_0^{\infty} t h(t) dt + \frac{s^2}{2} \int_0^{\infty} t^2 h(t) dt - \dots
 \end{aligned}$$

From this equality, it can be seen that the n^{th} moment of the impulse response is equal to $(-1)^n$ times the n^{th} derivative of the transfer function evaluated at $s=0$:

$$Y^n = (-1)^n \left. \frac{d^n}{ds^n} H(s) \right|_{s=0} \quad (3-22)$$

This identity can be used to match the moments to the poles and zeroes of the circuit directly. Using (3-4), (3-22), (3-9) and (3-19) it can be seen that:

$$\tau_{1,v} + \tau_{2,v} - \tau_{z,v} = \tau_{D_e}^v \quad (3-23)$$

$$(\tau_{1,v} + \tau_{2,v} - \tau_{z,v})(\tau_{1,v} + \tau_{2,v}) - \tau_{1,v}\tau_{2,v} = \left(\tau_{G_e}^v\right)^2 \quad (3-24)$$

Now additional information is necessary to solve for the three unknowns in (3-23) and (3-24). If the reciprocal pole sum is designated as τ_{sum} , these two equations can be combined to form the following quadratic, which yields two time constants:

$$\tau^2 - \tau_{sum}\tau + \tau_{D_e}^v\tau_{sum} - \left(\tau_{G_e}^v\right)^2 = 0 \quad (3-25)$$

Other than τ_{sum} , the other metrics in the equation, the first and second moment, are with reference to the victim. At this point, it is helpful to look at the physical interpretation of the first and second moments of the impulse response. The first moment always considers resistances of the switching line, and either all capacitances connected to the switching line (in the case of the victim driver switching) or capacitances connecting it to a particular line (for the switching of an aggressor driver). The second moment propagates outwards another level, and considers the resistances and capacitances of immediately adjacent lines as well. This intuition is valuable in generating a solution with minimum computational complexity; namely, equation (3-25) can be used to generate the pole time constants for *all* switching events, by using the appropriate reciprocal pole sum.

3.5.1 Guaranteeing Stability

Now first, since (3-25) can in general yield complex poles or a positive pole, some care is necessary to ensure stability. Potential instability can take one of two forms: if the sign under the radical in the solution for the roots of (3-25) is negative, complex poles can result; if the magnitude of the square root is greater than the reciprocal pole sum, a negative time constant results. Using these as limiting conditions, a methodology that always yields stable and accurate results can be formulated. The time constants are:

$$\tau_{1,2} = \tau_{sum} \pm \sqrt{\tau_{sum}^2 - 4\left[\tau_{sum}\tau_{D_e}^v - \left(\tau_{G_e}^v\right)^2\right]} \quad (3-26)$$

One limiting condition is that the sign under the radical should be positive. This leads to:

$$\tau_{sum}^2 + 4 \left[\left(\tau_{G_e}^v \right)^2 - \tau_{sum} \tau_{D_e}^v \right] > 0 \quad (3-27)$$

This inequality is satisfied if:

$$\left(\tau_{G_e}^v \right)^2 > \tau_{sum} \tau_{D_e}^v \quad (3-28)$$

However, this would violate the second condition, which is that the magnitude of the square root should be greater than the reciprocal pole sum:

$$\tau_{sum} > \sqrt{\tau_{sum} - 4 \left[\tau_{sum} \tau_{D_e}^v - \left(\tau_{G_e}^v \right)^2 \right]} \quad (3-29)$$

If (3-28) is true, (3-29) will never be true. Hence the stability condition has to be more stringent. It can be guaranteed that (3-29) is true if the following holds:

$$\left(\tau_{G_e}^v \right)^2 < \tau_{sum} \tau_{D_e}^v$$

or:

$$\tau_{sum} > \left(\tau_{G_e}^v \right)^2 / \tau_{D_e}^v \quad (3-30)$$

That is to say, the reciprocal pole sum must be large enough. However, when (3-30) is fulfilled, the second term in (3-27) is negative. Rewriting the Left Hand Side of it gives:

$$LHS = \tau_{sum}^2 - 4 \tau_{D_e}^v \tau_{sum} + 4 \left(\tau_{G_e}^v \right)^2 \quad (3-31)$$

The function designated *LHS* is a quadratic in τ_{sum} . By considering the first and second derivatives, this parabola can be shown to have a minimum at $2\tau_{D_e}^v$. The zero-crossing points are given by:

$$\tau_{sum} = 2 \left[\tau_{D_e}^v \pm \sqrt{(\tau_{D_e}^v)^2 - (\tau_{G_e}^v)^2} \right] \quad (3-32)$$

Obviously, both of these points are on the right hand side of the vertical axis. Hence we have the shape of the parabola (Fig. 3.5). Now first, if the sign under the radical in (3-32) is negative, its roots are complex, or in other words *LHS* will never become negative and (3-27) is always true. Hence for potential instability to occur, the following must always be true:

$$(\tau_{D_e}^v)^2 > (\tau_{G_e}^v)^2 \quad (3-33)$$

Now it can be proved that the line corresponding to the equality of (3-30) should appear to the left of the first zero-crossing as shown in the figure. Since (3-33) has to be true for potential instability to occur, where

$$\tau_{D_e}^v > 0 \quad (\tau_{G_e}^v)^2 > 0$$




Figure 3.5: Variation with τ_{sum} of quadratic which determines stability

the following is true:

$$4\left(\left(\tau_{D_e}^v\right)^2 - \left(\tau_{G_e}^v\right)^2\right) + \left(\tau_{G_e}^v\right)^4 / \left(\tau_{D_e}^v\right)^2 > 4\left(\left(\tau_{D_e}^v\right)^2 - \left(\tau_{G_e}^v\right)^2\right)$$

i.e.:

$$\left(2\tau_{D_e}^v - \left(\tau_{G_e}^v\right)^2 / \tau_{D_e}^v\right)^2 > \left(2\sqrt{\left(\tau_{D_e}^v\right)^2 - \left(\tau_{G_e}^v\right)^2}\right)^2$$

when:

$$2\tau_{D_e}^v - \left(\tau_{G_e}^v\right)^2 / \tau_{D_e}^v > 2\sqrt{\left(\tau_{D_e}^v\right)^2 - \left(\tau_{G_e}^v\right)^2}$$

Rearranging the terms results in:

$$\left(\tau_{G_e}^v\right)^2 / \tau_{D_e}^v < 2\left[\tau_{D_e}^v - \sqrt{\left(\tau_{D_e}^v\right)^2 - \left(\tau_{G_e}^v\right)^2}\right]$$

Hence the equality of (3-30) is always to the left of the first zero-crossing point of *LHS*.

Then for stability, τ_{sum} has to appear in the lightly hatched area, or to the right of the second zero-crossing point. If τ_{sum} is too small, the sign under the radical is positive, but we end up with one negative time constant. If τ_{sum} is situated between the zero-crossing points, we get complex poles. Finally if τ_{sum} is to the right of the second zero-crossing point, represented by the darkly hatched area, again a stable solution results. Hence from the zero-crossing points, the next condition is obtained:

$$\begin{aligned} \tau_{sum} < 2\left[\tau_{D_e}^v - \sqrt{\left(\tau_{D_e}^v\right)^2 - \left(\tau_{G_e}^v\right)^2}\right] \quad \text{or} \\ \left(\tau_{sum} > 2\left[\tau_{D_e}^v + \sqrt{\left(\tau_{D_e}^v\right)^2 - \left(\tau_{G_e}^v\right)^2}\right]\right) \end{aligned} \quad (3-34)$$

3.5.2 Switching of Victim Driver

Now the stability conditions have been identified, the values for τ_{sum} that give the best response for different switching events can be derived. Firstly, for the case of the victim driver switching, since all aggressors are grounded the metric that gives the best solution is the sum of the open circuit time constants with reference to the victim driver, which shall be defined as τ_p^* . This is simply the summation of the products of all capacitors connected to the victim line with the driving point resistance to each of those capacitors:

$$\tau_p^* = \sum_{k \in v} \left[R_{kk}^v C S_k^v + \left(R_{kk}^v + R_{jj}^{a_1} \right) C C_k^{va_1} + \left(R_{kk}^v + R_{jj}^{a_2} \right) C C_k^{va_2} + \dots \right] \quad (3-35)$$

Again, using expression (3-6), this can be simplified to:

$$\tau_p^* = \tau_D^v + \tau_D^{a_1v} + \tau_D^{a_2v} + \dots \quad (3-36)$$

This is a good approximation for the sum of the pole time constants [Cochrun73], giving:

$$\tau_{1,v} + \tau_{2,v} = \tau_p^* \quad (3-37)$$

Substituting (3-37) for τ_{sum} in (3-23) and (3-26) result in the zero time constant, and pole time constants respectively, for the victim switching. Now an inspection of (3-20) and (3-35) shows that $\tau_p^* > \tau_{D_e}^v$. Since (3-33) has to be true for instability to occur, this means that:

$$\tau_p^* > \left(\tau_{G_e}^v \right)^2 / \left(\tau_{D_e}^v \right)^2 \quad (3-38)$$

Therefore (3-30) is always true, and the only possible stability violation in this case is (3-34); i.e. very occasionally, using τ_p^* can result in complex poles. The physical interpretation of such an occurrence is that the sum of the open circuit time constants underestimates the reciprocal pole sum, which has been unusually escalated by an aggressor or aggressors with exceptionally high parasitics. Because both exponential waveforms are either additive or subtractive unlike when an aggressor switches (where

one is additive and the other is subtractive), the higher frequency pole does not have a significant impact. In fact, this form of instability is usually an indication of a very low frequency pole which makes the prediction of the waveform straightforward. The simplest remedy therefore is to consider a single pole response, with the pole time constant being given by $\tau_{D_e}^v$. This results in good accuracy as shall be shown in the results section.

3.5.3 Switching of an Aggressor Driver

Secondly, to solve for the poles and zeros associated with an aggressor switching, (3-5), (3-22), (3-20) and (3-21), are combined to give:

$$\tau_{D_e}^{a_i} = \tau_{z, a_i} \quad (3-39)$$

$$\left(\tau_{G_e}^{a_i}\right)^2 = \tau_{z, a_i}(\tau_{1, a_i} + \tau_{2, a_i}) \quad (3-40)$$

Now the zero time constant is available immediately from (3-39), and dividing (3-40) by (3-39) results in the reciprocal pole sum:

$$\left(\tau_{G_e}^{a_i}\right)^2 / \tau_{D_e}^{a_i} = \tau_{1, a_i} + \tau_{2, a_i} \quad (3-41)$$

The pole time constants can be obtained by substituting (3-41) as τ_{sum} in (3-26). It can be seen from an inspection of the relevant expressions, that potentially either of (3-30) or (3-41) can be violated. The solution without generating extra information about the circuit, is to accept the next best approximation. That is to say, if τ_{sum} is so small that it violates inequality (3-30), the simplest and most logical remedy is to increase τ_{sum} so that it is in the lightly hatched area. When inequality (3-34) is violated, if τ_{sum} is less than the minima, it should be decreased so that it falls into the lightly hatched region; if it is greater than the minima, it should be increased so that it falls into the darkly hatched region. Since the equality will generate coincident poles which is not acceptable, the exact value should be chosen so that it is slightly greater than or less than the equality, which can be achieved with a percentage factor, such as 1%. Using this approach, the values that τ_{sum} should take in the different cases are summarized in Table 3.1.

Table 3.1 Values that τ_{sum} should take

Condition		Value of τ_{sum}
Victim switching	no violation	τ_p
	(3-34) violated	N/A (use a single pole response)
Aggressor switching	no violation	$\left(\frac{a_i}{\tau_{G_e}}\right)^2 / \tau_{D_e}^{a_i}$
	(3-30) violated: $\tau_{sum} < \left(\frac{\tau_{G_e}^v}{\tau_{D_e}^v}\right)^2$	$0.99\left(\frac{\tau_{G_e}^v}{\tau_{D_e}^v}\right)^2 + 0.02\left[\tau_{D_e}^v - \sqrt{\left(\tau_{D_e}^v\right)^2 - \left(\tau_{G_e}^v\right)^2}\right]$
	(3-34) violated: $2\left[\tau_{D_e}^v - \sqrt{\left(\tau_{D_e}^v\right)^2 - \left(\tau_{G_e}^v\right)^2}\right] < \tau_{sum} < 2\tau_{D_e}^v$	$0.01\left(\frac{\tau_{G_e}^v}{\tau_{D_e}^v}\right)^2 + 1.98\left[\tau_{D_e}^v - \sqrt{\left(\tau_{D_e}^v\right)^2 - \left(\tau_{G_e}^v\right)^2}\right]$
(3-34) violated: $2\tau_{D_e}^v < \tau_{sum} < 2\left[\tau_{D_e}^v + \sqrt{\left(\tau_{D_e}^v\right)^2 - \left(\tau_{G_e}^v\right)^2}\right]$	$2.02\left[\tau_{D_e}^v + \sqrt{\left(\tau_{D_e}^v\right)^2 - \left(\tau_{G_e}^v\right)^2}\right]$	

Of the two, (3-30) being violated is by far the more common form of instability. This occurs when the dominant poles for the victim and the particular aggressor are very far apart on the frequency axis. Physically, this translates to a situation where the receiver node is charged extremely rapidly by a very strong aggressor (i.e. through a relatively very small time constant), and decays with a very long tail, dictated by the much larger time constant of the victim. Such behaviour is common for far-end coupling, as shown in Fig. 3.6. The instability in the solution predicted by (3-26) occurs because the pole sum given by (3-41) accurately reflects the high frequency nature of the poles in the aggressor's charging path, but $\tau_{D_e}^v$ and $(\tau_{G_e}^v)^2$ reflect the much lower frequency content of the victim's dominant poles, and the gap is too much to bridge.

The remedy proposed to this situation is to increase the reciprocal pole sum just beyond the threshold of the equality. Now this yields accurate results, because the intention is to generate the *best two-pole-one-zero model*; in other words the poles and zero need not equate to actual poles and zeros of the system, and indeed should differ for a second-order approximation. Using the factor of 1% beyond the threshold which yields coincident poles ensures that both the high and low frequency behaviour is matched.

It must be emphasized that conditions (3-30) and (3-34) are violated infrequently, and when they do, the values proposed above result in a simple yet accurate solution, which requires no extra information. The expressions for the reciprocal pole sum in body rows three to five of Table 3.1 represent the best approximations that guarantee stability when the first choice approximations in rows one and two prove to be incompatible with the quadratic (3-25).

3.6 Physical Basis of the Model

The physical basis of the Elmore delay is that the waveform is estimated by the *area* of the actual waveform, which consists of a sum of exponential terms. Since the integral of an exponential is also an (scaled) exponential, this means that the time constant




Figure 3.6: Far end coupling

of the estimated waveform is the time domain area of the real waveform. The area of the real waveform can be found from the network topology by means of the first moment of the impulse response.

Computing the first *and* second moments of the impulse response of the circuit, and using them to generate a transfer function with two poles and one zero results in the matching of the boundary conditions at time zero and infinity, and geometric properties -namely the area and first moment- of the actual waveform (step response) with the estimated waveform. The boundary conditions are already considered in the particular formulation of the transfer function (i.e. that the waveform starts and ends on a specific rail). Hence matching the first and second moment of the impulse response does not define a unique solution but a family of curves, as a two-pole-one-zero transfer function has three unknowns. The necessary third equation is obtained by matching circuit components to the reciprocal pole sum.

For the switching of the victim driver with the other inputs grounded, the sum of the open circuit time constants provides a good approximation to the reciprocal pole sum, and combining it with the moments of the circuit for the victim driver switching has a straightforward intuitive motivation. For the switching of an aggressor driver, the geometric properties of the actual waveform (via the first and second moments of the impulse response for an aggressor driver switching) are used to obtain the precise *reciprocal pole sum*. Since the quadratic (3-25) obtained from the moments of the impulse response for the victim driver switching contain relevant information about the victim net, combining it with the reciprocal pole sum for an aggressor switching gives a good approximation to the best two-pole-one-zero model. This is a procedure that works for the vast majority of circuits; however some adjustments are necessary to the reciprocal pole sum for certain pathological cases, which was analysed in a systematic manner, resulting in Table 3.1.

3.7 Computational Complexity

3.7.1 Background: Incremental Computation of the Elmore Delay

The Elmore delay as mentioned before, has been and is used extensively as a delay metric in VLSI interconnection circuits modelled by a tree where all capacitors are




Figure 3.7: Simple tree for demonstrating incremental computation property of Elmore Delay

grounded, which is termed a simple tree. Simple trees are characterized by nodes which may have multiple children but only one parent. The Elmore delay is defined as:

$$\tau_{D_e} = \sum_{k \in \text{tree}} R_{ke} C_k$$

Consider the simple tree given in Fig. , where the output node is designated as e . According to the definition, the Elmore delay is:

$$\begin{aligned} \tau_{D_e} = & R_1 C_1 + (R_1 + R_2) C_2 + (R_1 + R_2 + R_3) C_3 \\ & + (R_1 + R_2) C_4 + (R_1 + R_2) C_5 \end{aligned}$$

This can be rearranged so that the expression is in terms of the product of the downstream capacitance and resistance at each node on the path from the source to the sink:

$$\tau_{D_e} = R_1(C_1 + C_2 + C_3 + C_4 + C_5) + R_2(C_2 + C_3 + C_4 + C_5) + R_3(C_3)$$

Now the tree can be traversed once, and the downstream capacitances stored at each node. Hence after one traversal of the complete tree, the computation of the Elmore delay at any node requires only that the path from source to sink for that particular instance be traversed, with the product of the resistance and downstream capacitance at each node being summed up. Because of this property, any changes to the capacitance values at any node in the tree, require only that those changes be propagated upstream of those nodes where the changes took place. This is known as incremental computation, as only those cached values that are stale need to be updated. Any change to a resistance need only be considered when the metric with respect to a particular node is required, and the path from the root to that node is traversed. Incremental computation

bestows considerable savings, and is one of the principal reasons for the popularity of the Elmore delay.

3.7.2 Computational Complexity of Proposed Metrics

Altogether five metrics that depend on the circuit topology are required for the proposed models. The first order metrics are (3-9), (3-20), and (3-36), and the second order metrics are (3-19) and (3-21).

First order metrics

An inspection of the first order metrics (3-9) and (3-20) clearly shows their similarity to the Elmore delay. These can be rearranged so that the expressions are formulated as the sum of the products of resistance and downstream capacitance at each node on the path from source to sink. Because of the extra complexity introduced by the coupling capacitances, it is necessary to keep track of individual coupling capacitances at each node. This can be achieved by caching the sum of the downstream self (or total) capacitances, and the sum of the individual downstream coupling capacitances with associated root information at each node. Hence similar to the Elmore delay, all downstream capacitances are cached from a full tree traversal, and then the output with respect to a particular node e only requires a traversal from the source to e . Also similar to the Elmore delay, any changes to the tree require only that the capacitance changes be propagated to the upstream nodes, resulting in incremental computation being possible.

The final first order metric (3-36), the sum of the open circuit time constants, requires that at each node in the summation, that node should be treated as the output. Since the output node is therefore always defined for a given victim net (unlike in the previous metrics where the output can be any node in the tree), the incremental components of the summation in τ_p^* can be cached along with the downstream capacitance. For example, in Fig. 3.3, node 4 should have CS_5 as downstream self capacitance, and $R_5 \cdot CS_5$ as downstream τ_p^* information. Therefore this metric requires no extra traversals at all, but instead can be computed along with the downstream capacitances. Again, changes to the tree require only that the changes be propagated to upstream nodes.

Second order metrics

The second order metrics require the capacitances at each node be weighted individually by a first order time constant, which is basically expression (3-6) (in one of the

three forms used) for the path defined from the root of the relevant simple tree to the current node, or its coupled counterpart. There are now three issues related to the complexity;

1. How much work is needed to calculate the weights for the original tree?
2. When the weights are known, how much work needs to be done to calculate the second order metrics with respect to a particular node?
3. The third is, how much work needs to be done to recalculate all the weights once a change or changes have been made to the tree?

Calculation of the weights are demonstrated on the victim net of Fig. 3.3. The weights required are different for the two expressions, and also different for types of capacitances (i.e coupling capacitance between two trees, or the total capacitance, at a particular node), but always characterized in a generic sense by the expression (3-6). Hence any technique that works for one will always work for all the weights. For the sake of explanation, let us assume that the weight consists of $\tau_{D_k}^v$ where only self capacitances are considered, and that the weights at nodes 1, 2 are τ_1, τ_2 etc. Then:

$$\tau_1 = R_1(CS_1 + CS_2 + CS_3 + CS_4 + CS_5)$$

and:

$$\tau_2 = R_1(CS_1 + CS_2 + CS_3 + CS_4 + CS_5) + R_2(CS_2 + CS_3 + CS_4 + CS_5)$$

The rest of the metrics are calculated in a similar manner. Now since the weight is always with respect to the root, it is necessary to visit all the nodes once *after* the downstream capacitance information has been stored on the initial pass. (It is useful also, to store the upstream resistance at each node on this pass, so that in future visits to the node, the τ information can be updated instantly as will be shown later.) All weights can be calculated in one pass by using the property that:

$$\tau_{D_n}^v = \tau_{D_m}^v + \tau_{D_{m \rightarrow n}}^v \quad (3-42)$$

where node m is situated on the path between the root and node n . At branch points a depth first traversal of all child branches preserves the linearity of the traversal. Hence

the weights for *all* nodes can be calculated by one full tree traversal once the downstream capacitance information has been stored.

The answer to the second question is straightforward; an inspection of (3-21) and (3-20) shows that the form the outer (second order) summation takes is exactly similar to the inner (first order) summation, which is characterized in a generic way by the expression (3-6). Therefore it is possible to cache the downstream τC information (just as the downstream C information was cached for the first order metrics) and obtain the metrics from the root to a particular node by visiting *only* the nodes along the path from the root to that node.

So far two complete traversals have been necessary, one bottom-up pass to store the downstream capacitance information, and one top-down pass, beginning at the root to store the τ information (and the upstream resistance information, which is necessary later, to minimise computation when changes are made). Now to calculate the second order metric to any node, rearranging the terms in the summation exactly as in the first order calculation allows the downstream τC to be cached in one full traversal. Subsequently, the second order metric to any node can be calculated simply by visiting all the nodes on the path from the root to that node. Again, if an imaginary second order metric is defined to consist only of the self capacitances for simplicity of explanation, the value that would be cached at node 5 on the third (bottom-up) traversal would be $T_5 = \tau_5 \cdot CS_5$, that at node 4 would be $T_4 = T_5 + \tau_4 \cdot CS_4$, and so on.

Hence three full traversals are necessary, one bottom-up traversal to store the downstream capacitance information, one top-down traversal to store the weights, and a final bottom-up traversal to store the downstream τC information. None of these passes can be combined as the necessary order is bottom-up, top-down and bottom-up.

The only remaining question is also the most important; if it is necessary to traverse the entire tree three times each time a change is made, the incremental computation property is lost. However, after a modification to a component, since only the resulting changes in the stored values need to be accounted for, the calculations that required three traversals for the original tree can be accomplished in one traversal. Consider for example that the component value CS_2 is changed to CS_2' . This immediately causes:

1. the downstream capacitance values cached at node 2 and all nodes *upstream* of node 2 to be stale.
2. the cached weight (τ) information at *all* nodes to be stale.
3. the cached downstream τC information at *all* nodes to be stale.

In node 5 for example, the stored downstream capacitance is current (since the changed capacitor is upstream of it), but the weight and downstream τC information is stale. The old weight is:

$$\tau_5 = R_1(CS_1 + CS_2 + CS_3 + CS_4 + CS_5) + R_2(CS_2 + CS_3 + CS_4 + CS_5) + R_4(CS_4 + CS_5) + R_5(CS_5)$$

The new weight is:

$$\tau'_5 = R_1(CS_1 + CS'_2 + CS_3 + CS_4 + CS_5) + R_2(CS'_2 + CS_3 + CS_4 + CS_5) + R_4(CS_4 + CS_5) + R_5(CS_5)$$

The change is:

$$\tau'_5 - \tau_5 = (R_1 + R_2)(CS'_2 - CS_2)$$

Therefore:

$$\tau'_5 = \tau_5 + (R_1 + R_2)(CS'_2 - CS_2)$$

This is simply the change in the capacitance multiplied by the resistance that is upstream of the changed capacitance. This is true of all nodes downstream of node 2. At the nodes upstream of node 2, the capacitance change is multiplied by the upstream resistance from *that* node. Similarly, the downstream τC information can also be calculated and stored. Hence all stale information can be updated by doing a single bottom-up traversal by considering the difference introduced by the change to the component. First the changed component is located, and its upstream resistance which has been stored earlier, $(R_1 + R_2)$ is noted. Now starting from a leaf node, say node 5 for example, a bottom up traversal is initiated, where both the weight information, and the downstream τC information is updated at once. From node 2 upwards, the downstream capacitance also needs to be updated. Hence the original requirement of three passes for the virgin tree has been reduced to a single pass. This principal also applies for resistor changes, and also multiple component changes. That is, the effect of multiple changes can be considered in one pass.

Summary

It was shown that all the metrics have a very simple and small core, which exactly resembles the Elmore delay. The second order expressions can be described as a weighted Elmore delay; each term in the summation is weighted by either (3-9) or (3-20) for that particular node. These are similar to the second moment of the impulse response proposed for simple trees in [Horowitz84]. Just as the models of [Horowitz84] represent the minimum computational complexity for second order models for the class of circuits that were called simple trees, these models represent the minimum complexity for coupled trees. In fact, if the coupling capacitance terms are put to zero (the entire capacitance is lumped into a ground component) the model for the victim tree reverts to the model proposed in [Horowitz84].

One of the major attractions of the Elmore delay is its incremental computational property. This is a very useful feature, and is the mainstay of several interconnect optimisation algorithms. It should be noted that this is independent of the output node. Whatever node is chosen as e in the tree, this hierarchical property holds true. Now since the constituent summations in the proposed metrics have exactly the same form as the Elmore delay, which format is basically that of (3-6), incremental computation is possible for the proposed metrics.

3.8 Explicit Noise Models

So far the chief concern has been the generation of the transfer function, which is the most important aspect of the modelling. Choice of input waveform, driver modelling, and subsequent processing of the waveform depend on the application, and are not covered for the most part. Explicit expressions are however derived for step inputs, which are sufficiently accurate for quite a number of applications. As mentioned in [Bhavnagarwa00], the *50%-50% delay* for ramp inputs is almost independent of rise time.

First, the step response at node e when the victim driver switches is given by:

$$V_e^v(t) = 1 - \frac{\tau_{1,v} - \tau_{z,v}}{\tau_{1,v} - \tau_{2,v}} e^{-t/\tau_{1,v}} - \frac{\tau_{z,v} - \tau_{2,v}}{\tau_{1,v} - \tau_{2,v}} e^{-t/\tau_{2,v}} \quad (3-43)$$

The step response when an aggressor driver switches is:

$$V_e^{a_i}(t) = \frac{\tau_{z, a_i}}{\tau_{1, a_i} - \tau_{2, a_i}} \left(e^{-t/\tau_{1, a_i}} - e^{-t/\tau_{2, a_i}} \right) \quad (3-44)$$

It is not possible to solve a two pole (or higher order) waveform explicitly for the delay at a given threshold. Closed form heuristics for a two pole waveform can be derived - such as in [Vittal99]- but since the complete response for an m driver system will consist of $2m$ exponential waveforms, some iterative procedure needs to be adopted in the general case.

However (3-44) can be solved explicitly for the peak noise, and the time at which it occurs. Equating the first time derivative to zero and doing some trivial algebra results in:

$$t_{pk} = \frac{(\tau_{1, a_i} \tau_{2, a_i})}{(\tau_{1, a_i} - \tau_{2, a_i})} \ln \frac{\tau_{1, a_i}}{\tau_{2, a_i}} \quad (3-45)$$

Substituting (3-45) for t in (3-44) results in the peak noise from the switching of the driver of aggressor a_i :

$$V_{e, pk}^{a_i} = \frac{\tau_{z, a_i}}{\tau_{1, a_i} - \tau_{2, a_i}} \left(e^{-t_{pk}/\tau_{1, a_i}} - e^{-t_{pk}/\tau_{2, a_i}} \right) \quad (3-46)$$

Having this information available through closed-form equations is very useful for efficient implementation of iterative algorithms to solve for delay and peak noise in multiple-time-constant waveforms.

3.9 Results

The proposed metrics were tested on several different test beds which cover a wide range of topologies, by comparing the step response against a circuit simulator, Spectre, and other moment-based models. The other moment-based techniques are the two-pole-one-zero model from three moments described in [Tutuianu96], the *gamma* and *hgamma* probability distribution models from three moments described in [Celik02],

and *gamma* probability distribution model from two moments, also described in [Celik02]. It should be noted that the probability distribution models are valid only for the victim switching.

Shown here are the results pertaining to three which illustrate all the corner cases; the tree of Fig. 3.8 consisting of the victim, three primary aggressors, and three secondary aggressors (representing an arbitrarily-coupled circuit, where inequality (3-34) is violated when solving for the poles of the victim switching), the circuit of Fig. 3.10 with four primary and four secondary aggressors (representing global distributed interconnects) and the circuit of Fig. 3.12, (with far end coupling where inequality (3-30) is violated when solving for the poles of the aggressor switching). Shown in Fig. 3.9 are the waveforms at the receiver node e of the circuit in Fig. 3.8, for each driver switching. It can be seen that the model prediction is very close to the Spectre simulation, and always better than all probability distribution estimations. It has an accuracy comparable to the more expensive three-moment two-pole-one-zero model.

Since the actual and predicted delay at a single threshold can agree very well, and still result in significant deviations along the full waveform, the accuracy was tested at three points along the waveform. For the victim switching, the thresholds are 10%, 50% and 90%, while for the aggressors they are 25%, 100% and 25% of the peak amplitude. This is to ensure that three points, with two being on either side of the peak, are tested. For the aggressors, the error at different thresholds is given as a fraction of the pulse width between the first and last threshold. The waveforms for the circuit of Fig. 3.8 are given in Fig. 3.9, those of Fig. 3.10 in Fig. 3.11 and finally those of Fig. 3.12 are shown alongside.

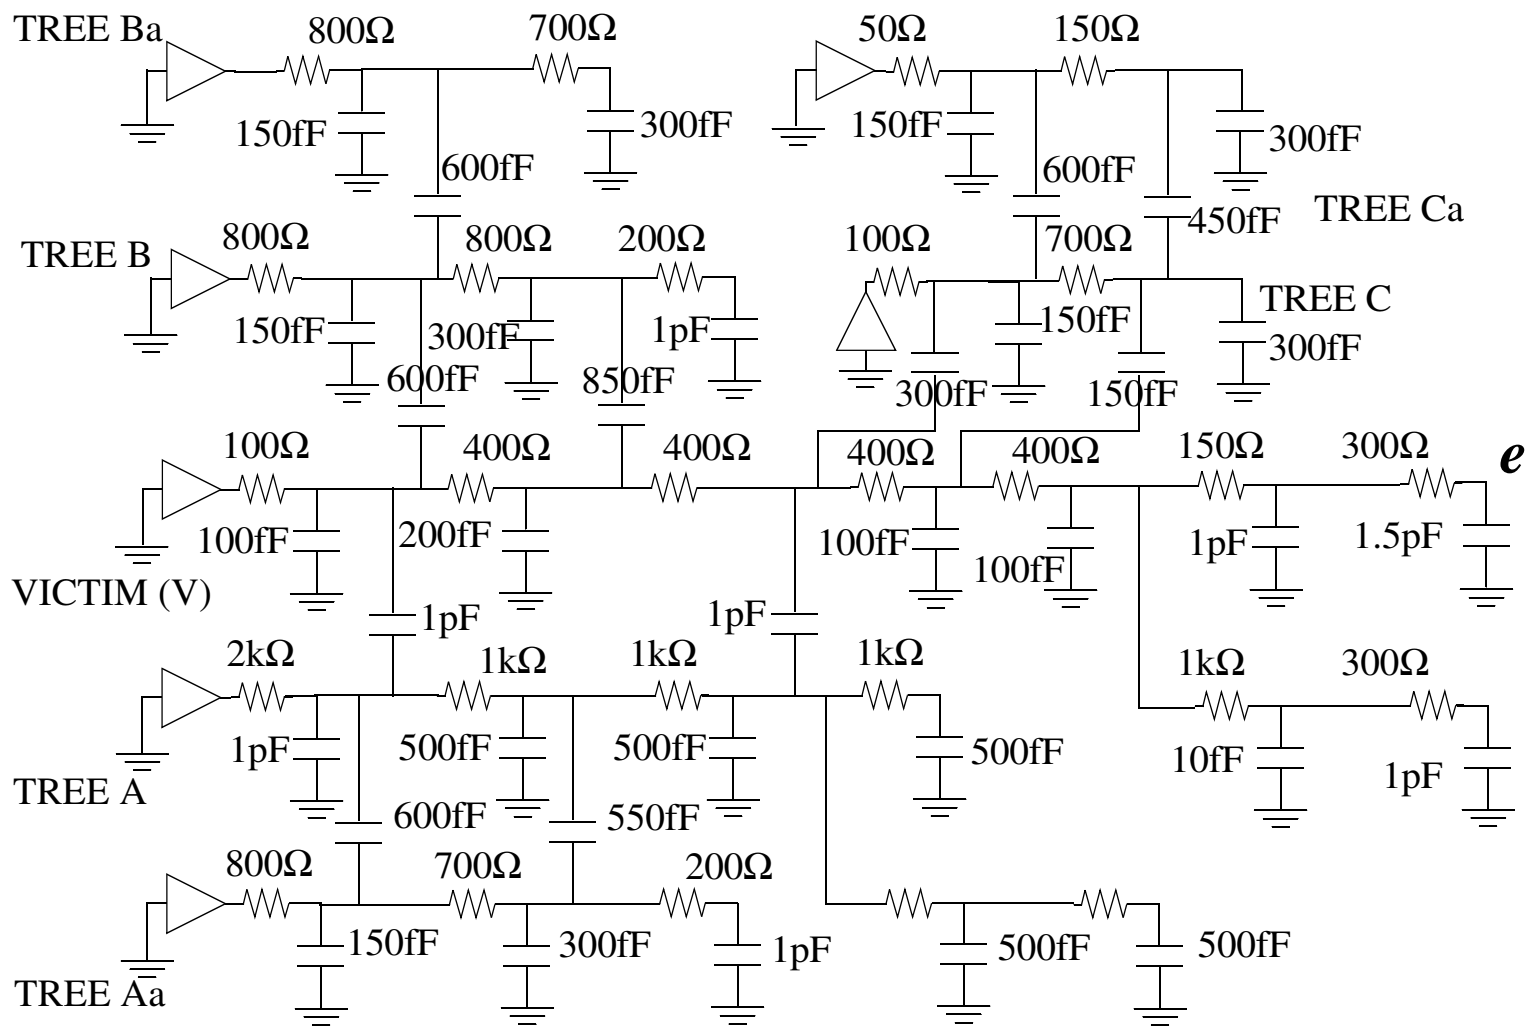


Figure 3.8: Test Bed 1: arbitrarily coupled RC tree

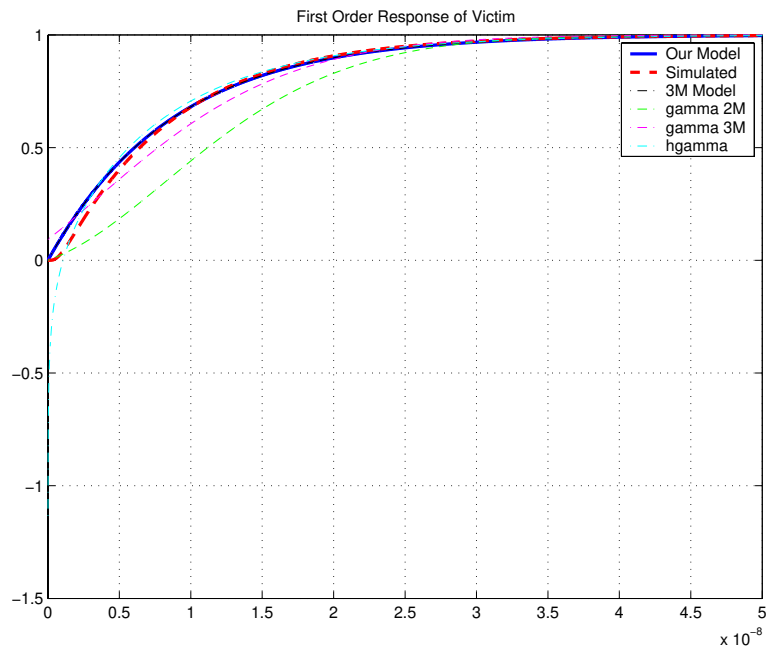


Figure 3.9.a: Response at node e in Test Bed 1 (Fig. 3.8) for the switching of the **Victim** driver




Figure 3.9.b: Response at node e in Test Bed 1 (Fig. 3.8) for the switching of the driver of **Aggressor Tree B**

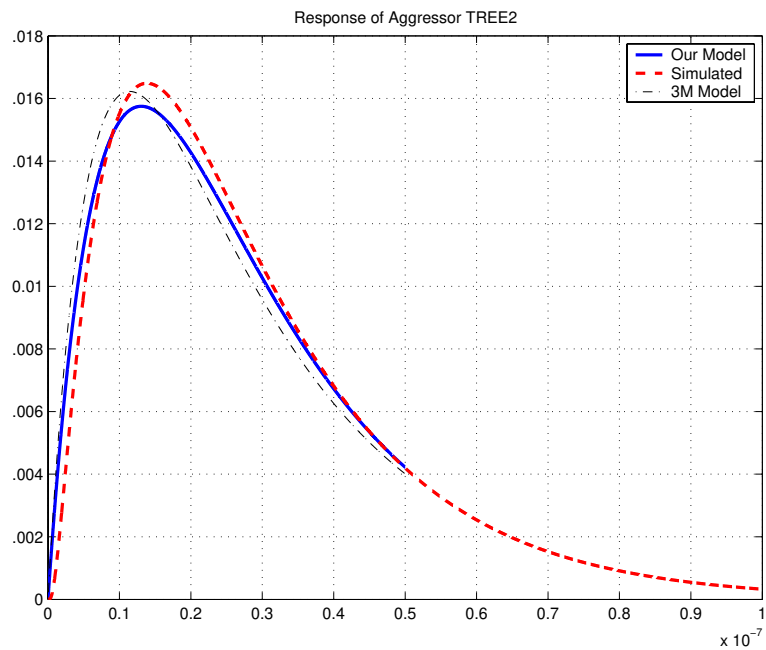


Figure 3.9.c: Response at node e in Test Bed 1 (Fig. 3.8) for the switching of the driver of **Aggressor Tree A**




Figure 3.9.d: Response at node e in Test Bed 1 (Fig. 3.8) for the switching of the driver of **Aggressor Tree C**

Figure 3.9: Waveforms for Testbed 1 (Fig. 3.8)

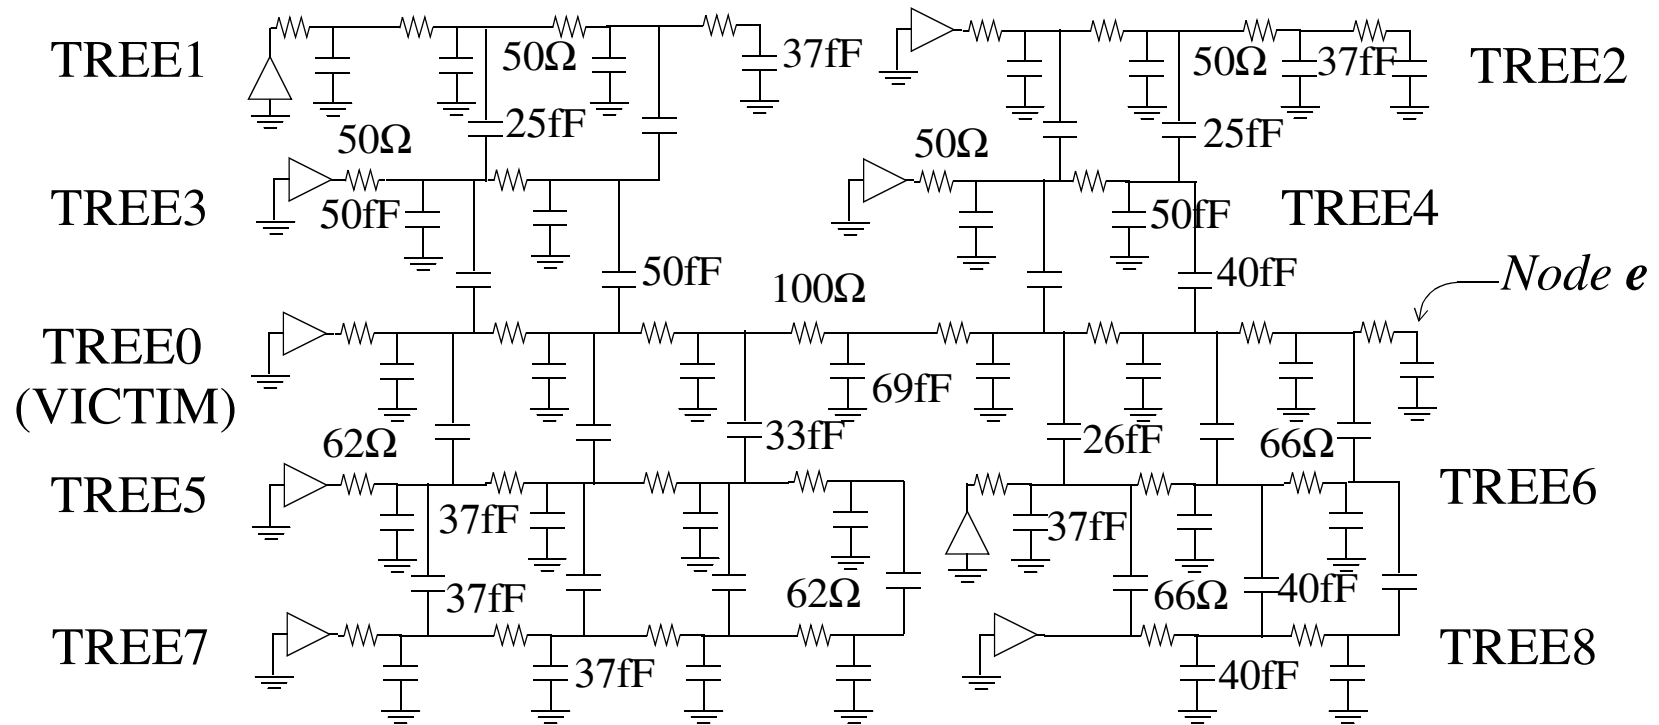


Figure 3.10: Testbed 2: System of distributed coupled interconnect (component values repeated within simple trees)

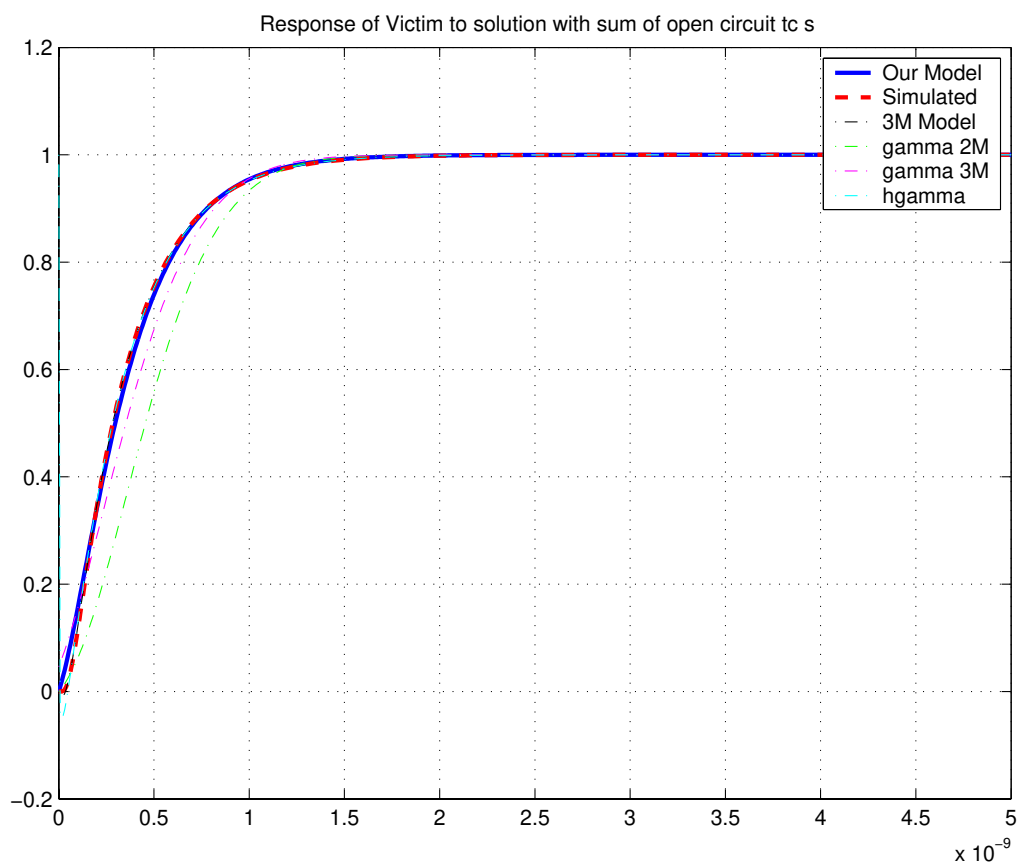


Figure 3.11.a: Response at node e in Test Bed 1 (Fig. 3.10) for the switching of the driver of **Aggressor Tree 3**




Figure 3.11.b: Response at node e in Test Bed 1 (Fig. 3.10) for the switching of the driver of **Aggressor Tree C**

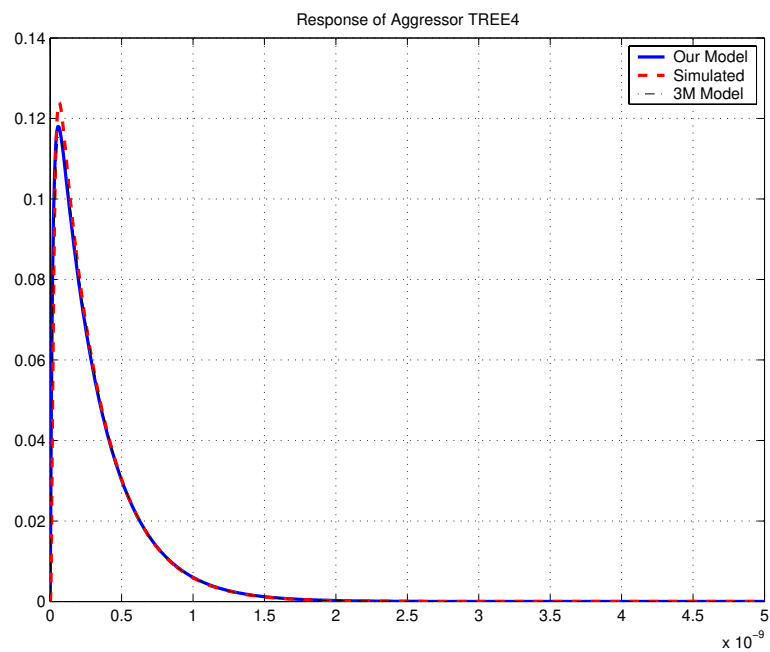


Figure 3.11.c: Response at node e in Test Bed 1 (Fig. 3.10) for the switching of the driver of **Aggressor Tree 4**




Figure 3.11.d: Response at node e in Test Bed 1 (Fig. 3.10) for the switching of the driver of **Aggressor Tree 5**

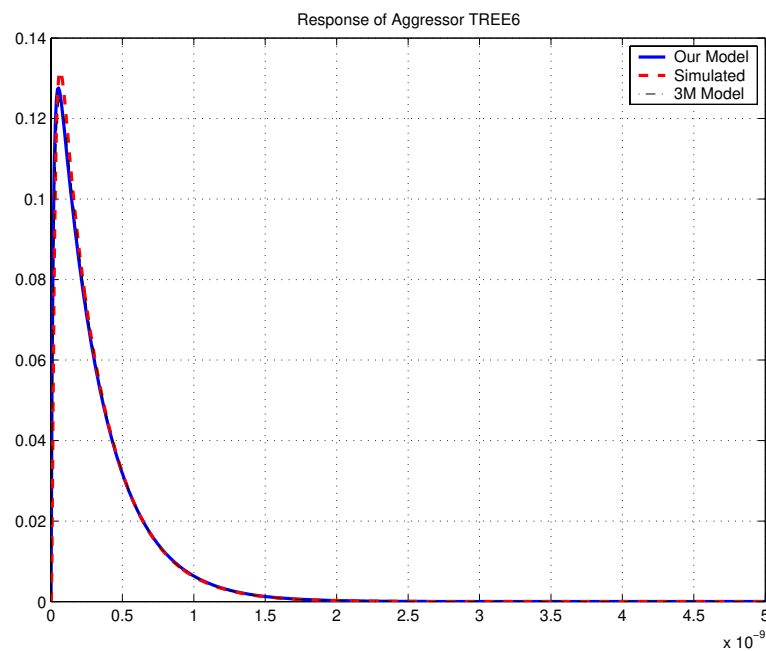


Figure 3.11.e: Response at node e in Test Bed 1 (Fig. 3.10) for the switching of the driver of **Aggressor Tree 5**

Figure 3.11: Waveforms for Testbed 1 (Fig. 3.10)




Figure 3.12.a: Test Bed 3 (component values repeated within simple trees)

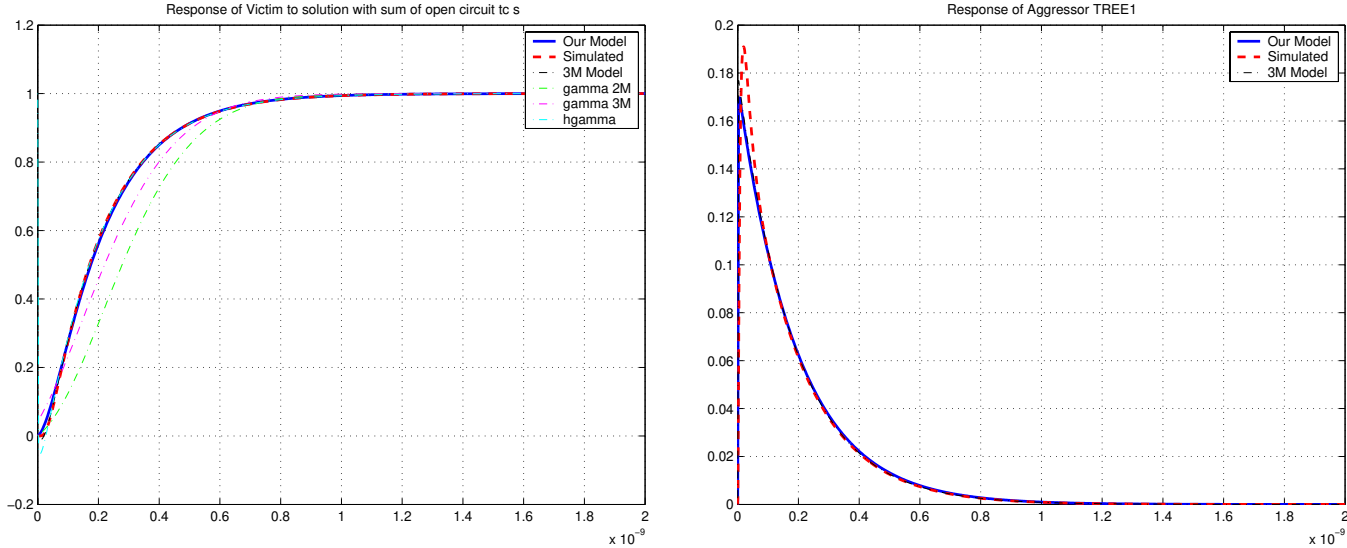


Figure 3.12.b: Output Waveforms

Figure 3.12: Test Bed 3: Circuit with far end coupling

3.10 Summary

Simple yet accurate models for estimating delay and noise are necessary for timing and signal integrity analyses in VLSI circuits in nanometer technologies. For initial analyses early in the design flow cycle, even generic model-order-reduction techniques can present too much of an overhead. In this context, new second-order models for the transfer function for any switching event in general arbitrarily-coupled RC trees with multiple drivers, based on closed form expressions for the first two moments of the impulse response and the sum of the open circuit time constants with respect to the victim driver were proposed. This allows the effect of multiple switching aggressors on a victim to be estimated, with a minimum of computational complexity. The summation of all waveforms result in the complete response at the node of interest.

Since the models have guaranteed stability, they do not compromise on generality, and in fact subsume a lot of models that address simplified structures. For the case of an aggressor driver switching, the methodology for generating the poles can be described as an averaging of the dominant poles along the victim and that aggressor. This procedure provides the best opportunity to match the waveform along the charging and discharging paths.

Two exponential waveforms are the minimum required to model a spike that starts and ends on the same rail. For a system with m drivers, the complete response will potentially consist of $2m$ exponential waveforms. Hence to find the delay at a particular threshold, an iterative procedure has to be adopted; similarly, to find the maximum noise for a given set of switching events with no restriction on aggressors arrival times requires processing of the waveforms; ramp inputs and non-linear modelling of drivers requires some form of convolution in the time domain, or its equivalent in the frequency domain. However the most important aspect in all of this is generating an accurate transfer function from the circuit elements with minimum complexity - which is what has been addressed here. Once the transfer function is known, driver modelling, choice of input waveform, and processing of the final output waveform can be as simple or as detailed as the application warrants.

The point of doing early signal integrity analyses is that where problems are identified, some change in the circuit graph is required. Once the changes are carried out the outputs need to be calculated again, with minimum effort. This is one reason why the Elmore delay is so efficient in interconnect optimization algorithms; because it is possible to model a stub of interconnect by its input capacitance and Elmore delay, changes to the circuit graph require only minimal re-calculations. Because the proposed closed form metrics have an Elmore like flavour, similar stub characterization is pos-

sible for each one; three of the metrics are very similar to the Elmore delay, while the other two have a higher order of complexity but possess a similar form.

When the technique of “path tracing” as described in [Ratzlaff94] -which basically calculates the moments in a hierarchical manner- is adopted, similar savings are possible in the moment computation. However the model proposed here uses one less moment than any other published model, and hence represents a saving on computation. Also, the structure of the expressions describing the metrics extracted from the circuit is such that partitioning of the interconnect system is very easy.

For testing purposes, the models were used to derive the time domain waveform for the step response. For the delay at a given threshold, the accuracy was found to be more than 90% on average, even for complex circuits such as shown in Fig. 3.8. The time at which the peak noise occurs was predicted with even better accuracy. The peak noise itself was predicted with an accuracy of about 85% or higher in general. These figures cannot be claimed as being hard bounds for all possible circuit topologies as it is always possible to create a circuit which is poorly represented by a two pole response. However the models did perform very well when tested over a wide range of circuits that are representative of coupled interconnect structures in nanometer technologies. Also, the accuracy is comparable to and often better than more expensive three-moment models. The simplicity and accuracy of the models combined with their generality should make them useful in delay and noise estimations in complex systems, early in the design flow.

3.11 Limitations and Future Work

In many applications, it is useful to know the upper and lower bounds of a waveform. The methodology presented here only estimates the waveform with no information on bounds. Future work should focus on developing upper and lower bounds for the waveforms induced by switching aggressor drivers.

4. Repeater Modelling

The modelling of repeaters is discussed, including ways of optimising the size to minimise delay and power. Modifications to existing models which allow optimization for coupled nets are considered.

4.1 Introduction

Repeater insertion in VLSI circuits, like delay and noise modelling, has been extensively researched. Historically loads in VLSI circuits have resembled lumped capacitive elements, as wires were relatively short. With increasing die sizes, interconnection lengths have increased to the point that multiple section Π or T models are necessary for accurate modelling of wires. For purposes of reviewing the literature, it is convenient to separate the material along the lines of those that address lumped and distributed models. This chapter looks at proposed methodologies and analysis methods and presents some contributions by the author in the area. It should be noted that the words buffer and repeater are used interchangeably to mean an amplifier, which consists either of an inverter or a pair of cascaded inverters. The distinction is made where necessary.

4.1.1 Background

For lumped capacitive loads, initial analyses modelled a buffer which consisted either of inverters or coupled inverters, with a resistor-capacitor combination [Lin75], [Jaeger75], [Mead80], [Nemes84]. It was established that a horn of repeaters, each larger than its predecessor results in the minimum delay from input to output. Assuming a linear scaling of delay with output capacitance, Mead and Conway in [Mead80] derived the well known condition that sizing each buffer to be e (base of the natural logarithm) times the size of the previous buffer results in an optimal solution with regard to delay. In [Hedenstierna87] the authors consider ramp inputs and also the effect of the intrinsic delay of the buffer which leads to a different optimal tapering factor.

The basic analysis of inserting repeaters in long resistive lines was first presented in [Bakoglu85]. The fundamental idea is that signal propagation on long resistive interconnect lines is a function of the product of the line resistance and capacitance, the RC delay. Since both the resistance and capacitance show a linear increase with length,

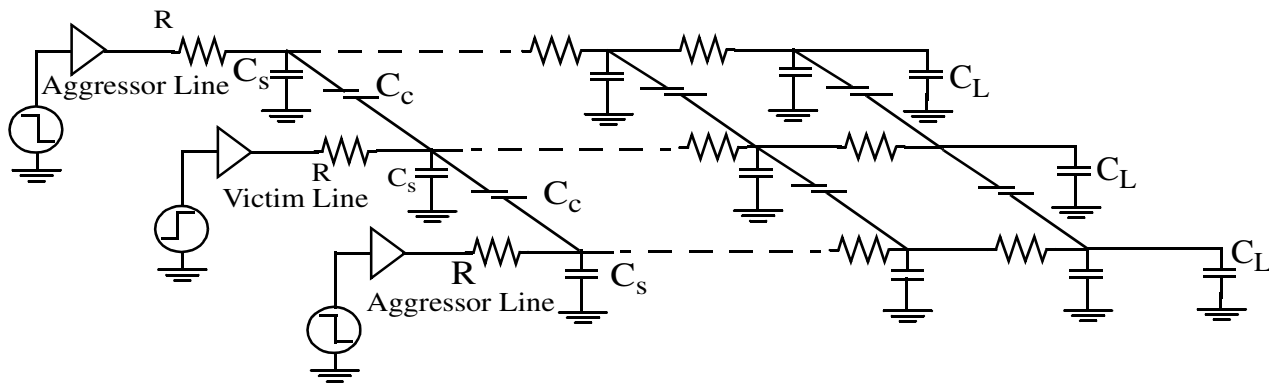


Figure 4.1: True distributed RC line with uniform coupling




Figure 4.2: Repeaters inserted in long uniformly coupled nets to reduce delay

the delay increases quadratically with length. Breaking up the line into smaller segments preserves the overall linearity of the delay. Because the repeaters introduce their own delay, which in turn depends on their driving strength or size in addition to the line load, there is an optimal number and size for minimising delay. Bakoglu and Meindl characterized a repeater by an input capacitance (C_{min}) and output resistance (R_{min}) which was typical of a minimum sized driver for the particular technology, and a number (h) which represented its size in terms of multiples of the minimum sized repeater's W/H ratio. The input capacitance of the particular repeater was then hC_{min} and the output resistance R_{min}/h . This allowed them to neatly quantify the effect of repeater scaling on driving strength and delay, which though a linear model, was and is, extremely useful for system level modelling.

Subsequently researchers have improved on both the resistor-capacitor repeater model and the wire-load model. Wu and Shiau in [Wu90a] and [Wu90b] use a linearised form of the Schichmann-Hodges equations at a particular operating point to model the response of an inverter. In [Dar91] Dhar and Franklin present an elegant mathematical treatment of area constrained optimisation for a two-pin net. The authors of [Nekili92] consider optimal repeater insertion with different inputs and wire models, and propose the use of parallel regenerators in [Nekili93]. In [Adler98] Adler and Friedman use Sakurai's alpha power model [Sakurai90] to include the effect of velocity saturation in short channel devices. Ismail and Friedman in [Ismail00] present an analysis which models inductance in the interconnect for the first time.

Quite a few works have addressed the problem of constrained optimisation of repeaters for a set of interconnected nets. A landmark paper is [Ginneken90] where van Ginneken proposes a hierarchical programming algorithm to find the optimal buffer placement under the Elmore delay model. A large amount of research has concentrated on extensions and modifications to this algorithm, two of which are [Alpert97] and [Lillis96].

In current and future generation of VLSI circuits when the feature size shrinks to a fraction of a micro meter, it is important to consider the effect of cross talk between lines. Cross talk couples a noise voltage onto the victim net, and has an effect on the delay (see Chapter 2 and Chapter 3). In the literature, it is easy to find algorithms which insert buffers to combat both effects [Alpert99], [Menezes99]. They basically iteratively check for noise and delay constraint violations, and insert repeaters where necessary, optimising the placement in the process. The delay calculations in these methodologies are made with the switch factor based models, or higher order numeric models such as AWE described in [Pillage90]. There have been relatively few works which address the issue of driver modelling and optimising the repeater *size* and *number* to combat

the cross-talk on delay effect. One such work is [Sirichotiyakul01] where the authors model the driver with a “transient” resistance which is calculated numerically.

The main interest here is in cross-talk induced delay, and further in a parallel line configuration, where the nets are laid out alongside each other for a relatively long distance as would occur in an intermediate or global level bus. Recently there has been a profusion of research into block level architectures with each block containing 50k to 100k gate modules [Sylvester99a]. These blocks communicate with each other via global level interconnects, either through buses or dedicated links. Regardless of the exact high-level signalling protocol, the parallel net topology in Fig. 4.1 will occur very often.

Capacitive coupling between lines can result in speeding up of the signal or cause delay- depending on the correlation between the data on the different lines. One of the most widely used techniques is to use a coefficient for the coupling capacitance which takes the value of either 1, 0 or 2 to differentiate between a quiet aggressor, an aggressor switching in the same direction or one switching in the opposite direction respectively. Using a similar approach, simple first-order expressions for a variety of switching patterns giving accurate measures of average, best- and worst-case delay for buffered lines are derived. These delay models show how repeater insertion can be optimised to compensate for dynamic effects, and are suitable for initial timing estimates.

4.2 Signal Delay in Long Uniformly Coupled Nets

From now on, delay refers to the 50% delay, since this is the delay to the switching threshold of an inverter. Also in all cases the victim line is assumed to switch from zero to one, without loss of generality. When a line switches up (down) from zero (one) it is assumed to have been zero (one) for a long time. The line model is one with coupling on two sides as shown in Fig. 4.1. The reason is that this is closest to the actual situation for an interconnect in a bus. Now to gain some insight into signal propagation on the distributed line, it is useful first to analyse the lumped model which consists of the *first section* of Fig. 1. For simultaneously switching lines, six different switching scenarios can be identified.

- (a). Both aggressors switch from one to zero
- (b). One switches from one to zero, the other is quiet
- (c). Both are quiet
- (d). One switches from one to zero, the other switches from zero to one

- (e). One switches from zero to one, the other is quiet
- (f). Both switch from zero to one

Consider (c) above as the reference delay, where the driver of the victim line charges the entire capacitance. Cases (a) and (b) slow down the victim line, (d) is equivalent to (c), and (e) and (f) speed up the victim. Now given in (4-1) is the complete response of the victim line.

$$V = 1 + A_i e^{-\frac{t}{R(C_s + 3C_c)}} + B_i e^{-\frac{t}{RC_s}} \quad (4-1)$$

Depending on how the aggressor lines switch, the coefficients A_i and B_i take the values given in Table 4.1.

Table 4.1 Complete response for different switching patterns

i	Switching pattern	A_i	B_i
1	(a)	-4/3	1/3
2	(b)	1	0
3	(c)	-2/3	-1/3
4	(d)	-2/3	-1/3
5	(e)	-1/3	-2/3
6	(f)	0	1

In cases (b) and (f), the response is a single decaying exponential with a time constant of $R(C_s + 3C_c)$, while in the other cases this is the slower time constant. In cases (a), (c) and (d), this slower time constant is also associated with the larger coefficient, and hence becomes a truly dominant time constant. This is especially so in case (a). Typically in current and future submicron technologies with high aspect ratio interconnect, C_c is close to C_s and often greater. The accuracy of the single time constant is compromised only when $C_c \ll C_s$ when there is no need to distribute the capacitance anyway.

Now to state some well known results, a lumped RC circuit has a single pole response and the delay is as given in (4-2).

$$t_{lump} = 0.7RC \quad (4-2)$$

Signal propagation along a distributed RC line is governed by the diffusion equation which does not lend itself readily to closed form predictions for the delay at a given threshold. However it turns out that a simple exponential is a very good predictor [Bakoglu90]. The reason is that a distributed line (which comprises cascaded RC sections in the limit where the number of sections tends to infinity) is a degenerate version of an RC tree. The transfer function in consequence has a dominant pole which can be well approximated by the reciprocal of the first moment of the impulse response. The first moment of the impulse response is $0.5RC$ which leads to (4-3) as the model for the 50% delay of a distributed RC line to a step input.

$$t_{dist} = 0.4RC \quad (4-3)$$

This is a very good approximation and is reputed to be accurate to within 4% for a very wide range of R and C .

In general whenever the response of the lumped model corresponding to a single section of the distributed line is or can be approximated by a waveform containing a single exponential, most of the response of the distributed line can also be approximated by a waveform with a single exponential. Hence the delay of the *distributed* lines corresponding to (a), (b), (c), (d) and (f) can be approximated with single time constant expressions. (In the case of (e) because the lumped model does not have a dominant time constant, the accuracy is not high enough to justify such an approach). Since the time constants in question are linear combinations of R , C_s and C_c , changing coefficients are sufficient to distinguish between the different cases. The delay is as given in (4) where λ_i take the values in Table 1.

$$t_{vic} = 0.4RC_s + \lambda_i RC_c \quad (4-4)$$

These constants were obtained by running sweeps with the circuit analyser SPECTRE. Now the total delay of the line is affected by the driver strength, and the load at the end of the line. The simplest characterization of the driver is to consider it as a voltage source in series with an output resistance R_{drv} , with a capacitive load of C_{drv} at the

Table 4.2 Coefficients of the heuristic delay model for distributed line with different switching patterns

i	Switching pattern	λ_i	μ_i
1	(a)	1.51	2.20
2	(b)	1.13	1.50
3	(c)	0.57	0.65
4	(d)	0.57	0.65
5	(e)	<i>na</i>	<i>na</i>
6	(f)	0	0

input. The linear approximation of the buffers allows the use of superposition to find the delay:.

$$t_{T, vic} = 0.7R_{drv}(C_s + C_{drv} + \mu_i \times 2C_c) + R(0.4C_s + \lambda_i \times C_c + 0.7C_{drv}) \quad (4-5)$$

The lumped resistance R_{drv} combines with all the capacitances (lumped and distributed) to produce delay terms with a coefficient of 0.7. Similarly the distributed resistance of the line combines with various capacitances to produce different delay terms (it is assumed that the load at the end of the line is an inverter which is the same size as the driving inverter). The terms which model cross-talk are shown in bold. The coefficient μ_i is a second empirical constant to model the Miller effect. Together, these two coefficients make the expression for total delay more accurate than using a single coefficient of 2 for the coupling capacitance to model the worst-case. For $i=1$, the above expression reduces to (4-6).

$$t_{T, vic} = 0.7R_{drv}(C_s + 4.4C_c + C_{drv}) + R(0.4C_s + 1.5C_c + 0.7C_{drv}) \quad (4-6)$$

If a universal factor of 2 is used for the coupling capacitance, the expression takes the form given in (4-7).

$$t_{T, vic} = 0.7R_{drv}(C_s + 4C_c + C_{drv}) + R(0.4C_s + 1.6C_c + 0.7C_{drv}) \quad (4-7)$$

Hence with the empirical constants that are proposed, factors of 4.4 and 1.5 appear before C_c , while in a conventional worst-case analysis they would be 4 and 1.6. If the driver impedance is set to zero, the difference between the two expressions is very small, but with non-zero driver impedances, the difference is significant. The accuracy of (4-6) and (4-7) was checked against simulated values, and the results are presented in Table 4.3, which is divided into three sections. The first section has parasitic values that can be said to represent those of global or semi global wires, the second has values that are more typical of narrower wires, while the third has a much wider variation of all three parameters. The values corresponding to R_{drv} and C_{drv} were set to 1k Ω and 0, 3k Ω and 0, and 5k Ω and 100fF for the three sections respectively. The comparison is also plotted in Fig. 4.3. It can be seen that in all cases, the proposed model contains the error to under 5%, while the traditional worst-case model is more sensitive to the value of the driver impedance and has errors of up to 10% for certain cases. It should be noted that this is not just the result of shifting the error so that it is distributed above and below zero, rather than being purely pessimistic or optimistic. A model with an error which varies between +5% and -5% is no better, is indeed worse, than a model with an error between 0 and 10%. Instead, it can be seen that there is a genuine improvement in the accuracy.

4.3 Repeater Insertion

To reduce delay the long lines in Fig. 4.1 are broken up into shorter sections, with a repeater (an inverter) driving each section. Let the number of repeaters including the original driver be K , and the size of each repeater be H times a minimum sized inverter (all lines are assumed to be buffered in a similar fashion). The output impedance of a minimum sized inverter for the particular technology is $R_{drv,m}$ and the output capacitance $C_{drv,m}$ both of which are assumed to scale linearly with size similar to the modelling in [Bakoglu85]. This arrangement is sketched out in Fig. 4.2, where the symbol $\overline{\overline{w}}$ refers to a capacitively coupled interconnect as shown in Fig. 4.1. In general, the line segments corresponding to the gain stages would not be equal in length, as repeat-




Figure 4.2.a: Comparison of empirical and traditional switch factor based analyses; Points correspond to entries 1 through 27 of Table 4.3

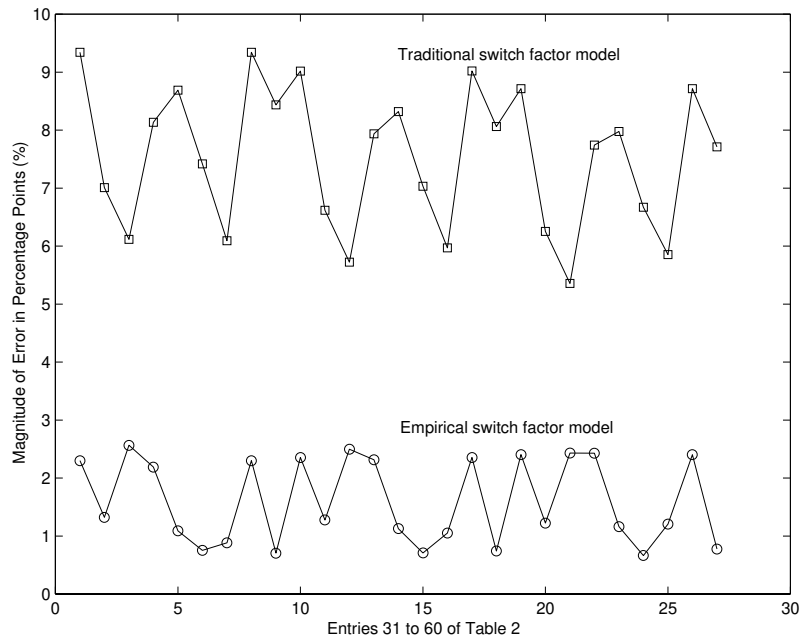


Figure 4.2.b: Comparison of empirical and traditional switch factor based analyses; Points correspond to entries 28 through 54 of Table 4.3

Figure 4.3: Model Verification




Figure 4.2.c: Comparison of empirical and traditional switch factor based analyses; Points correspond to entries 28 through 54 of Table 4.3

Figure 4.2: Model Verification

$$\begin{aligned}
 t_{uneq} = \sum_{i=1}^K \left[0.7 \left(\frac{R_{drv_m}}{h_i} + R_{via} \right) (c_s l_i + H_i C_{drv_m} + \mu_i \times 2c_c l_i) \right. \\
 \left. + r l_i (0.4c_s l_i + \lambda_i c_c l_i + 0.7H_i C_{drv_m}) \right] + \frac{t_r}{2}
 \end{aligned} \tag{4-8}$$

ers are typically situated in “repeater stations”, the locations of which are determined by overall layout considerations. Then the delay is given by (4-8).

It is assumed that the load C_L is equal to the input capacitance of an H sized inverter. Also the signal rise time has been included¹. For the long lossy lines that are considered here, usually the delay of the line is much greater than the rise time of the signal with which the driving inverter is gated, and the 50%-50% delay from buffer input to output

1. This is assuming that zero time is when the driving inverter starts to switch. If zero time is considered to be the point at which the ramp to the first driver starts, the entire rise time should be added.

interconnect node is independent of rise time [Bhavnagarwa00]. Now the minimum delay is obtained when the repeaters are *equalized* over the line, when the above expression reduces to (4-9) where the via resistance R_{via} has been omitted as it is negligible in comparison with R_{drv} .

$$t_{eq} = K \left[0.7 \frac{R_{drv_m}}{H} \left(\frac{C_s}{K} + HC_{drv_m} + \mu_i \frac{2C_c}{K} \right) + \frac{R}{K} \left(0.4 \frac{C_s}{K} + \lambda_i \frac{C_c}{K} + 0.7 HC_{drv_m} \right) \right] + \frac{t_r}{2} \quad (4-9)$$

In order to find the optimum H and K for minimizing delay, the partial derivatives of (4-9) with respect to K and H are equated to zero, resulting in (4-10) and (4-11).

$$K_{i,opt} = \sqrt{\frac{0.4RC_s + \lambda_i RC_c}{0.7R_{drv_m} C_{drv_m}}} \quad (4-10)$$

$$H_{i,opt} = \sqrt{\frac{0.7R_{drv_m} C_s + 1.4\mu_i R_{drv_m} C_c}{0.7RC_{drv_m}}} \quad (4-11)$$

When a number corresponding to a certain case is substituted for i in the two equations, the number and size of repeaters to minimize the delay for that particular switching pattern results. Note that when the coupling capacitance term C_c is set to zero (i.e. the entire capacitance is lumped into the term C_s), (7) and (8) simplify to the Bakoglu equations [Bakoglu90]. Thus it can be seen that this is a simple way to distribute the capacitance and take the effect of switching aggressors into account. These equations and their ramifications for repeater insertion strategies are examined in more detail later.

Table 4.3 Accuracy of delay model with empirical constants measured against Spectre and traditional worst-case model.

R Ω	C _s pF	C _c pF	t _T (sim) ns	t _T (model) ns		Magnitude of error %		R Ω	C _s pF	C _c pF	t _T (sim) ns	t _T (model) ns		Magnitude of error %		R Ω	C _s pF	C _c pF	t _T (sim) ns	t _T (model) ns		Magnitude of error %	
				old	new	old	new					old	new	old	new					old	new	old	new
10	1	0.1	0.992	0.986	1.014	0.65	2.16	100	0.1	0.1	1.180	1.070	1.153	9.34	2.30	10	0.01	0.1	0.0020	0.0018	0.0019	9.55	2.48
10	1	0.2	1.310	1.267	1.323	3.29	0.97	100	0.1	0.3	2.992	2.782	3.031	7.00	1.32	10	0.01	1	0.0153	0.0144	0.0158	5.62	3.54
10	1	0.3	1.639	1.549	1.633	5.53	0.42	100	0.1	0.5	4.787	4.494	4.910	6.11	2.56	10	0.01	10	0.1478	0.1405	0.1545	4.91	4.55
10	1.1	0.1	1.060	1.056	1.084	0.42	2.21	100	0.2	0.1	1.398	1.284	1.367	8.13	2.18	10	0.1	0.1	0.0023	0.0021	0.0022	8.24	2.14
10	1.1	0.2	1.376	1.338	1.393	2.82	1.24	100	0.2	0.3	3.281	2.996	3.245	8.68	1.08	10	0.1	1	0.0157	0.0147	0.0161	6.27	2.63
10	1.1	0.3	1.704	1.619	1.703	4.98	0.07	100	0.2	0.5	5.085	4.708	5.124	7.41	0.75	10	0.1	10	0.1482	0.1409	0.1549	4.98	4.45
10	1.2	0.1	1.129	1.126	1.154	0.24	2.23	100	0.3	0.1	1.595	1.498	1.581	6.09	0.88	10	1	0.1	0.0053	0.0053	0.0054	0.42	2.22
10	1.2	0.2	1.443	1.408	1.464	2.42	1.45	100	0.3	0.3	3.541	3.210	3.459	9.34	2.30	10	1	1	0.0198	0.0179	0.0193	9.55	2.47
10	1.2	0.3	1.769	1.690	1.773	4.48	0.25	100	0.3	0.5	5.375	4.922	5.338	8.43	0.70	10	1	10	0.1527	0.1440	0.1580	5.66	3.49
20	1	0.1	0.998	0.991	1.019	0.65	2.14	300	0.1	0.1	1.220	1.110	1.191	9.01	2.35	100	0.01	0.1	0.0020	0.0018	0.0019	9.48	2.52
20	1	0.2	1.318	1.274	1.330	3.28	0.94	300	0.1	0.3	3.090	2.886	3.130	6.61	1.27	100	0.01	1	0.0154	0.0146	0.0159	5.55	3.47
20	1	0.3	1.648	1.558	1.641	5.51	0.45	300	0.1	0.5	4.945	4.662	5.069	5.72	2.49	100	0.01	10	0.1491	0.1420	0.1559	4.79	4.52
20	1.1	0.1	1.067	1.062	1.090	0.43	2.18	300	0.2	0.1	1.447	1.332	1.413	7.93	2.31	100	0.1	0.1	0.0023	0.0021	0.0023	8.20	2.19
20	1.1	0.2	1.384	1.345	1.401	2.8	1.21	300	0.2	0.3	3.390	3.108	3.352	8.32	1.12	100	0.1	1	0.0159	0.0149	0.0163	6.20	2.56
20	1.1	0.3	1.713	1.628	1.712	4.97	0.10	300	0.2	0.5	5.253	4.884	5.291	7.03	0.70	100	0.1	10	0.1496	0.1423	0.1562	4.86	4.43
20	1.2	0.1	1.136	1.133	1.161	0.24	2.21	300	0.3	0.1	1.653	1.554	1.635	5.96	1.05	100	1	0.1	0.0053	0.0053	0.0055	0.43	2.17
20	1.2	0.2	1.451	1.416	1.472	2.42	1.42	300	0.3	0.3	3.660	3.330	3.574	9.01	2.35	100	1	1	0.0199	0.0181	0.0194	9.46	2.48
20	1.2	0.3	1.779	1.699	1.783	4.47	0.22	300	0.3	0.5	5.554	5.106	5.513	8.06	0.74	100	1	10	0.1541	0.1455	0.1594	5.55	3.47
30	1	0.1	1.003	0.997	1.025	0.65	2.11	500	0.1	0.1	1.260	1.150	1.229	8.71	2.40	1k	0.01	0.1	0.0022	0.0020	0.0022	8.79	2.88
30	1	0.2	1.325	1.282	1.337	3.27	0.91	500	0.1	0.3	3.189	2.990	3.229	6.25	1.22	1k	0.01	1	0.0168	0.0161	0.0174	4.62	3.15
30	1	0.3	1.658	1.566	1.650	5.50	0.48	500	0.1	0.5	5.103	4.830	5.228	5.35	2.43	1k	0.01	10	0.1626	0.1565	0.1696	3.77	4.28
30	1.1	0.1	1.073	1.068	1.096	0.42	2.16	500	0.2	0.1	1.496	1.380	1.460	7.74	2.42	1k	0.1	0.1	0.0026	0.0024	0.0025	7.73	2.63
30	1.1	0.2	1.392	1.353	1.408	2.81	1.18	500	0.2	0.3	3.499	3.220	3.459	7.97	1.16	1k	0.1	1	0.0173	0.0164	0.0177	5.25	2.30
30	1.1	0.3	1.723	1.638	1.721	4.95	0.12	500	0.2	0.5	5.422	5.060	5.458	6.66	0.66	1k	0.1	10	0.1631	0.1568	0.1699	3.84	4.18
30	1.2	0.1	1.142	1.139	1.167	0.24	2.19	500	0.3	0.1	1.710	1.610	1.690	5.85	1.20	1k	1	0.1	0.0059	0.0059	0.0060	0.52	1.69
30	1.2	0.2	1.459	1.424	1.479	2.41	1.39	500	0.3	0.3	3.779	3.450	3.689	8.71	2.40	1k	1	1	0.0218	0.0199	0.0212	8.67	2.67
30	1.2	0.3	1.789	1.709	1.792	4.46	0.19	500	0.3	0.5	5.732	5.290	5.688	7.71	0.77	1k	1	10	0.1679	0.1603	0.1734	4.53	3.26




Figure 4.3: Graphs showing how a repeater insertion strategy optimised for a particular switching pattern performs for other switching patterns. The x-axis shows different aggressor switching patterns, and the y-axis the delay for different repeater sizes and numbers. Case (i) refers to a repeater insertion strategy where K and H are optimal for minimizing delay for pattern (i).

4.3.1 Minimum Delay

Equations (4-10) and (4-11) give the K and H values for minimizing delay for different switching patterns. The obvious question is, how will a repeater insertion strategy optimised for a particular switching pattern work for other patterns? Given in Fig. 4.3 are the delays for different patterns, when the repeater insertion strategies are optimised for cases (a) through (f), excepting (e). The net considered here has a resistance of $1 \text{ k}\Omega$ and capacitances of 100 fF to ground and to each of the adjacent wires. R_{drv} and C_{drv} are set to $7.7 \text{ k}\Omega$ and 9.5 fF to match the $0.35 \mu\text{m}$ technology used for testing. The legend termed *single* refers to the conventional delay minimization strategy that would be carried out by treating the total capacitance as a single lumped component. These values are also given for comparison in Table 4.4. Now as expected, for each switching pattern, the delay is minimum for the H and K that is optimised for that particular pattern and is not optimal for other patterns. It can also be seen that the optimal strategy for minimizing the worst-case delay is indeed $H_{1,opt}$ and $K_{1,opt}$. *Although this particular number and size of repeaters performs sub-optimally for cases (b) through (f), they do not perform so badly that the delay for any one of these patterns is greater*

Table 4.4 Shows how a repeater insertion strategy optimised for a particular switching pattern performs for other switching patterns (data corresponds to the graphs in Fig. 4.3)

Optimisation Strategy			Line Delay (pSec)				
Case	K	H	Case 1	Case 2	Case 3	Case 4	Case 6
1	2	21	476	421	349	349	287
2	2	18	479	418	340	340	272
3	1	14	546	453	330	330	222
4		14	546	453	330	330	222
6	1	9	625	504	346	346	211
Single	1	13	556	458	330	330	218

than the delay corresponding to case 1. Hence when a repeater insertion strategy is referred to as optimal, it means that H and K take the values $H_{1,opt}$ and $K_{1,opt}$ respectively.

As mentioned before, in general, for design purposes the delay of a line would refer to the worst-case as any switching pattern can occur. What is interesting here is the fact that a repeater insertion strategy that is more aggressive than the optimal predicted by the conventional analysis for an isolated line, helps in reducing the worst-case delay. The exact numbering and sizing would of course depend on the timing constraints, and the resources available for repeaters in terms of maximum allowable area and power. These models are thus an aid in deciding upon a repeater insertion strategy to match the particular application.

4.4 Model Verification

4.4.1 Aggressor Alignment

The effect of aggressor alignment (the times at which the aggressors switch relative to the victim net) on delay is a much researched topic [Gross98]. For a three net arrangement such as was considered here, it has been shown that when the slew rates are

unequal, the worst-case delay is caused by aggressors which switch at different times [Kahng00]. Since the models are built up by considering simultaneously switching nets, and case (a) is presented as being very near worst-case, it is interesting to check the inaccuracy introduced by the assumption.

Since uniformly coupled data lines are being analysed, it is reasonable to make the simplifying assumption that the rise times of the input signals are the same. Even for this simplified case, it is not simultaneous switching, but both aggressors switching slightly after the victim that causes the worst delay. This is however a very small difference and is really negligible. The effect of different aggressor alignment on delay can be seen by inspection of the eye diagram at the output of the victim net, built up over hundreds of cycles, with different pseudo-random bit streams (PRBS) being fed to the three lines. Consider a net with $R=600\ \Omega$, $C_s=550\ \text{fF}$ and $C_c=100\ \text{fF}$, where $R_{drv}=208\ \Omega$ and $C_{drv}=351\ \text{fF}$ ¹. The worst-case delay predicted by (4-6) is 277.5 ps. Now shown in Fig. 4.4 is the eye diagram built up with 1000 bits of 1ns period having 100ps rise and fall times where PRBSs with different seeds have been fed to the three lines. The worst-case delay is indicated by the intersection of the markers, and is 274.9 ps,




Figure 4.4: Eye diagram at the output of the victim net showing the effect of aggressor alignment on delay

1. These figures are the Thevenin resistance and input capacitance of an appropriately sized MOS driver.

which is very close to that predicted by the model. The exact error depends very much on the rise times used. Obviously the smaller the rise time, the more accurate is the model.

4.4.2 Testing with Real Repeaters

Table 4.5 Comparison of model for buffered net with worst-case cross-talk against actual delay with real inverters

R Ω	C_s fF	C_c fF	K	H	T_d (actual) ps	T_d (model) ps	Error%
600	550	100	2	37	607.0	555	8.5%
800	100	100	2	23	534.0	477	10.6%
1k	100	100	2	21	581.0	526	9.5%
600	550	550	3	63	1061	918	13.4%
800	1000	100	3	38	918.0	757	17.6%
1k	550	100	3	28	864.0	704	18.6%
600	550	1000	4	82	1300	1165	10.4%
800	550	550	4	55	1223	1047	14.4%
1k	100	550	4	45	1181	1072	9.2%
600	1000	1000	5	86	1435	1216	15.3%
1k	550	550	5	49	1395	1168	16.3%
1k	1000	550	5	53	1524	1251	17.9%

The accuracy of the models was investigated with an actual 0.35 μm technology. The input capacitance of a minimum sized inverter in that technology is approximately 9.5f F while its output impedance is 7.7k ohm. Signal rise and fall times of 100p seconds were used. In the same technology, a 1 cm long wire in metal 3 has a total capacitance to substrate of 720f F, a coupling capacitance of 850f F to an adjacent wire with minimum spacing, and a resistance of 800 ohm. Hence the loads in Table 4.5 are chosen to represent global or semi-global length wires. The repeater insertion strategy that is shown here is $H_{I,opt}$ and $K_{I,opt}$, and the accuracy is tested for case (a).

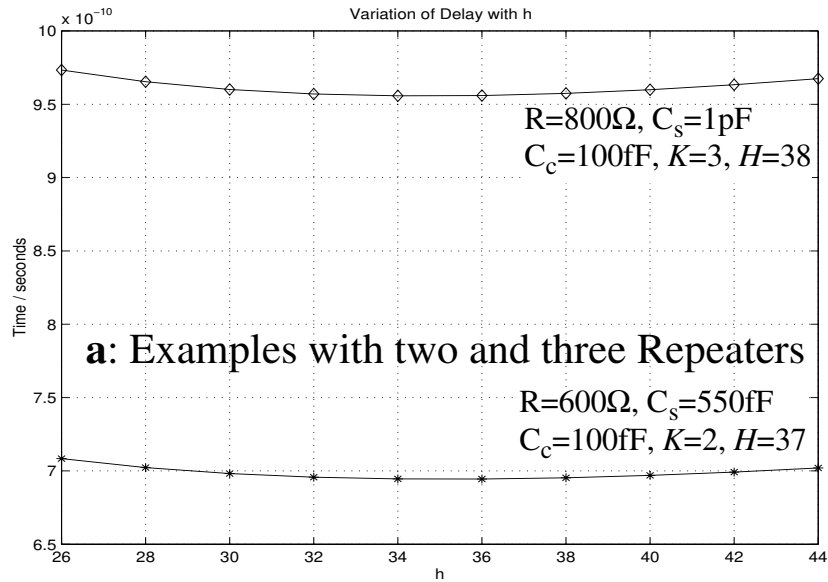


Figure 4.5.a: Top graph shows the delay for a net of $R=800 \Omega$, $C_s=1 \text{ pF}$, $C_c=100 \text{ fF}$, $K=3$, $H=38$ and the bottom for a net of $R=600 \Omega$, $C_s=550 \text{ fF}$, $C_c=100 \text{ fF}$, $K=2$, $H=37$




Figure 4.5.b: top graph shows the delay for a net of $R=600 \Omega$, $C_s=1 \text{ pF}$, $C_c=1 \text{ pF}$, $K=5$, $H=86$ and the bottom for a net of $R=600 \Omega$, $C_s=550 \text{ fF}$, $C_c=1 \text{ pF}$, $K=4$, $H=82$

Figure 4.5: Graphs showing how the delay varies with repeater size around the optimal size for minimum delay




Figure 4.6: Driver characterization

The drop in accuracy seen here is due more to the effect of resistive shielding, i.e. poor driver modelling than a weakness in the delay models. The practice of treating the inverter as a voltage source-resistor-capacitor combination where the parasitics scale linearly with size, and ignoring all second order effects, though poor, is often the only option available for timing driven layout optimization.

As seen in Table 4.5, the accuracy of the total delay predicted by (4-6) is accurate within 82% and 92%. The *fidelity* of (4-10) and (4-11) appear however to be much greater. Fidelity refers to the closeness of the solutions predicted by (4-10) and (4-11) to the optimal solutions. This is evident from Fig. 4.5, where the results of simulations for a range of H situated either side of the value predicted by (4-6) are shown. It can also be seen that the delay curves are quite flat, and K and H can be relaxed with little loss in performance.

4.5 Estimating Device Characteristics

A driver is characterized by a capacitance and an effective resistance which scale linearly with size. The estimation of the input capacitance is straightforward:

$$\begin{aligned}
 C_{in} &= C_{ox} + C_{ovrlp} \\
 &= \frac{3.9 \times 8.85 \times 10^{-14} \times w \times l_{eff}}{t_{ox}} + C_{gdo}w
 \end{aligned} \tag{4-12}$$

The driver resistance can be calculated by considering the equivalence between a lumped RC network and the actual driver as seen in Fig. 4.6. The left hand side shows the inverter which uses the actual (Spectre or Spice) transistor model supplied by the technology vendor, and includes the junction capacitance and all other capacitances and non-linear effects. The average delay t_{dj} corresponding to a load capacitance C_{Lj} is calculated by running transient simulations (with Spectre in this case) for a range of W/L ratios depicted in terms of multiples of the minimum ratio, H_i , chosen such that it matches the relevant range. Then the average delay corresponding to a minimum sized inverter is calculated as shown below.

$$t_{d_j} = \frac{1}{n} \sum_{i=1}^n H_i \times t_{d_{j,i}}(C_{L_j}, h_i) \quad (4-13)$$

Finally the load capacitance C_{Lj} is also averaged out over the relevant range, resulting in an output impedance for a minimum sized driver as given below.

$$R_{drv_m} = \frac{1}{0.69m} \sum_{j=1}^m \frac{t_{d_j}}{C_{L_j}} \quad (4-14)$$

Depending on the circumstances, both averages can be over as wide or as narrow a range as necessary. For very wide ranges, different values can be used. Characterization for future technologies is done by scaling from an existing technology. A similar technique is proposed in [Sylvester99b].

Quite a lot of information is lost in the linearisation, but the effect is minimised by running simulations with detailed transistor models, when the effect of junction and other parasitic MOS capacitances, and non-linear effects are encapsulated in some sense, in the output impedance.

4.6 Summary

This chapter has focused on developing buffer models for aiding in the system level analyses of long coupled interconnects. The models were derived by considering different equivalent capacitive loads for different switching patterns of victim and aggressor drivers. Although very simple, they are useful for system level analyses, and in particular for investigating the performance of buses, as carried out in the next chapter.

5. Optimal Signalling Over On-chip Buses

Optimisation of wires and repeaters for maximising bandwidth over a parallel wire topology is discussed. Analytic guidelines are derived, for designing the wires and repeaters. An alternative to repeater insertion for long lossy lines in the form of error-control coding is investigated.

5.1 Introduction

Moore's law has held remarkably true over the years and challenges at the device level have been and are being met with solutions of great ingenuity. It seems reasonable to assume that Moore's law will continue to hold true over the next eight to ten years. The ability to put hundreds of millions of transistors on a single chip has however created new challenges for the systems engineer of dealing with the complexity in such a way that potential bottlenecks such as timing closure, power distribution and input/output requirements are not allowed to dictate the ultimate size and hence the functionality of the chip. A potential solution is an on-chip packet switched network, which has been proposed by a number of authors [Sylvester99a], [Dally01], [Sgroi01] and [Benini01]. Whether of a regular tiled nature or otherwise, the inter-block communication link in all of these schemes will consist of a large number of parallel wires, with uniform coupling over most of the wire length in all probability.

This chapter examines signalling techniques and conventions over such relatively long coupled lossy lines, with emphasis on minimizing delay and maximizing bandwidth over multi-net structures. A key question that is posed is, given a fixed area in which to distribute the interconnect, what is the best arrangement of the wires to obtain the highest bandwidth? Is it to have a few fat wires and a high signalling frequency, or a large number of small wires with a lower signalling frequency, or anything in between? How does the wire spacing affect overall bandwidth? What effect does repeater insertion have? How many repeaters should there be, and how should they be sized?

In particular the exact manner in which the total capacitance is distributed into a ground component, and a component consisting of the capacitance to the adjacent wires is important, as this dictates the charging/discharging time. In the following sections an analysis is carried out for optimising bandwidth which maps the wire geometry to the parasitics, and uses the models derived in the previous chapter. It is shown that for a given metal resource in terms of a fixed total width, there is a clear global optimum consisting of a particular number of wires having a particular wire width and

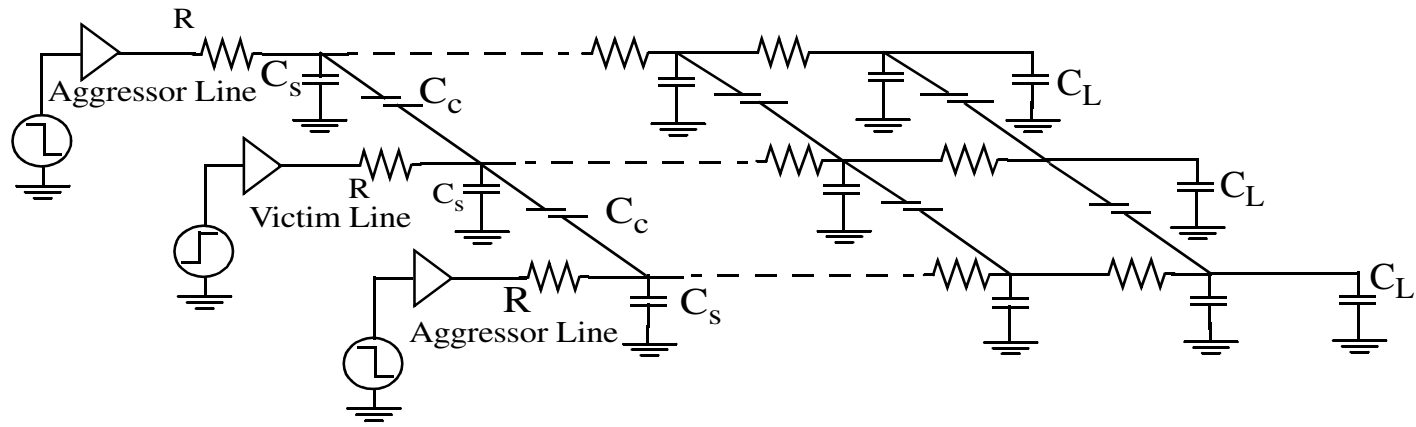


Figure 5.6.a: True distributed RC line with uniform coupling

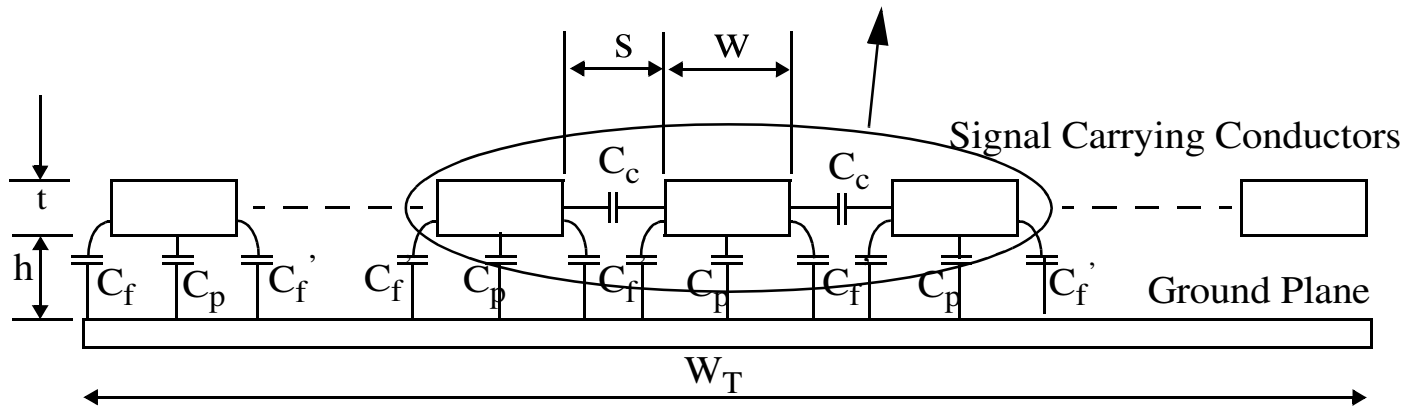


Figure 5.6.b: Geometry of the parallel multi-net structure

spacing. This optimum configuration does not necessarily translate to the maximum parallelism allowed by the technology, and in fact deviates considerably from it when the resources available for repeater insertion are limited. For wide buses, this optimal wire width and spacing is mostly independent of the total area.

The line model considered here is the same as was considered in the previous chapter, a long RC line uniformly coupled for its entire length on both sides to aggressor lines. The metrics proposed in [Ismail99] are used to verify that the RC model is valid. Closed form equations that match the wire geometry to the parasitics are required for the bandwidth analysis. From a review of the existing literature (see Chapter 2) it appears that there are only a few choices for DSM technologies. Reported models include [Lewis84], [Chern92] and more recently [Lee98] and [Zheng00]. The equations in [Lewis84] have been widely used in the past, but drop in accuracy when the aspect ratio of the wires increase to DSM proportions. The methodology proposed in [Lee98] uses numerous technology dependent constants, which render the models rather difficult to use without familiarity with their derivation. The formulae proposed in [Chern92] and [Zheng00] can be conveniently used for parasitic extraction of DSM geometries. The models in the latter reference use a single technology dependent constant, derived by generating a database of values with a field solver for different geometries in a particular technology, and then using curve fitting techniques. They are in effect a modification of Sakurai's equations to render the partitioning between self and coupling capacitances more accurate (see Chapter 2, Section 2.3). In the bandwidth analysis these latter models will be used for mapping the wire geometry to the capacitive parasitics.

5.2 Interconnect Modelling and Delay Analysis

The electrical model for investigating delay is shown in Figure 5.6.b. Each line, except the two peripheral lines are coupled on both sides to aggressors. The reason is that this is closest to the actual situation for an interconnect in a bus. This is a lossy capacitive model which does not include inductance, and is valid for the thin wires which are typical of DSM technologies. The driver modelling and repeater insertion is identical to that carried out in the previous chapter.

5.2.1 Parasitic Modelling

To calculate the capacitance terms shown in Figure 5.1 the models proposed in [Zheng00] are used. They use a technology dependent constant β which is calculated

from a database of values generated by a field solver, and are defined in equations (5-1) through (5-6).

$$C_f = \epsilon_k \left[0.075 \left(\frac{w}{h} \right) + 1.4 \left(\frac{t}{h} \right)^{0.222} \right] l \quad (5-1)$$

$$C'_f = C_f \left[1 + \left(\frac{h}{s} \right)^\beta \right]^{-1} \quad (5-2)$$

$$C_p = \epsilon_k \frac{wl}{h} \quad (5-3)$$

$$C_{s, mid} = C_p + 2C'_f \quad (5-4)$$

$$C_{s, corn} = C_p + C_f + C'_f \quad (5-5)$$

$$C_c = C_f - C'_f + \epsilon_k \left[0.03 \left(\frac{w}{h} \right) + 0.83 \frac{t}{h} - 0.07 \left(\frac{t}{h} \right)^{0.222} \right] \left(\frac{h}{s} \right)^{1.34} l \quad (5-6)$$

Typical values of β range from 1.50 to 1.75, and 1.65 may be used for most DSM technologies. The equations are reported to be accurate to over 85% when the following inequalities are satisfied.

$$0.3 < (w/h) < 30 \quad (5-7)$$

$$0.3 < (t/h) < 10 \quad (5-8)$$

$$0.3 < (s/h) < 10 \quad (5-9)$$

For DSM circuits, typical geometries are well within this range. Additionally, the DC resistance is given by:

$$R = R_{SQ} \frac{l}{w} \quad (5-10)$$

5.2.2 Line Delay and Repeater Insertion

The analysis carried out in the previous chapter is valid for the boundary conditions considered here. Reproduced here is the delay for a line with equalised repeaters for case (a):

$$t_{eq} = K \left[0.7 \frac{R_{drv_m}}{H} \left(\frac{C_s}{K} + HC_{drv_m} + 4.4 \frac{C_c}{K} \right) + \frac{R}{K} \left(0.4 \frac{C_s}{K} + 1.51 \frac{C_c}{K} + 0.7 HC_{drv_m} \right) \right] + \frac{t_r}{2} \quad (5-11)$$

and the optimal H and K for minimizing delay.

$$K_{1,opt} = \sqrt{\frac{0.4RC_s + 1.51RC_c}{0.7R_{drv_m} C_{drv_m}}} \quad (5-12)$$

$$H_{1,opt} = \sqrt{\frac{0.7R_{drv_m} C_s + 3.08R_{drv_m} C_c}{0.7RC_{drv_m}}} \quad (5-13)$$

5.3 Optimal Signalling Over Parallel Wires

For the wire arrangement show in Fig. 5.1, the *worst-case* delay of a line is defined as t_{WC}^1 . Since in general it has to be assumed that the worst-case aggressor-victim

1. The delays of the two corner conductors differ as they are coupled to only one line, and the distribution of the capacitance changes slightly. Considering this difference would be an unnecessary refinement for most applications.

switching pattern will occur on a given line, any calculation of bandwidth has to consider the worst-case delay as the *minimum* delay over a line. This minimum delay, as was shown, depends on the resources available for repeater insertion, but always corresponds to the switching pattern in case (a) (see Chapter 4). Hence for all delay calculations, equation (5-11) is used. The line delay is matched to the minimum pulse width T , by allowing a sufficient margin of safety. The exact mapping depends on the type of line [Deutsch01], but it is generally accepted that three propagation delays are sufficient to let the signal cross the 90% threshold for RC lines [Ho01]. Since the worst-case delay is already considered with good accuracy, a factor of 1.5 is deemed to be sufficient¹, resulting in (5-14).

$$T = 1.5t_{WC} \quad (5-14)$$

The total bandwidth in terms of bits per second is now given by:

$$BW = \frac{N}{T} \quad (5-15)$$

This expression changes if pipelining is carried out so that at any given time, more than 1 bit -up to a maximum of one bit per each gain section- is on the line. Since each repeater will refresh the signal and sharpen its rise or decay, the mapping between the propagation delay and the pulse width needs to be carried out for *each section*. Theoretically it is possible to gain an increase in bandwidth by introducing repeaters up to the limit where the bit width is determined by considerations other than the delay of a single stage, or where the delay of the composite net is greater than its constraint. In practice one rarely sees repeaters introduced merely for the sake of pipelining, when the total delay of the net, and power consumption increases as a result. If pipelining is carried out, it is a simple matter to multiply (5-15) by the appropriate factor.

The number of signal wires N , that can be fitted into a given area depends on whether shielding is carried out or not. In general, shielding individual lines is only useful against capacitive cross-talk. The magnetic field will in all probability permeate the entire breadth and length of the bus, and can only be contained by very fat wires. Hence for the shielded case it is assumed that the shielding wires are the thinnest permitted by the technology, regardless of the size of the signal wires, as this serves the intended purpose while minimizing area for non-signal wires. From the geometry of Fig. 5.1,

1. The constant used to match the 50% delay to the pulse width depends on the application and is irrelevant in the context of the methodology. The value given here is used merely to be able to talk in terms of numbers.

the relation given in (5-16) is obtained for unshielded wires, and (5-17) for shielded wires.

$$W_T = NW + (N - 1)S \quad (5-16)$$

$$W_T = NW_{signal} + (N - 1)(2S + W_{shield}) \quad (5-17)$$

The problem definition is to maximize the bandwidth for a constant width W_T . Depending on whether or not the designer has freedom over wire sizing, the analysis is different. These two cases are covered in sections 5.3.1 and 5.3.2.

5.3.1 Fixed Wire-Width and -Pitch

When the wire-width and -pitch is fixed, optimising bandwidth reduces to the simple task of designing the repeaters to minimize delay over each individual line. The issue is complicated when the resources available for repeater insertion are limited. The area of a minimum sized inverter can be modelled as the sum of two components, one of which is dependent on the W/L ratio of the transistors, and one which is independent of it. Now since the repeaters are H times a minimum sized inverter and are K in number, minimizing the area is equivalent to minimizing the product HK . The dynamic power consumption of an inverter is $0.5C_{load}V_{dd}^2f$ (where f refers to frequency), and hence for a given frequency, power consumption is minimized by minimizing C_{load} . Since the output capacitance of an inverter is proportional to H , minimizing power consumption is also equivalent to minimizing HK .

The problem of repeater optimization for uniform coupled nets can take two forms. Either the maximum acceptable delay for the net is specified, and the objective is to minimize area subject to the constraint $t \leq t_{max}$, or the maximum acceptable area is specified and the objective is to minimize the delay subject to the constraint $A \leq A_{max}$. Consider Fig. 5.7 which shows the variation of delay with H and K where the line parasitics correspond to row 1 of Table 4.3. The plane shows a delay constraint of 1.3 ns for that net, and any of the K and H combinations which lie below this and on the curved surface showing the delay is acceptable to meet the delay constraint. Also shown is an appropriately scaled plot of HK .

Because HK is quasi-concave in the quadrant of positive H and K , it is not possible to find an analytical solution to the first optimization problem, which has to be solved numerically. However it is possible to analytically solve the second optimization prob-

lem because its objective function $t_{0.5}$ as given in (6), is concave as seen in the figure. The optimum solution can be found by solving the Karush-Kuhn-Tucker conditions [Karush39], [Kuhn51] given by the equations (5-18) through (5-22) where L_i refer to the Lagrangian constants.

$$0.7 \frac{R_{drv}}{H^2} (C_s + 4.4C_c) - 0.7RC_{drv} + L_1K + L_2 = 0 \quad (5-18)$$

$$\frac{R}{K^2} (0.4C_s + 1.5C_c) - 0.7R_{drv}C_{drv} + L_1H + L_3 = 0 \quad (5-19)$$

$$L_1(HK - A_{max}) = 0, \quad HK \leq A_{max}, \quad (L_1 \geq 0) \quad (5-20)$$




Figure 5.7: Variation of delay with H and K for a net having $R=600 \Omega$, $C_s=550$ fF and $C_c=100$ fF. The plane at 1.3 ns describes the delay constraint for that net, while the third surface is an appropriately scaled plot of HK . Any of the H, K coordinates corresponding to the points on the curved convex surface below the plane are acceptable to meet the delay constraint, and the particular point among all these points that gives the minimum HK product is the most desirable solution.

$$L_2(H - 1) = 0, \quad H \geq 1, \quad L_2 \geq 0 \quad (5-21)$$

$$L_3(K - 1) = 0, \quad K \geq 1, \quad L_3 \geq 0 \quad (5-22)$$

5.3.2 Variable Wire-Width and -Pitch

In this section it is considered additionally that the wire size and spacing can also change. Typically in a process the wires in a certain layer are limited to tracks determined by the minimum feature size of the technology. Within this frame, the designer has freedom to vary the spacing and the width of the wires. Now the problem definition can be stated as follows: for a constant width W_T , what are the N (number of conductors), s (spacing between conductors), and w (width of an conductor) values that give the optimum bandwidth? The variables are discrete as s and w are dictated by the process as well, and there are geometrical limits which cannot be exceeded. The optimal arrangement depends very much on the resources allocated for repeaters, and is investigated by simulations first. Then approximate analytic equations are developed that give close to optimal solutions, and can be used as guidelines to quickly obtain the true solution.

Simulations

The simulations are carried out for a future technology with parameters estimated from guidelines laid out in [ITRS01]. The minimum feature size is 50 nm, and copper wires are assumed with the technology dependent constant β being 1.65, height above substrate h being 0.2 μm , and wire thickness t being 0.21 μm . The minimum wire width and spacing are each assumed to be 0.1 μm and the output impedance of a minimum sized inverter estimated to be 7k Ω and its input capacitance 1f F. In all cases the constraint for the wires is set to a total width of 15 μm . Of the three variables N , s and w , only two are linearly independent, as the third is defined by (5-16) or (5-17) for any values that the other two may take. We choose to vary N and s , and assume that w and s are variable in multiples of the minimum pitch. In the subsequent sections different constraints on the repeaters are considered.

Ideally driven line

Although ideal sources are never present in practice, the wire arrangement for the optimum bandwidth is interesting as it serves as a point of comparison for later results.

Given in Fig. 5.8.a is the plot of how the bandwidth varies with N and S . It can be seen that there is a clear optimum of 16 conductors which is far from the maximum number of 150 conductors allowed by the technology constraints.

Unshielded lines with optimal buffering

The bandwidth for changing N and s where the repeaters are optimally sized is plotted in Fig. 5.8.b. It can be seen that the maximum bandwidth is obtained when the parallelism is the maximum allowed by the physical constraints of the technology, of $w=s=0.1\mu\text{m}$. This result is logical because the buffers which are optimally sized for each configuration compensates for the increased resistance and cross-talk effect. The values of H and K are 52 and 7 respectively, while the maximum bandwidth is 345.5 Gbits/sec.

Unshielded lines with constant buffering

Optimal repeater insertion results in a large number of huge buffers. Also, as is the case with optimal buffering in general whether the load is lumped or distributed, the delay curve is quite flat, and the sizes can be reduced with little increase in delay. Instead of optimal repeater insertion, if a constraint is imposed on the number and size of buffers for each line, the optimal configuration does *not* equate to the maximum number of wires. Given in Fig. 5.8.c is a plot of the bandwidth when a constraint of $K=1$ and $H=20$ is laid down for each line. The optimal configuration corresponds to $w=0.16\mu\text{m}$, $s=0.2\mu\text{m}$ and $N=42$, so that the NHK product is 840. The maximum bandwidth is now 171.1 Gbits/sec.

Unshielded lines with constrained buffering

Typically the constraint would be on the total area occupied by the buffers, and hence K and H would be affected by N . If (28) describes the area constraint on the buffers, the optimum configuration is the solution to the constrained optimization problem of maximizing (5-15) subject to (5-23).

$$NKH \leq A_{max} \tag{5-23}$$

This adds a third independent variable to the objective function (21), of either K or H since A_{max} is a constant. It is a simple matter to incorporate all the relevant equations presented here into an iterative algorithm that can be used to obtain a computer generated solution. As an example, assume that A_{max} is set to 500 for the same boundary conditions. It turns out that the optimal configuration is when $K=1$, and shown in Fig.




Figure 5.8.a: Ideal Drivers




Figure 5.8.b: Optimal Repeater Insertion for unshielded lines




Figure 5.8.c: Fixed size and number of buffers for each line for unshielded lines




Figure 5.8.d: Total resources for repeaters (HK product) is fixed for unshielded lines

Figure 5.8: Variation of bandwidth with number of conductors (N), and spacing between conductors (s) over a fixed metal resource for different repeater configurations

5.8.d is a plot of the bandwidth where $K=1$ and H changes according to N . The optimal wire arrangement turns out to be $w=0.26 \mu\text{m}$, $s=0.4\mu\text{m}$ and $N=23$.

Shielded lines with optimal buffering

In general, shielding each signal wire results in a *drop* in the overall bandwidth. The reason is that although shielding reduces the delay over each individual line, the reduction in the number of signal lines more than negates this effect. Shown in Fig. 5.8.e is a plot of the bandwidth where every other wire is a minimum sized shielding wire, and the signal wires are buffered optimally. The total bandwidth of 261.3 Gbits/sec is less than in the unshielded case. This reduction is however accompanied with a saving in repeater size, and shielding can be considered as an option to reduce area and power consumption for repeaters.

Shielded lines with constrained buffering

This plot also offers a straight comparison with the unshielded case. There is a drop in the bandwidth as can be seen, from 163 Gbits/sec to 160 Gbits/sec. With constrained

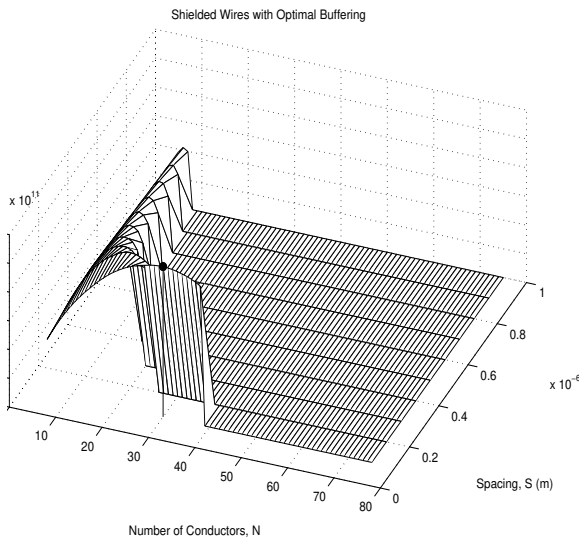


Figure 5.8.e: Fixed size and number of buffers for each line for unshielded lines




Figure 5.8.f: Total resources for repeaters (HK product) is fixed for shielded lines

Figure 5.8: Variation of bandwidth with number of conductors (N), and spacing between conductors (s) over a fixed metal resource for different repeater configurations

buffering, the repeater area and power consumption are the same in the two cases, as the maximum available resources are utilized.

Validity of Analysis

Since inductance is not considered in the timing model, the question arises of how close the prediction is to the true optimum for real wires which always have non-zero inductance. Inductance as mentioned before, depends on the signal return path, and hence is relatively insensitive to the wire width. Typical values range from 2-4 nH/cm [Ismail01]. If the metrics defined in (2-48) and (2-49) are applied with a signal rise time of 50 ps and the very conservative inductance value of 5 nH/cm, it can be seen that the inductive effects are not important even for the fattest wires in the plots, which are in an unimportant region, and far away from the optimal point.

Analytic Guidelines

An analysis for the optimum bandwidth with the exact capacitance equations proves to be intractable rather quickly. However an approximate solution can be derived by recognizing certain characteristics in the fringe capacitance terms. An inspection of (5-1) shows that dependence of C_f on the width w is rather weak. (*In fact this is the main reason that increased wire width results in reduced delay; the total capacitance of a wire is dominated by the fringe component, which is insensitive to w . Hence it is possibility to increase the width by a certain factor and reduce the resistance proportionally, and benefit from the fact that the parallel plate capacitance which increases by the same proportionate factor is only a small portion of the total capacitance, which reduces the overall RC product.*) The contribution from the term proportional to w is much less than the term proportional to the t/h ratio. Hence an approximate expression for C_f is given in (5-24) which is constant in the face of changing w and s .

$$C_{f_{app}} = 1.4\epsilon_k \left(\frac{t}{h}\right)^{0.222} l \quad (5-24)$$

Similarly, the term proportional to w can be neglected for the expression for C_c , leading to (5-25):

$$C_{c_{app}} = C_{f_{app}} - C'_{f_{app}} + a\epsilon_k \left(\frac{h}{s}\right)^{1.34} l \quad (5-25)$$

where a is a unitless constant defined by (5-26):

$$a = 0.83 \frac{t}{h} - 0.07 \left(\frac{t}{h} \right)^{0.222} \quad (5-26)$$

Now the approximate pulse width T which is defined as $1.5t_{WC}$ can be expressed in the form given in (5-27)

$$T = T_1 \frac{w}{h} - \frac{T_2}{1 + (h/s)^{1.65}} - \frac{T_3 l}{[1 + (h/s)^{1.65}] w} \quad (5-27)$$

$$+ T_4 (h/s)^{1.34} + T_5 \frac{l}{w} + T_6 \frac{l}{w} (h/s)^{1.34} + T_7$$

where the time constants are defined in (5-28) through (5-35).

$$T_1 = 0.7 \varepsilon_k R_{drv} l \quad (5-28)$$

$$T_2 = 1.68 R_{drv} C_{f_{app}} \quad (5-29)$$

$$T_3 = 0.71 \frac{R_{SQ} C_{f_{app}}}{K} \quad (5-30)$$

$$T_4 = 3.08 a \varepsilon_k R_{drv} l \quad (5-31)$$

$$T_5 = 1.51 \frac{R_{SQ} C_{f_{app}}}{K} + 0.7 C_{drv} R_{SQ} \quad (5-32)$$

$$T_6 = 1.51 \frac{a \varepsilon_k R_{SQ} l}{K} \quad (5-33)$$

$$T_7 = 0.4 \frac{\epsilon_k R_{SQ} l^2}{Kh} + 3.08 R_{drv} C_{f_{app}} \quad (5-34)$$

$$R_{drv} = \frac{R_{drv_m}}{H} \quad C_{drv} = HC_{drv_m} \quad T_{drv_m} = R_{drv_m} C_{drv_m} \quad (5-35)$$

Now substituting for N in (5-15) from (5-16) (since it was shown that unshielded wires result in the greater bandwidth over a constant metal resource, only this case is considered) results in expression (5-36) for bandwidth.

$$BW = \frac{W_T + s}{(w + s)T} \quad (5-36)$$

This is a concave function in w and s with a global maximum as was shown in the simulated plots. At this maximal point, the numerators of the partial derivatives of BW with respect to w and s are zero. Recognizing that $s \ll W_T$ close to the optimal point allows the following equations to be derived from these two conditions.

$$T + (w + s) \frac{\partial T}{\partial w} = 0 \quad (5-37)$$

$$T + (w + s) \frac{\partial T}{\partial s} = 0 \quad (5-38)$$

Substituting for T from (5-27) in (5-37) and doing some rather unpleasant number crunching allows w to be written as an explicit function of s , as defined in (5-39).

$w =$

$$\sqrt{\frac{[T_3 - T_5 - T_6(h/s)^{1.34} - T_5(h/s)^{1.65} - T_6(h/s)^3]ls}{T_2 - T_7 + T_1(h/s)^{-1} + T_1(h/s)^{0.65} - T_4(h/s)^{1.34} - T_7(h/s)^{1.65} - T_4(h/s)^3}} \quad (5-39)$$

Also the partial derivative of T with respect to s is as given in (5-40).

$$\frac{\partial T}{\partial s} = -\frac{1.65}{h} \left(T_2 + T_3 \frac{l}{w} \right) \frac{(h/s)^{2.65}}{[1 + (h/s)^{1.65}]^2} - \frac{1.34}{h} \left(T_4 + T_6 \frac{l}{w} \right) (h/s)^{2.34} \quad (5-40)$$

Now in (5-38), $\partial T/\partial s$ is replaced by (5-40), T replaced by (5-27), and w replaced by (5-39) in the resulting expression. This results in a single variable function of s in the form of $f(s) = 0$. Given that the initial expressions were rather complex and unwieldy, this is a fairly simple equation, in so much that it is a function of a single variable with constants completely defined in terms of easily obtained technological parameters and the design constraints K , H and l . The coordinates of the optimal point is given by the roots of (5-38) and (5-39). Since BW is a well behaved function with a single maximal point in the regime of interest, (5-38) usually has only one root. This root can easily be found either by a simple iterative algorithm such as a binary search, or by inspection of a plot.

To demonstrate this, consider the first example in the simulation, which consisted of optimally buffered lines, when $K=7$ and $H=52$. Shown in Fig. 5.9.a is a plot of (5-38) shown against the different s values considered in the simulation. The only possible root is $s=0.1\mu\text{m}$, when (5-39) gives $w=0.1\mu\text{m}$, which is exactly the values given by the simulation. To consider a second example, simulations showed that for constrained buffering the optimal point is when $w=0.27\mu\text{m}$ and $s=0.4\mu\text{m}$, when $N=23$ and K and H are 1 and 21 respectively. The function (5-38) for these values of H and K are plotted in Fig. 5.9.b. The solution predicted by the roots is $w=0.27\mu\text{m}$ and $s=0.3\mu\text{m}$, when $N=27$. This is very close to the true optimum, and in fact, checking the values predicted by the exact equations with the values on either side of the s value predicted by (5-38), that of $s=0.2\mu\text{m}$ and $s=0.4\mu\text{m}$ results in the correct solution. Finally for the case with ideal drivers, when $R_{drv} = 0$ and $C_{drv} = 0$, the simulation showed that the optimal point was when $w=0.56\mu\text{m}$, $s=0.4\mu\text{m}$, and $N=16$. The plot of (5-38) shown in Fig. 5.9.c predicts the optimal to be $w=0.56\mu\text{m}$, $s=0.3\mu\text{m}$, when $N=18$. Again checking just the two values on either side of the approximate s value results in the correct solution. Hence (5-38) and (5-39) can be used to garner values that can either serve as the starting point for simulations with the exact equations to yield the true optimal point, or even be used unchanged, as they are quite close to the true optimum.

There is a rather important ramification of these approximate analytic equations for designing buses. An inspection of (5-37) and (5-38) reveals that the optimal bus width




Figure 5.9.a: Optimal Buffering




Figure 5.9.b: Constrained Buffering

Figure 5.9: Roots of the function $f(s) = T_+(w+s) \partial t/\partial s$

and spacing is independent of the total width, W_T . The only approximation made in deriving these two expressions was that the optimal spacing s is very small in comparison to W_T , which is valid for buses with word length greater than or equal to 8. This makes the design much less complicated, and the optimal wire width and pitch for maximizing bandwidth can easily be derived by estimating an initial solution with the analytic formulae, and then running a few simulations with the exact capacitance equations. The following guidelines can be followed in this process.

- (a). The maximum bandwidth across a metal resource can be achieved by fitting the maximum number of wires, each optimally buffered according to (5-12) and (5-13) with $i=1$. This defines the upper bound on the repeater resources or the HK product.
- (b). Depending on the design bandwidth requirement and area and power constraints, H and K are chosen for each line.
- (c). The single-variable function (5-38) is plotted against the values of s that are allowed by the technology, and the value that most closely resembles a zero represents the approximate optimal inter-wire spacing. This value is substituted in (5-39) to yield the matching wire width.

- (d). With these approximate values as a starting point, a few simulations are carried out with the exact capacitance equations for T in (5-36), to find the true optimal solution.

It must also be stated that the validity of the lossy capacitive line model must be established at the start of the analysis, which can easily be checked by any of the metrics proposed by a number of authors [Deutsch97], [Ismail99], [Lin00], [Krauter99], [Banerjee01]. In the experiments carried out by the authors, it was evident that inductive effects could be safely ignored for the wire widths that were close to the optimal point, and indeed even for those wires much fatter than the wires in this region.

5.4 Error-Control Coding for Lossy Lines

So far well established techniques of reducing delay over long resistive interconnect, namely repeater-insertion and wire-sizing, have been examined. The data-dependent cross-talk and delay suggests that some coding may possibly be used to gain an improvement in bandwidth. In this section the possibility of using error-control coding, specifically binary BCH encoding, is investigated. For this investigation the inductive parasitics of the lines are included, and the boundary conditions are different from




Figure 5.9.c: Ideal Buffering

Figure 5.9: Roots of the function $f(s) = T + (w+s) \frac{dt}{ds}$

the previous example used. Also power-supply noise is modelled by corrupting the ground and power rails by an appropriate noise spectrum.

5.4.1 Noise Analysis and Modelling

A distributed *RLC* line (shown in Fig. 5.10) which models the effects of delay, capacitive and inductive cross-talk and power-supply noise, is used for the simulations. A pseudo-random-bit-sequence (PRBS) is encoded, and fed into the central line while two other PRBSs are fed to the adjacent lines. Sampling at the output is by means of threshold detection at half amplitude. The output from the detector is then decoded to give the final output.

Power-supply noise is principally the difference in voltage caused by the drop across the parasitic impedances of the power supply network and is essentially a deterministic signal, since its spectrum depends to a large extent on the current profile of the switching logic blocks. The charging/discharging currents create a peak considerably higher than the average, causing inductive drops. A statistical approximation similar to that given in [Dally98] is used to determine the logic current profile. The chip is assumed to be composed of 50k gate modules (in agreement with the premise of several block-oriented schemes such as [Dally01], [Sylvester98]), and a certain fraction of these gates is assumed to switch in any given clock cycle. These switching gates are distributed in a triangular manner over several stages, resulting in a triangular current profile, the dimensions of which depend on the rise time and the switching load. The power supply network can also be subject to *LC* ringing. Typical values result in peaks in the spectrum at around 85 MHz corresponding to the package resonance frequency and at around 370 MHz corresponding to tank ringing.

In the test channel, the effects due to the transient load currents -calculated as explained above, and *LC* ringing are modelled by adding components at the relevant frequencies to random noise. Hence the white spectrum of random noise is effectively coloured by the spectral content of the noise caused by the charging and discharging transients and by the package resonance frequency components. Different composite noise files corresponding to different random distributions corrupt the ground taps.

5.4.2 Boundary Conditions

The hardware comprising the encoding and decoding circuitry has a complexity that is the order of hundreds of transistors rather than the two transistors of an inverter. As such the line lengths and conditions under which error-control coding can be considered for on-chip signalling, just as with overdriving repeaters, are limited. Consider a




Figure 5.10: Victim capacitively and inductively coupled to 2 aggressor lines switching in random fashion




Figure 5.11.a: Aggressors switch in same direction




Figure 5.11.b: Adjacent interconnect grounded




Figure 5.11.c: Aggressors switch in random fashion

Figure 5.11: Eye diagrams for 700 MHz signalling Over 1.5 mm metal 1 interconnect for parasitics extracted for a 0.05 micron technology with minimum spaced wire geometry.

1.5 mm line in a hypothetical 50 nm technology with $w=100\text{nm}$, $t=210\text{nm}$ and $s=100\text{nm}$. If the repeaters are sized according to (5-12) and (5-13) with $i=1$, the resulting numbers are $k=2$ and $h=28$. If the traditional analysis is used with the entire capacitance being considered as a component to ground, $k=1$ and $h=16$. Shown in Fig. 5.11 are the eye diagrams built up with spice simulations for the victim line with two aggressors where all lines are modelled by distributed *RLC* lines as given in Fig. 5.10. All lines are driven by a single inverter sized to 30 times a minimum sized inverter (i.e. $h=30$). The first diagram is for case (f), while the second is for case (c) and the third is when the aggressors switch in *random* fashion. That is, PRBSs with different seeds are fed to the three lines. This third eye diagram corresponds to what would occur in an actual situation. It can be seen that the combination of the cross-talk and power-supply noise almost completely closes the eye. To investigate the performance of ECCs a line that is 2mm long is chosen, where the eye diagram is completely closed. The objective here is to see what impact if any, ECCs, specifically binary BCH codes, will have in the lossy environment that is peculiar to semi-global length on-chip interconnect.

5.4.3 Genesis of Binary BCH Codes

One of the most important families of ECCs is BCH codes, which are a powerful and popular class of linear cyclic block codes well suited to coping with random errors. Their popularity is due in fairly large part to the existence of computationally efficient and easily implementable decoding procedures. A background in finite field arithmetic is necessary to understand the genesis of BCH codes [Berlekamp68], [Lin83], [Rorabaugh96]. The basic idea is that the field consists of q elements, and the operations of addition and multiplication are defined to conform to certain rules. The field of q elements is depicted as $\text{GF}(q)$ where *GF* stands for *Galois Field*. When q is the integer power of 2, $\text{GF}(2^m)$ are formed and these extension fields are the basis for working with BCH codes. Now the elements are defined in terms of a primitive element α which is a root of a primitive polynomial.

For a linear block code to be capable of correcting t errors per block it must have a minimum Hamming distance of $2t+1$. Binary BCH codes are built by constructing the generator polynomial so that its roots contain $2t$ consecutive powers of β where t is the number of errors per block to be corrected and β is an element of order n from $\text{GF}(2^m)$. Here n is the number of bits per code word. The encoder takes the input word and imparts the necessary redundancy. The decoder receives the transmitted word and calculates its syndrome, and maps it to an error pattern. It is the second step in the decoding procedure which is computationally heavy and poses area and speed constraints. It involves formulating a set of simultaneous equations from the syndrome, the solution of

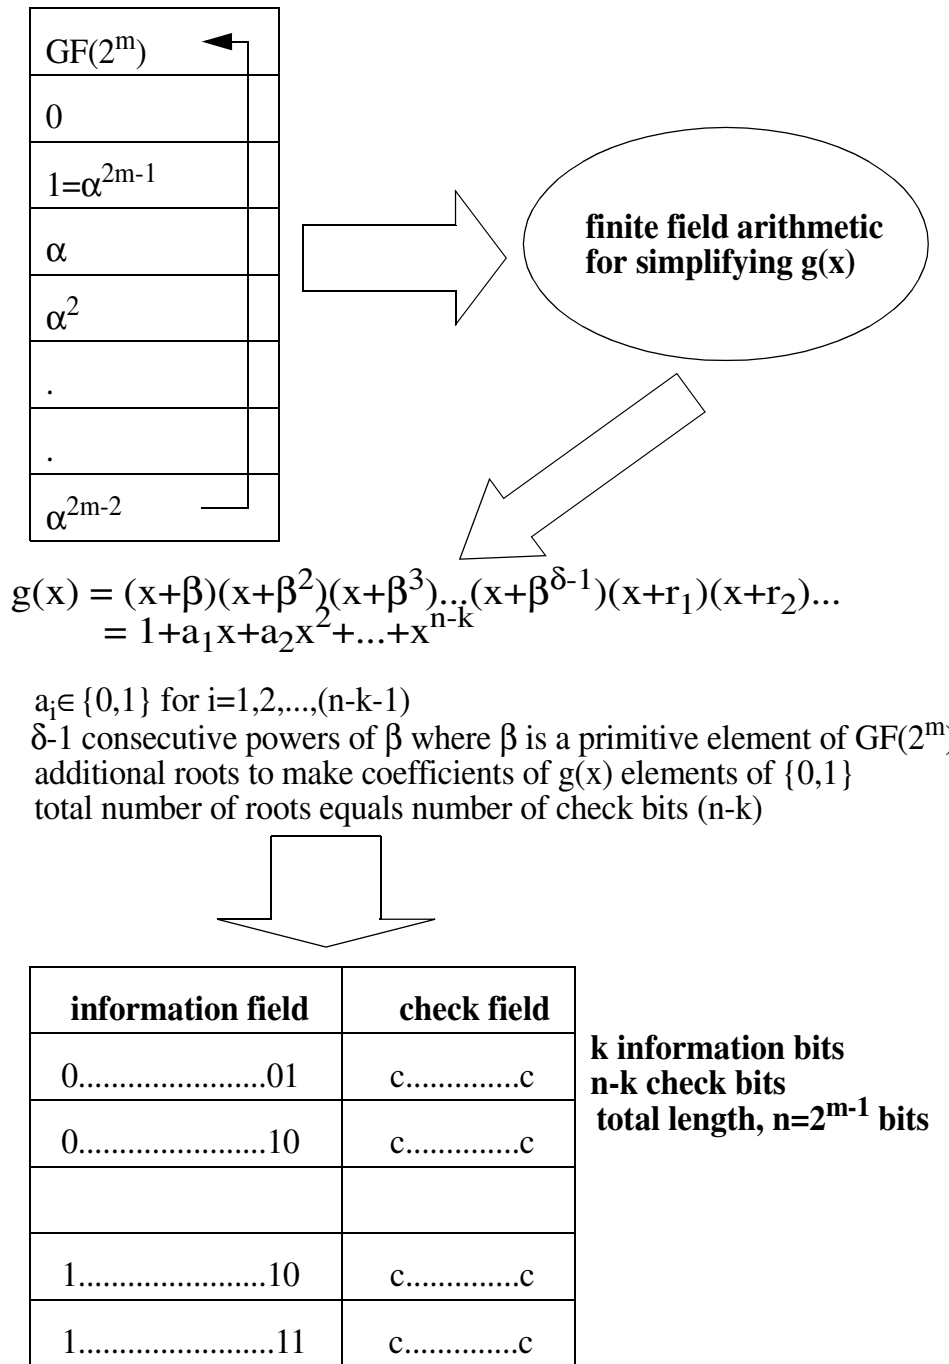



Figure 5.12: Genesis of Binary-Narrow-Sense-Primitive BCH Codes



A: 4/7 single ECC; **B:** 7/15 double ECC; **C:** 16/31 triple ECC; **D:** uncoded data; **E:** 21/31 double ECC; **F:** 63/51 double ECC

Figure 5.13: Curves of BER with different codes in the face of ISI, cross talk and power supply noise.

which gives a connection polynomial, which in turn gives the location of the errors by its roots. Given in Fig. 5.12 is a diagram depicting the manner in which the codes are constructed, which is reproduced from [Rorabaugh96].

5.4.4 Coding Gain

The increase in the number of bits introduced by the code reduces the overall information throughput. To maintain the same effective information rate it is necessary to increase the energy per information bit E_b by an amount proportional to $1/\rho$, the code rate. In Fig. 5.13, curves of Bit Error Rate (BER) for different generator polynomials are plotted against the peak-to-peak power supply noise for a normalized E_b .

It can be seen from an inspection of this figure that a properly designed code can result in significant coding gains. The BER for uncoded data over a 2mm long inter-

connect in a $0.05\mu\text{m}$ technology with minimum spacing at a frequency of 500 MHz, is approximately $4 \text{ in } 10^4$. When encoding with a rate $4/7$ single ECC generated from an extension field of degree $m=3$, which is basically a Hamming code, it is necessary to transmit at a bit rate of 500MHz times $7/4$ (reciprocal of ρ), to maintain the same information rate. This results in a much worse performance, with the BER increasing to almost $3 \text{ in } 10^2$. A rate $7/15$ double ECC over $\text{GF}(2^4)$ results in a marginal improvement, but is still much worse than the uncoded BER. However a rate $21/31$ double ECC over $\text{GF}(2^5)$ results in an improvement of the BER to approximately $1.5 \text{ in } 10^4$. This is a three-fold improvement over the uncoded BER. This can be explained by the fact that the increase in ρ caused by the increase in complexity of the extension field allows the same information rate to be maintained at a signalling frequency lower than with the rate $7/15$ code. The most dramatic improvement is given by the rate $51/63$ double ECC which causes a drop in the BER to almost $3 \text{ in } 10^5$. This is again a consequence of the increase in ρ .

It is also interesting that a rate $16/31$ triple ECC performs better than a rate $7/15$ double ECC with approximately the same ρ . This seems to indicate that errors are more likely to occur in groups of three rather than two and serves to emphasise the fact that it is rather difficult to formulate a mathematical model for the channel and accordingly select a code, as is usually done in communication applications. The nature of the errors depends very much on the layout and on the correlation between bit streams. Simulation is thus an invaluable tool in selecting a proper code. Deciding on the appropriate code depends very much on the application.

5.5 Summary and Conclusions

In this chapter signalling issues over lossy capacitive lines that are representative of global and semi-global length interconnect in DSM circuits have been discussed. The expressions introduced in Chapter 2 to calculate the parasitics, and the repeater and delay analysis formulae outlined in Chapter 4, were then used to investigate the optimum arrangement of wires to yield the maximum throughput for a given metal resource. For a parallel wire configuration, several factors combine to affect the delay in various ways. Increased parallelism is desirable in general, but when the total area that is allowed for the wires is constrained, this results in smaller, more tightly coupled wires, increasing cross-talk, and causing greater line delay. Repeater insertion and especially area constrained repeater insertion further complicates the issue. However a method of analysis that takes into account all these factors has been demonstrated, and used to show that there is a clear optimum configuration. Because of the closed form nature of

the expressions presented, this optimum can be predicted easily by means of an iterative algorithm.

Additionally simplified versions of the equations have been employed to produce a single-variable function of inter-wire spacing s , and a companion function for wire width w , the roots of which give a solution that is quite close to the true optimum. This approximate solution can be used as a starting point for simulations with the exact equations to provide the correct solution with 1 or 2 iterations. It was also shown that for wide buses, the optimal wire width and spacing depends on the repeater constraints and length, but is independent of the total width. The results presented here can conveniently be used to optimise on-chip buses.

Finally, an alternative to repeater insertion, and other traditional methods, namely error-control-coding, was investigated for performance in a typical on-chip environment. The lossy environment was carefully modelled, and performance metrics in the form of BER were extracted by means of simulations. The curves show that a coding gain is possible. Further investigation is necessary however, both in modelling the environment and analysis of the performance, and also in quantifying the hardware cost against that of conventional repeaters.

6. Designing SoC Communication Networks

The physical details related to the implementation of packet-switched Networks-On-Chip are discussed. All major issues are identified, and the feasibility of implementation in a typical 65nm DSM technology is investigated.

6.1 Introduction

In the DSM regime, electrical level issues that affect signalling, timing, power and noise are challenging established design procedures. Hence architectures that exploit locality and standardise on-chip communication via shared protocols are receiving a lot of attention. Such architectures are referred to as *Network-on-Chip (NoC)* architectures and have been proposed by several research groups ([Dally01], [Sgroi01], [Benini01], [Sylvester99a], [Hemani00]). A NOC architecture can consist of a few dozens to several hundreds or thousands of resources, communicating with each other via the on-chip network. A resource may be a processor core, a DSP core, an FPGA block, a dedicated HW block, an analog or mixed signal block, or a memory block of any kind such as RAM, ROM or CAM.

This chapter examines in detail the physical issues related to the implementation of such mesh-based architectures. Timing, power and area issues are considered. Some of the issues raised are: is it actually wireable without imposing too many constraints on resource IPs? If so, what are likely wiring schemes, and what kind of performance can be expected? What are the area and power penalties for the overhead of the on-chip network and its switches, and what are the trade-offs involved? How does link bandwidth trade-off against the link power consumption? To answer these questions and others, a case study of two likely architectures in a typical 65 nm technology is carried out.

6.1.1 Background

In the deep sub-micrometer (DSM) regime, electrical level issues that affect signalling, timing, power and noise are challenging established design procedures. Expected trends in the future evolution of VLSI systems can be codified into the following points:

- (a). Moore's law will continue to hold for another ten years [ITRS01].

- (b). Single processors will not be able to utilize the transistors of an entire chip, and a single synchronous clock region will span only a small fraction of the chip area [Sylvester98], [Hemani99].
- (c). Applications will be modelled as a large number of communicating tasks, where the tasks may have very different characteristics (such as control or data flow dominated) and origins (IP re-use from earlier products or external sources) [Szyperski98], [Gajski99].

If as is widely predicted these premises hold true, a large number of different kinds of blocks -each of the size of a few hundred thousand gates- will constitute the computational resources. For acceptable performance they have to be connected efficiently, and several authors have argued eloquently that a regular on-chip network based on packet switching is the most likely scenario for chip architectures in five to ten years [Dally01], [Sgroi01], [Benini01]. It eases the expected bottlenecks of complexity and wire delay in nanometre technologies, and promotes extensive re-use of design cores through standardization of on-chip communication. The objective is to achieve physical-level and architectural-level design integration. This implies that physical layout and implementation issues need to be kept in mind while taking architectural decisions, or in other words, the architectural design has to be carried out within the constraints of a floor plan. Such a standardised framework increases the fractional non-recurring component of the cost, and also eases the pressures of meeting time-to-market demands imposed by very competitive and fickle markets.

6.1.2 Feasibility Study

The exercise of carrying out a feasibility study for NoC implementations in the DSM regime is basically twin faceted; firstly, it is necessary to build or modify models that accurately capture the behaviour of active and passive devices in the given technology node. Secondly, it is necessary to decide upon an architecture that is representative in a general sense.

The first is the science of *technology extrapolation* which has received a great deal of attention over the past few decades. Its importance is due to the fact that not only does predicting future trends give us an idea of what is achievable, but also has a strong influence on the evolution of future VLSI systems. A major collaboration of both industry and academia has resulted in the roadmapping venture of the International Technology Roadmap for Semiconductors (ITRS), the latest version of which is [ITRS01]. This a very important source of information, setting targets and guidelines for future evolution.

Influential technology extrapolation systems developed 10-15 years ago, are described in [Bakoglu87] and [Sai-Halasz95]. More recent second-generation systems include [Eble96], [Geuskens97] (available as a web-based tool at [RIPE97]) and [Sylvester99b] (available as a web-based tool at [BACPAC98]), along with Roadmap-related efforts [ITRS01]. A collective resource which gathers together these models and others is [GTX00]. Typically, each system provides estimates of chip area, maximum clock frequency, power dissipation and other parameters based on a small set of descriptors spanning device and interconnect technology through system architecture.

The parameters of the 65 nm technology are obtained by following guidelines outlined in the ITRS, and by scaling from an existing technology. Some assumptions about unknowns are made in order to facilitate analysis. This may make for a somewhat rough and ready approach on occasion, but it is adequate to provide representative figures for likely future CMOS technologies.

The second part of the problem obviously has no right or wrong solution. Two architectures are used for the study, and these were chosen from a study of the literature describing likely topologies, as being most likely to give representative figures for parameters of interest such as bandwidth and power consumption. This is discussed in the next section.

6.2 NoC backbone

6.2.1 Architecture

A perusal of the literature shows many works that have elucidated the NoC concept and a layered protocol [Dally01], [Sgroi01], [Benini01]. Of these and others, the most attention to the physical level is paid in [Dally01]. It describes a folded torus topology that fits well to VLSI implementations with a two-dimensional layout and limited wires. The proposed routing layout places the network wires on top of the resources in dedicated metal layers.

This is the most intuitive layout, but there are in general a myriad of ways to lay out this NoC backbone. Two extremes can be identified: the first is where, as mentioned, the network interconnects are routed over the resource (the “thin-switch architecture” shown in Fig. 6.1.b), and the second is where the wires run in dedicated channels (the “square-switch” architecture shown in Fig. 6.1.c). The former has no area overhead associated with the network wires, but routing the wires over the resource does impose a few restrictions on the design methodology of the resource. The placing of repeaters

for example may interfere with importing IP cores. Also, to avoid routing congestion over the resource it may be necessary to dedicate one or two metal layers to the network interconnects, which may pose problems in distributing the power and ground networks depending on the number of metal layers available. Even with dedicated metal layers, vias to I/O (power, ground and signal) pads will restrict the number of available wiring tracks. On the other hand, laying restrictions on the routing imposes area overheads, but routing the network and ensuring signal integrity over its wires is straightforward.

Both these two alternative global routing layouts are analysed here. The analysis is conducted for a simple mesh network with only direct neighbour connections but is also applicable to the folded torus topology of [Dally01].

6.2.2 Network Protocol

The topology and protocol of the NoC where the switches are connected to their direct neighbours only, is described in [Millberg02]. This section provides a brief overview of the protocol details relevant to the study. The NoC backbone consists of *Resources* and *Switches* organised in a Manhattan-like structure with a one-to-one correspondence (Fig. 6.1.a). All resources are equipped with a *Network Interface (NI)* to communicate between the resource core and the network. The NI handles all communication protocols to make the network as transparent as possible to the resources. To accommodate a reasonably sized network (more than 25 resources), a bus width of 128 bits in each direction for the switch-to-switch and switch-to-resource connection appears suitable [Dally01].

The signalling between switches and between switch and resource is *Packet* based. Each packet consists of *Header* and *User Data*. The user data is the actual message and the header is additional information needed for the switching of the packets over the network. The header is further divided into the following data fields: *Source Address (SA)*, *Destination Address (DA)*, *Packet Sequence Number (PSN)*, *Process ID (PID)* and *Hop Counter (HC)*.

Since the message can vary in size and the size of the packets are fixed, one or more packets will be needed to transmit one message. This is usually referred to as segmentation. The switching policy of the packets is datagram based, i.e. each packet is treated independently, with no reference to preceding packets [Stallings94]. This means that the switches must make an independent routing decision for each individual packet. As a result of this, packets with the same source and destination may possibly not follow the same route.




Figure 6.1.a: Logical mapping of network

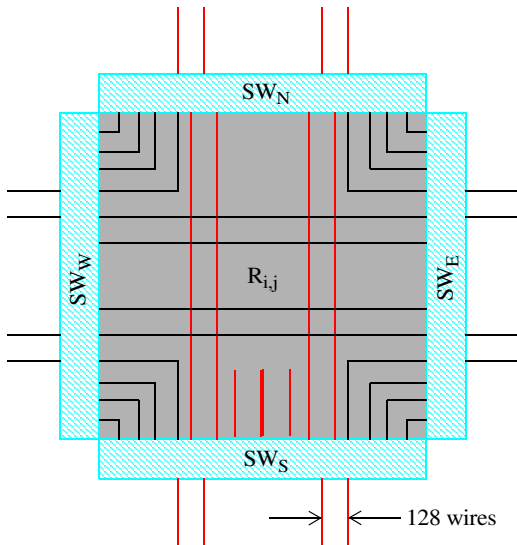


Figure 6.1.b: Thin-switch architecture

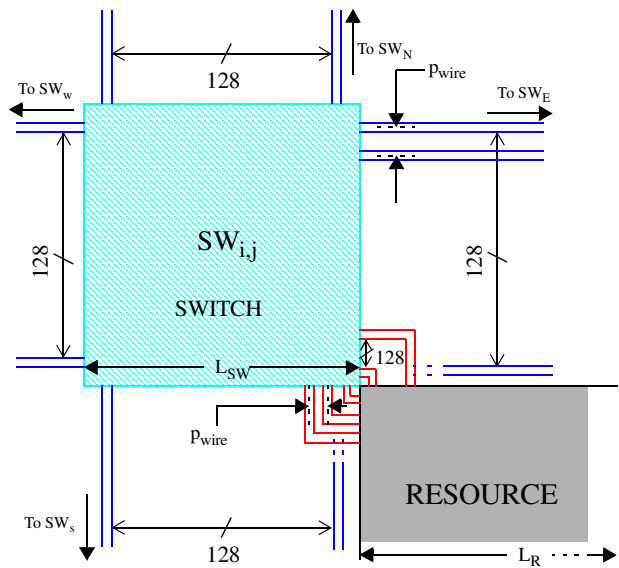


Figure 6.1.c: Square-switch architecture

Figure 6.1: Network-on-chip Backbone

As the Manhattan structure is a 2D-mesh it is always known beforehand where each resource is physically located, thus making the addressing of the resources easy. This property makes the design of the switches simple and the need for routing tables is reduced or eliminated. Hence the DA and SA fields in the packets can be considered to consist of two separate fields - the *Row Address (RA)* and *Column Address (CA)* fields. Since no pre-planned route is established before the packet is sent - as would be in the virtual circuit switching approach [Stallings94] - a local routing decision must be dynamically made in the switch for each packet. This decision is made based on the DA and HC (and possibly the SA). All these data fields reside in the header of the packet. The DA is the final physical destination address whereas the HC is a counter tracking the number of hops a packet has made, i.e. the HC is increased by one for every hop taken. In case of a competitive situation in the switch, the packet with the higher HC is given priority. The main reason for not choosing a virtual-circuit approach is that the datagram approach more easily adapts to changes in the network such as congestion and dead links.

6.3 Modelling issues

6.3.1 Technology Scaling

Based upon the technology roadmap [ITRS01] the following properties of a representative 65nm technology available in the year 2008 have been obtained. The process supports from 8 to 10 layers of metal made of a copper-alumina alloy having a resistivity of $2.5\mu\Omega\text{cm}$. The wires of the lower levels are 210nm thick, giving a sheet resistance of $0.12\Omega/\text{square}$. They have a minimum width of 100nm and a minimum pitch of 200nm. It is expected that area array bonding techniques will be available for power and I/O connections, and pads comprising solder balls of $40\mu\text{m}$ at $100\mu\text{m}$ centre distances are assumed. The power supply is surmised to be 0.9 V, and the local clock frequency to be 3 GHz.

The Thévenin equivalent output impedance R_{drv} of a minimum sized inverter driving a similar load is estimated to be $7\text{ k}\Omega$ while its input capacitance C_{drv} is 1 fF. A representative gate (2-input nand gate) will occupy $1.6\mu^2\text{m}$ of area. A typical gate load is usually assumed to be four gates, which results in a load capacitance of 4 fF.

The number of gates that can be accommodated in a given area can be estimated by using the following scaling equation.

$$N_{new} = N_{old} \times \left(\frac{A_{new}}{A_{old}}\right) \times \left(\frac{\lambda_{old}}{\lambda_{new}}\right)^2 \times \alpha_s \quad (6-1)$$

Here N denotes the number of gates, A the area, and 2λ the feature size, with the subscripts *old* and *new* referring to the technology in which the design is currently implemented, and the future technology respectively. The factor α_s ($1 \geq \alpha_s > 0$) is used to account for integration losses or gains in the scaling.

6.3.2 Switches and Inter-Switch Links

Square Switch

According to the communication protocol, 128 wires come into and go out of the switch in each direction. Also an additional 128 wires go into and out of the resource to handle the resource's communication with the network. For the thin switch these wires translate to two extra links (Fig. 6.1.b), while for the square switch they are situated on the two sides of the switch that are closest to the resource (Fig. 6.1.c). One possible arrangement is that each incoming and outgoing wire in the switch is latched by a flip-flop, and each outgoing wire is fed from a 5-to-1 multiplexer. Since each switch has 10 sets of 128 wires, this translates to a total of 14,720 gates. Inside the switch, there is also some additional decision circuitry consisting of adders, subtractors and comparison units, adding up to an estimated 2000 gates. Thus the control logic for the switch is approximately 17,000 gates [Nilsson02], with each gate occupying approximately $1.6 \mu\text{m}^2$. Leaving a routing overhead of 30% for the control logic gives a total of $36,000 \mu\text{m}^2$, which translates to an approximate switch size of 0.2 mm X 0.2 mm.

Thin Switch

The logic in the thin switch has the same functionality as the square switch, but is distributed across a much larger distance, the length and breadth of a tile. This has implications for power consumption and latency, as the arbiter logic is situated only on one side of the tile, as suggested in [Dally01]. All header data (12 bits in our protocol) need to be forwarded to this side for arbitration to take place. Then the decision is forwarded to the side where the data is actually multiplexed into the outgoing flip-flops. This means that the actual delay is twice the line delay, unless the arbiter logic is du-




Figure 6.2: Wire model for inter-switch network links

plicated at each side, which would increase the switch area and power consumption. Another possibility is to pipeline the decision procedure into two stages, when the latency is twice the line delay, but the throughput is unchanged. This impacts the power consumption in a different way; because of the distributed nature of the switch, the average power consumption of the switches has to include the power consumed in forwarding the header from any of three sides of the tile to the arbiter on the fourth side. Since this utilizes the network wires, the power consumption of the switches varies for different repeater insertion strategies. In this study, pipelining is assumed to allow a direct comparison with the square-switch architecture.

The area of the switch logic on one of the sides of a tile is determined by wiring considerations. This logic is spread across the length of a side, and a breadth of approximately $15\ \mu\text{m}$ on the western side, and about $10\ \mu\text{m}$ on the other sides. This translates to an area overhead of under 5%.

Network Links

An essential part of this analysis is the physical modelling of the network links. Many different signalling techniques have been proposed in the literature including low-swing, differential- and current-mode techniques, but the most common and robust is full swing CMOS signalling with inverters as repeaters. This is the signalling convention adopted here, with analysis techniques similar to those used in Chapter 5. The capacitively coupled RC model that was introduced in Chapter 4 is used for the wires as they are thin enough and long enough that this model is valid. The notation adopted is the same; namely that k refers to the number of repeaters (inverters) on a single line including the first driver, and h to the size of the inverter in terms of multiples of the W/L ratio of a minimum sized inverter. This arrangement is sketched out in Fig. 6.2. The inverters are modelled as resistor-capacitor combinations -with parameters as given in section 6.3.1- that scale linearly with size.

That neglecting inductance is justified for the wire pitches and lengths of interest can be verified by applying the metrics described in a number of well known works ([Ismail99], [Deutsch97]) as was done in the bandwidth analysis in Chapter 5. Also in mapping the geometry to the parasitics, the same models that were used in Chapter 5

and reported in [Zheng00] are used. The delay analysis uses timing models that consider worst-case cross-talk and was introduced in Chapter 4, to consider the merits of fat wires with no shielding against thinner wires with shielding.

$$t_r = k \left[0.7 \frac{R_{drv_m}}{h} \left(\frac{C_s}{k} + h C_{drv_m} + 4.4 \frac{C_c}{k} \right) + \frac{R}{k} \left(0.4 \frac{C_s}{k} + 1.51 \frac{C_c}{k} + 0.7 h C_{drv_m} \right) \right] + \frac{t_{r,i}}{2} \quad (6-2)$$

6.3.3 Resources

Now in the square-switch architecture, the dedicated communication channels are the same width as the side of a switch. Hence the tile size has to be large enough that the area overhead is not too high. A 2 mm X 2 mm tile size gives an overhead of 20% for the network, which would seem to be an upper limit. In order to be able to compare between architectures, the same tile size is used for both. This choice may seem somewhat arbitrary, but it is a reasonable one, allowing good sized resources to be housed in a single tile, as shown below.

A high performance ASIC in a 0.35 μm technology comprises approximately 1M gates on a 16 mm X 16 mm die [Dally98]. Using (6-1) to scale to a 2 mm X 2 mm area for the same integration efficiency ($\alpha_s=1$) results in a gate count of 450k. A more representative gate count for a tile is obtained by relaxing α_s to account for standard cell-based designs, to approximately 200k gates. This value is used for the power estimations. A resource consisting of significantly more gates can occupy multiple tiles, when the analysis is still valid.

6.3.4 Power Estimations

The power consumption in the chip is composed of the portions dissipated in the resources, and the network. Most of the power consumed in the resource is due to switching logic. The average current consumed in a cycle [Dally98] is

$$I_{avg} = \frac{N_s C_{ld} V_p}{t_{clk}} \quad (6-3)$$

where N_s is the fraction of gates switching in one direction in one clock cycle, C_{ld} the average capacitive load of a gate, V_p the positive rail voltage, and t_{clk} the clock period. A typical gate load in the 65 nm technology is estimated to be 4 fF, the power supply to be 0.9 V, and the local clock frequency to be 3 GHz (see section 6.3.1). The number of switching gates can be computed by assuming that half the gates switch in any given cycle, with equal numbers in each direction. This means that for our representative 200k gate resource, $N_s = 50,000$, giving an average current of 0.64 A. The power is then VI , or 0.52 W per resource. Assuming a likely die size of 3 cm [ITRS01], the total power consumed in the resources is roughly 120 W.

The average power consumed in a single wire in an inter-switch link is calculated by the following expression

$$P_w = k(h \times C_{drv} + C/k) \times V_p^2 \times f_b \quad (6-4)$$

where k , h and C_{drv} are as defined in Section 6.3.1, and f_b is the bit period. The total average power consumed in the network links is given by

$$P_{nw} = P_w \times N_w \times \beta \times \delta \quad (6-5)$$

where β and δ are coefficients that represent the fraction of bits that switch on a given link in one direction (assumed to be 0.25, i.e. half the bits switch, and half of the switching bits switch up, while the other half switch down), and the fraction of links that are active in any given cycle (assumed to be 0.5). By using an analysis similar to that used for resources, the power consumed in the logic of the switch can be calculated. Additionally the wire-load model plays a factor as explained in section 6.3.2.

6.4 Analysis and results

6.4.1 Square-Switch Architecture

Sketched out in Fig. 6.1.a is the switch and resource arrangement. As mentioned, the area overhead for this architecture is 20%. To reduce this, there are two possibilities:

- (a). let the resource area extend under the wire channels, creating a compromise between architectures 1 and 2;

- (b). let the switch extend into the wire channels (i.e., shape it like a '+' sign instead of a square).

The former will not be investigated as there are a great variety of different geometrical arrangements possible and the intent here is to investigate the two extremes. The latter trades off area overhead against link bandwidth. Consider that the switch extensions (peripheral legs of the '+') have linear dimensions equal to the square in the centre, when 5 times the area of the central square is available. For the $36,000 \mu\text{m}^2$ total area mentioned above, this translates to a switch that is $85\mu\text{m}$ or say $100\mu\text{m}$ square, with extensions of the same size into all four channels. Now the area overhead is only 10% but the channel width for the communication link is halved, resulting in a lower bandwidth. However because of the non-linear scaling of the parasitics with wire width, this is only a small percentage decrease (Chapter 5).

All metal layers can be utilised to route the network wires. If it is assumed that 10 metal layers are available for example, and that the top three are reserved for power, ground and clock distribution (which is conservative), 7 layers remain for routing the communication link. One possible implementation would be to route the 256 inter-switch wires in 4 layers, with explicit signal return planes between each signal layer. This also provides for shielding between metal layers, which is important as the wires on all layers are parallel. Then for the western and eastern sides of the switch, 64 wires need to be routed on a metal layer, giving a wire pitch of approximately $3\mu\text{m}$. Since the 128 wires between the resource and switch on the southern and eastern sides run only for a short fraction of the resource length, they can either occupy one of the shielding layers briefly, or be routed on the four signal layers along with the inter-switch wires at less than maximum pitch. Once all the wires have gone into the resource, the remaining wires (from the inter-switch link) spread out to take maximum advantage of the space available. The distance for which wires are congested will be a very small fraction of the total length of 2mm, and can be neglected for timing purposes.

Given that the maximum pitch is $3 \mu\text{m}$ for 64 wires on a single layer, an arrangement that appears to maximise the bandwidth for the fully-square switch is the following: the signal wires are $2.5 \mu\text{m}$ wide on $3.1 \mu\text{m}$ centres, and between the signal lines are thinner shielding lines of the minimum width of $0.1 \mu\text{m}$, on the same $3.1 \mu\text{m}$ pitch. The signal wires are also vertically shielded with clearly defined return paths. For the '+' shaped switch, the signal wire width is $0.9 \mu\text{m}$, while the pitch of the signal and shielding wires changes to $1.6 \mu\text{m}$. Locating repeater stations in the channel is ideal in terms of utilising space. Since the wire pitch is not sufficient to fit all repeaters horizontally, they can be placed in zig-zag fashion, so that the channel is packed with repeaters.

Given in Table 6.1 are the delays (rounded to the nearest decade) corresponding to

Table 6.1 Slew rates for wiring schemes of square-switch architecture

Switch shape	Area overhead	k	h	t_r (psecs)	pulse width
square	20%	3	301	70	$1.2t_r$
+	10%	4	123	85	$1.2t_r$

the two cases along with repeater details (rounded to the nearest half decade) for the maximum possible bandwidth. As a point of comparison, the line delay at the speed of light in SiO₂ is 13.3 ps. Since the environment is very tightly controlled, with shielding on all sides, and the worst-case delay is considered, a factor of 1.2 is deemed sufficient to map the rise-time to the pulse-width.

Shown in Fig. 6.3.a and Fig. 6.3.b are plots detailing the trade-off of bandwidth for power, for the network. The power consumption of the switches is also shown, as being invariant for the various repeater insertion schemes.

6.4.2 Thin-switch Architecture

The physical layout is that a resource block is surrounded by four thin strips of logic, each of which is connected to four other switches (Fig. 6.1.a). In addition, one side has wires going into the resource. The shorter dimension of the switch is defined by the layout of the logic contained in it. Since the control logic will be the same for both architectures, the total area is also the same, but as the logic will be distributed along the full side of the region, the area overhead for the switches is only about 5%.

Instead the main consideration will be routing congestion over the resource. The vertical and horizontal wires of the network need to be wired in two metal layers since the lines cross as shown, restricting either the number of metal layers to be used by the resource blocks or the wiring freedom in terms of the available fraction of a metal layer. Although the wires may occupy less than a quarter of each metal layer, having the resource share them for local signal wiring would impose fairly severe restrictions on the design methodology of the resource. It would appear to be more viable to share the metal layer between the network and power distribution grid, which would also provide some form of mutual decoupling. Devoting two entire layers to the network if possible would of course be the ideal solution.




Figure 6.3.a: Square-switch architecture: fully-square switch




Figure 6.3.b: Square-switch architecture: plus-shaped switch




Figure 6.3.c: Thin-switch architecture: dedicated metal layers




Figure 6.3.d: Thin-switch architecture: shared metal layers

Figure 6.3: Performance and Power Consumption for Different Wiring Schemes

Now pins to power, ground and signal I/O pads over the chip will limit the number of wiring tracks for all metal layers. Practice shows that between 20% and 30% of a metal layer will not be wireable. Two cases are considered:

Table 6.2 Slew rates for wiring schemes of thin switch architecture

Scheme	k	h	t_r (psecs)	pulse width
Dedicated metal layers	3	307	70	$1.6t_r$
Shared metal layers	3	176	80	$1.6t_r$

- (a). two metal layers are devoted to routing the network;
- (b). only a certain fraction of each metal layer is utilized for the network.

In the first case it is estimated that signal and power and ground pins to pads will render 20% of a metal layer unusable, while in the second that 70% is unusable due to additional utilisation for the power distribution grid. In both cases 512 wires need to fit on one side (of 2 mm length) on a single metal layer, comprising the links to two switches as can be seen from Fig. 6.3.c. As in the dedicated channel architecture, the wires into the resource may either occupy the same metal layer at the minimum pitch or a different metal layer as they are very short in comparison with the other wires.

In comparison to the overhead of the NI, that imposed by repeater stationing inside resource IPs would be relatively minor. Hence no restrictions are assumed on repeaters. Table 6.2 gives likely details for minimum slew rates. In both cases, the wire pitch is 4 μm . For the wiring scheme with dedicated metal layers, signal wires and shielding wires are 2.6 μm and 0.1 μm wide respectively. For the scheme where the network and power grid share metal layers, signal wires and power wires (which serve as mutually-decoupling shields for the signal wires) are 1.3 μm wide. The power-versus-bandwidth trade-off is detailed in Fig. 6.3.c and Fig. 6.3.d, as is the power consumption of the switches.

6.5 Discussion and Conclusions

This chapter considered the physical issues related to the implementation of a packet-switched NoC. By reviewing proposals in the literature from several research groups, two simple but representative architectures were identified. Based on the ITRS, parameters for a technology expected in 2007 were derived, and used to refine

the architectural details (size of a tile and switch) by considering the integration density of devices. Finally cost and performance metrics for two likely architectures in a future technology were extracted, for full-swing CMOS signalling. In particular, curves quantifying the power-bandwidth trade-off for the network were derived.

These are scatter-plots of the power consumed in the network for various repeater insertion strategies. The curve at the bottom represents the power consumed in the switch alone. In the thin-switch architecture, because the distribution of the switch is over a tile, its activity is dependent on the repeater insertion strategy of the network. Hence this is reflected in the power consumption plot, which increases with increasing network bandwidth. In the square-switch architecture, the activity in the switch is dependent only on the local clock frequency, and hence the power consumption is, to a first order, independent of the network bandwidth.

The curves at the top represent the total power consumed in the network (consisting of the power consumed in the wires, repeaters, and switch). The multiple lines correspond to different repeater insertion strategies. For example, $k=2, h=30$ and $k=3, h=20$ result in the same hk product, and approximately the same power consumption. However they result in different slew-rates, which define two separate coordinates on the plot.

Some characteristics of the plots also require clarification. It can be seen that the power consumption for a given bandwidth at the lower end of the spectrum is not the same for the different cases. For example, in the square-switch architecture, the power corresponding to 400 Tbits/sec is roughly 35 W for the plus-shaped switch, and around 60 W for the fully-square switch. This is simply a consequence of having chosen fatter wires in the latter arrangement to take advantage of the larger channel. These fatter wires allow a higher peak bandwidth, but are not the optimal width for lower values of bandwidth. This is very logical, as a designer would not choose the extra area overhead associated with the fully-square switch, unless he or she also required the extra bandwidth associated with it.

One interesting conclusion of this study is that the power consumed in the network is a considerable fraction of the power consumption of the full chip. As can be expected, it is possible to trade-off power for bandwidth, and for the maximum possible bandwidth (in the square-switch architecture) the power consumed in the network is more than twice the power consumed in the resource logic. Some of the advantages and disadvantages of the two architectures are given in Table 6.3.

The intuitive and obvious layout is the thin-switch architecture, but the square-switch architecture also recommends itself in certain aspects. Briefly the square switch architecture guarantees signal integrity and provides a higher link bandwidth at the cost

of a higher area overhead. The thin-switch architecture is more difficult to wire, has a lower link bandwidth, and has an intrinsically higher power consumption for the switch itself (because it is distributed across a tile and has longer wires), but has less area overhead. In both, the power expended in the network is a major fraction of the total power consumption.

The choice of architecture, whether one of the above or a hybrid of the two, will of

Table 6.3 Pros and Cons of the Two Schemes

Square-Switch Architecture	Thin-Switch Architecture
Area overhead of between 10% and 20% for network	Area overhead of roughly 5%
All metal layers can be freely utilized for resource	No. of available metal layers or available fraction of two metal layers reduced for resource
No routing/pin congestion over resource due to network	Routing/pin congestion introduced by network
Dedicated channel allows repeater insertion, shielding and explicit signal return planes, guaranteeing signal integrity	Repeater insertion and shielding more of a problem. More susceptible to noise coupling from above and below
Max. link bandwidth of 1.5 Tbits/sec in any direction	Max. link bandwidth of 1.1 Tbits/sec
Network power is a high fraction of the total power consumption of chip	Network power is a high fraction of the total power consumption of chip

course depend on the application. However this study has shown the feasibility of implementing the NoC concept under the physical constraints of interconnections in the DSM regime, and cost and performance estimates were extracted by considering the physical implementation in as much detail as possible.

This study was by necessity, conducted only to a first-order degree of accuracy. Also all numbers should be treated as ballpark figures, and be interpreted with regard to the boundary conditions considered. Additionally two specific points should be kept in mind. Firstly, the choice of architecture depends to a fair extent on the switching policy of the network. The thin-switch architecture for example is more suited to deterministic routing, while the square-switch architecture is tailored more to adaptive

routing. This study however considered the same protocol with adaptive routing for both architectures.

Secondly, although it is interesting and informative to see what the cost and performance of a NoC is for rail-to-rail, voltage mode CMOS signalling, it may not be the ideal choice for a NoC. As pointed out in [Dally01], the environment of the network can be tightly controlled so that more advanced techniques such as low-swing, booster-driven and current-mode signalling can be considered to reduce power consumption and reduce latency. Another possibility to reduce the power in the network is encoding of the data packets to minimise transitions. It is expected though that this study will provide a point of reference for future comparisons.

7. Conclusions

This chapter summarises the thesis and describes avenues for future work.

7.1 Summary and Conclusions

The Moore law scaling of device integration in IC manufacturing has led to major advances in chip functionality and complexity. It has also resulted in new challenges in managing complexity. The sheer profusion of devices and interconnects requires accuracy, and efficiency in computation. This thesis has addressed this issue by proposing new models that capture important effects in future technology nodes while still being computationally cheap, and proposing and examining new design methodologies.

The first contribution of the thesis is in proposing a general 2-pole-1-zero model for arbitrarily-coupled RC trees, that represents the minimum complexity for this class of circuits. It is an extension of a well known and widely used model for a simpler topology [[Horowitz84]], to a more complex topology, with refinements that guarantee stability. For the switching of the victim driver, it reverts to that model when the coupling capacitors are put to zero. The physical basis of the model is that geometric attributes of the actual waveform (step response), namely the area and first moment, are matched to the model response. Because there are three unknowns in a 2-pole-1-zero transfer function, this results in a family of curves. To obtain the particular solution for a given switching event, other attributes are used; for the victim switching, it is the sum of the open-circuit time-constants, for an aggressor switching, it is a combination of the moments for different switching events.

The accuracy of the models were tested against a circuit simulator, and also against more expensive moment-based models. They were found to be of an accuracy comparable to more expensive models for a number of representative test-beds.

The second part of this thesis addresses the important structure of long, multiple parallel wires, or data buses. A methodology is proposed to simultaneously optimise the repeaters and wires, and closed-form expressions that predict the optimal point under design constraints, and are dependent only on process parameters are derived. Because of the exhaustive possibilities available for simulations, these models should be very

useful in reducing the search space, cutting down the design time of complex systems, where long buses are ubiquitous.

Finally, a major new design paradigm being proposed by a number of research groups is the NoC. The NoC is expected to facilitate design re-use, and greatly reduce the turnaround time for large and complex designs by standardising on-chip communication, much as is done in board-based systems now. A lot of the published work in the literature describes higher level aspects of the NoC, and the third part of the thesis concentrates instead on the physical mapping. The implementation of a representative NoC architecture is considered in a future technology, and the general subject of implementing Networks-On-Chip under the physical constraints of interconnections in the DSM regime is addressed. Cost and performance metrics are extracted, which quantify design trade-offs such as network power and bandwidth. Quite a few assumptions are made to facilitate analysis, but it is expected that this study will serve as the basis for further investigations and comparisons, being the first of its kind to address the NoC architecture specifically.

7.2 Limitations and Future Work

The work described in the first part of the thesis, on modelling the response of arbitrarily-coupled *RC* trees, would be complete if bounds were derived, that limit each aggressor waveform between an upper and lower value, in the tradition of the landmark work describing bounds for simple trees [[Rubinstein83]]. An upper bound for the noise has been proposed in [[Devgan97]], but metrics describing an interval, and also giving indications of when a second order model can deviate appreciably from the true response, would be extremely useful. This is earmarked for future work.

The main issues to be investigated in the work described in the second part, on optimising on-chip buses, are the fidelity of the models that map the geometry to the parasitics, and the possibility of including inductive effects in the delay analysis. The first can truly be investigated only by manufacturing different test structures on a chip, and measuring the deviation of the true response from the predicted. The main challenge in the second is to accurately model the inductance of an isolated wire, which is currently being investigated.

The third part which describes the feasibility study, considers the most common and robust form of signalling, rail-to-rail CMOS signalling. As pointed out in [[Dally01]], the regular structure and controlled noise environment of the NoC has the potential to allow other, more complicated signalling methods such as low-swing, current-mode, booster- or accelerator-driven, to outperform rail-to-rail voltage-mode signalling,

without being prone to their usual vulnerability of noise susceptibility. A study that considered other forms of signalling for the boundary conditions of NoC architectures, would be potentially useful.

8. Bibliography

- [Acar99] E. Acar, A. Odabasioglu, M. Celik, and L. T. Pillage, "S2P: A stable 2 pole RC delay and coupling noise metric," in *Proc. GLSV-LSI*, 1999, pp. 60-63.
- [Adler98] V. Adler and E. B. Friedman, "Repeater Design to Reduce Delay and Power in Resistive Interconnect," *IEEE Trans. Circuits and Systems-II*, Vol. 45, No. 5, May 1998
- [Alpert01] C. J. Alpert, A. Devgan, and C. V. Kashyap, "RC Delay metrics for performance optimization," *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, vol. 20, no. 5, pp. 571-582, May 2001.
- [Alpert99] C. J. Alpert, A. Devgan, and S. T. Quay, "Buffer insertion for noise and delay optimization," *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, vol. 18, no. 11, pp. 1633-1645, Nov. 1999.
- [Alpert97] C. J. Alpert and A. Devgan, "Wire segmenting for improved buffer insertion," in *Proc. DAC*, 1997, pp. 588-593.
- [Anderson01] C. J. Anderson et.al., "Physical design of a fourth-generation POWER GHz microprocessor," in *Proc. ISSCC*, 2001, pp. 232-3.
- [BACPAC98] BACPAC: Berkeley Advanced Chip Performance Calculator. [Online]. Available: <http://www.eecs.umich.edu/~dennis/bacpac/>
- [Bakoglu90] H. B. Bakoglu, *Circuits, Interconnections, and Packaging for VLSI*, Reading, MA: Addison Wesley 1990.
- [Bakoglu87] H. B. Bakoglu and J. D. Meindl, "A system-level circuit model for multi- and single-chip CPUs," in *Proc. ISSCC*, pp. 308-9, 1987.

- [Bakoglu85] H. B. Bakoglu and J. D. Meindl, "Optimal Interconnection Circuits for VLSI," *IEEE Trans. Electron Devices*, vol. ED-32, no. 5, pp. 903-909, May 1985.
- [Banerjee01] K. Banerjee and A. Mehrotra, "Analysis of on-chip inductance effects using a novel performance optimization methodology for distributed RLC interconnects," in *Proc DAC*, Jun. 2001, pp. 798-803.
- [Bardeen48] J. Bardeen and W. Brattain, "The transistor, a semiconductor triode," *Phs. Rev.*, vol. 74, p.230, July 15, 1948.
- [Barke88] E. Barke, "Line-to-ground capacitance calculation for VLSI: a comparison," *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, vol. 7, no. 2, pp. 295-298, Feb. 1988.
- [Benini01] L. Benini and G. DeMicheli, "Powering Networks on Chip," in *Proc ISSS*, Oct. 2001, pp. 33-38.
- [Berlekamp68] E R Berlekamp, *Algebraic Coding Theory*, McGraw-Hill, 1968.
- [Bhavnagarwa00] A. J. Bhavnagarwala, A. Kapoor and J. D. Meindl, "Generic models for interconnect delay across arbitrary wire-tree networks," in *Proc. Interconnect Technology Conference*, 2000, pp. 129-131.
- [Celik02] M. Celik, L. Pileggi, and A. Odabasioglu, *IC Interconnect Analysis*, Kluwer Academic Publishers, May 2002.
- [Chan01] S. C. Chan and K. L. Shepard, "Practical considerations in RLCK crosstalk analysis for digital integrated circuits," in *Proc. ICCAD*, Nov. 2001, pp. 598-604.
- [Chen02] L. H. Chen and M. Marek-Sadowska, "Efficient closed-form cross-talk delay metrics," in *Proc. ISQED*, 2002, pp. 431-436.
- [Chen97] W. Chen, S. K. Gupta, and M. A. Breuer, "Analytic delay models for cross-talk delay and pulse analysis under non-ideal inputs," in *Proc. International Test Conference*, 1997, pp. 809-818.

-
- [Chern92] J. H. Chern, J. Huang, L. Arledge, P. C. Li, P. C. Lee, and P. Yang, "Multilevel metal capacitance models for CAD design synthesis systems," *IEEE Electron Device Lett.*, vol. 13, no. 1, pp. 32-34, Jan. 1992.
- [Chiprout98] E. Chiprout, "Interconnect and substrate modelling and analysis: an overview," *IEEE J. Solid-State Circuits*, vol. 33, no. 9, pp. 1445-1452, 1998.
- [Chiprout92] E. Chiprout and M. Nakhla, "Generalized moment-matching methods for transient analysis of interconnect networks," in *Proc. DAC*, 1992, pp. 201-206.
- [Chu87] C. Y. Chu and M. A. Horowitz, "Charge-sharing models for switch level simulation," *IEEE Trans. Computer-Aided Design*, vol. CAD-6, no. 6, pp. 1053-1061, Nov. 1987.
- [Cochrun73] B. L. Cochrun and A. Grabel, "On the determination of the transfer function of electronic circuits," *IEEE Trans. Circuit Theory*, vol. CT-20, pp.16-20, Jan. 1973.
- [Curran01] B. Curran et.al., "A 1.1 GHz first 64B generation z900 microprocessor," in *Proc. ISSCC*, 2001, pp. 238-9.
- [Dally01] W. J. Dally and B. Towles, "Route packets, not wires: on-chip interconnection networks," in *Proc. DAC*, Jun. 2001, pp. 684-689.
- [Dally98] W. J. Dally and J. W. Poulton, *Digital Systems Engineering*, New York, NY: CUP, 1998.
- [Dally97] W. J. Dally and J. W. Poulton, "Transmitter equalization for 4-Gbps signaling," *IEEE Micro*, vol 17, issue 1, pp. 48 -56, Jan.-Feb. 1997.
- [Dar91] S. Dar and M. A. Franklin, "Optimum Buffer Circuits for Driving Long Uniform Lines," *IEEE J. Solid-State Circuits*, vol. 26, pp. 32-40, Jan. 1991.

- [Davis00] J. A. Davis and J. D. Meindl, "Generic models for interconnect delay across arbitrary wire-tree networks," in *Proc. Interconnect Technology Conf.*, 2000, pp. 129 -131.
- [Deutsch01] A. Deutsch, P. W. Coteus, G. V. Kopcsay, H. H. Smith, C. W. Surovic, B. L. Krauter, D. C. Edelstein, and P. J. Restle, "On-chip wiring design challenges for gigahertz operation," in *Proc. IEEE Special Issue on Interconnections*, vol. 89, no. 4, pp. 529-555, April 2001.
- [Deutsch97] A. Deutsch, et al., "When are transmission lines important for on-chip interconnects," *IEEE Trans. Microwave Theory and Techniques*, vol. 45, no. 10, pp. 1836-1846, Oct. 1997.
- [Deutsch96] A. Deutsch, et al., "Design guidelines for short, medium and long on-chip interconnect," in *Proc. IEEE Topical Meeting on Electrical Performance of Electronic Packaging*, Oct. 1996, pp. 30-32.
- [Deutsch95a] A. Deutsch, et al., "Modelling and characterisation of long interconnects for high performance microprocessors," *IBM J. Research and Development*, vol. 39, no. 5, pp. 547-667, Sep. 1995.
- [Deutsch95b] A. Deutsch, A. Kopcsay, and G. V. Surovic, "Challenges raised by long on-chip wiring for CMOS microprocessors," in *Proc. IEEE Topical Meeting on Electrical Performance of Electronic Packaging*, Oct. 1995, pp. 21-23.
- [Deutsch90] A. Deutsch, et al., "High speed signal propagation on lossy transmission lines," *IBM J. Research and Development*, vol. 34, no. 4, pp. 601-615, Jul. 1990.
- [Devgan97] A. Devgan, "Efficient coupled noise estimation for on-chip interconnects," in *Proc. ICCAD*, 1997, pp. 147-153.
- [Eble96] J. C. Eble, V. K. De, D. S. Willis, and J. D. Meindl, "A generic system simulator (GENESYS) for ASIC technology and architecture beyond 2001," in *Proc. IEEE Intl. ASIC Conf.*, 1996, pp. 193-6.

-
- [Elmore48] W. C. Elmore, "The transient response of damped linear networks with particular regard to wideband amplifiers," *J. Appl. Physics*, vol. 19, pp-55-63, Jan. 1948.
- [Eo93] Y. Eo and W. R. Einstadt, "High speed VLSI interconnect modeling based on s parameter measurement," *IEEE Trans. on Components, Hybrids, and Manufacturing Technology*, vol. CHMT-16, no. 5, pp. 555-562, Aug. 1993
- [Feldmann95] P. Feldmann and R. W. Freund, "Efficient linear circuit analysis by Pade approximation via the Lanczos process," *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, vol. 14, no. 5, May 1995.
- [Franklin91] S. Dar and M. A. Franklin, "Optimum Buffer Circuits for Driving Long Uniform Lines," *IEEE J. Solid-State Circuits*, vol. 26, pp. 32-40, Jan. 1991.
- [Gajski99] D. Gajski, R. Dömer, and J. Zhu, "IP-centric methodology and design with the SpecC language," *System Level Design*, Nato Science Series, vol. 357, 1999.
- [Geuskens98] B. Geuskens and K. Rose, *Modeling Microprocessor Performance*, Kluwer, 1998.
- [Geuskens97] "Modeling the influence of multilevel interconnect on chip performance," Ph.D. dissertation, Rensselaer Polytechnic Institute, Troy, NY, 1997.
- [Ginneken90] L. P. P. P. van Ginneken, "Buffer placement in distributed RC-tree networks for minimal elmore delay," in *Proc. ISCAS*, 1990, pp. 865-68.
- [Gross98] P. D. Gross, R. Arunachalam, K. Rajagopal, and L. T. Pillegi, "Determination of worst-case aggressor alignment for delay calculation," in *Proc. ICCAD*, 1998, pp. 212-219.
- [Grover62] F. Grover, *Inductance Calculations: Working Formulas and Tables*, New York, NY: Dover, 1962.

- [GTX00] GTX: The GSRC Technology Extrapolation System. [Online]. Available: <http://public.itrs.net/Files/2001ITRS/Links/design/GTX/index.html>
- [Guerrier00] P. Guerrier and A. Greiner, "A generic architecture for on-chip packet switched networks," in *Proc. DATE*, 2000, pp. 250-6.
- [Hedenstierna93] N. Hedenstierna and K. O. Jeppson, "Comments on the optimum CMOS tapered buffer problem," *IEEE J. Solid-State Circuits*, vol. 29, no. 2, pp. 155-8, Feb. 1993.
- [Hedenstierna87] N. Hedenstierna and K. O. Jeppson, "CMOS circuit speed and buffer optimization," *IEEE Trans. Comp.-Aided Design*, vol. cad-6, no. 2, pp. 270-81, Mar. 1987.
- [Hemani00] A. Hemani, A. Jantsch, S. Kumar, A. Postula, J. Öberg, M. Millberg, and D. Lindqvist, "Network on Chip: An architecture for the billion transistor era," in *Proc. Norchip*, Nov. 2000, pp. 166-173.
- [Hemani99] A. Hemani, T. Meincke, A. Postula, T. Olsson, P. Nilsson, J. Öberg, P. Ellervee, and D. Lundqvist, "Lowering power consumption in clock by using globally asynchronous locally synchronous design style," in *Proc. DAC*, Jun. 1999, pp. 873-878.
- [Ho01] R. Ho, K. W. Mai, and M. A. Horowitz, "The future of wires," in *Proc. IEEE Special Issue on Interconnections*, vol. 89, no. 4, pp. 490-504, April 2001.
- [Horowitz84] M. A. Horowitz, "Timing models for MOS circuits," Ph.D. dissertation, Stanford University, Stanford Electronics Laboratory, Stanford, CA, Jan, 1984.
- [Hubing91] T. H. Hubing, "Survey of electromagnetic modeling techniques," Technical Report TR91-1-001.2, Electromagnetic Compatibility Lab, University of Missouri-Rolla, Rolla MO, 1991. Online. Available: <http://www.emclab.umn.edu/pdf/TR91-1-001.pdf>

-
- [Ismail01] Y. I. Ismail and E. G. Friedman, *On-chip Inductance in High Speed Integrated Circuits*, Massachusetts, Kluwer Academic Publishers, 2001.
- [Ismail00] Y. I. Ismail and E. G. Friedman, "Effects of inductance on the propagation delay and repeater insertion in VLSI circuits," *IEEE Trans. VLSI Systems*, vol. 8, pp. 195-206, Apr. 2000.
- [Ismail99] Y. I. Ismail, E. G. Friedman, and J. L. Neves, "Figures of merit to characterize the importance of on-chip inductance," *IEEE Trans. VLSI Systems*, vol. 7, pp. 442-449, Dec. 1999.
- [ITRS01] International technology semiconductor roadmap (ITRS) 2001. [Online]. Available: <http://public.itrs.net/files/2001ITRS/Home.htm>
- [Jaeger75] R. C. Jaeger, "Comments on an optimized output stage for MOS integrated circuits," *IEEE J. Solid-State Circuits*, vol. sc-10, no. 2, pp. 185-6, Jun. 1975.
- [Jarvis63] D. B. Jarvis, "The effects of interconnections on high-speed logic circuits," *IEEE Trans. Microwave Theory and Techniques*, vol. MTT-42, no. 8, pp. 476-487, Oct. 1963.
- [Jeppson96] K. O. Jeppson, "Comments on the metastable behaviour of CMOS latches," *IEEE J. Solid-State Circuits*, vol. 31, no. 2, pp. 275-7, Feb. 1996.
- [Kahng01] A. B. Kahng, S. Muddu, N. Pol, and D. Vidhani, "Noise model for multiple segmented coupled RC interconnects," in *Proc. ISQED*, 2001, pp. 145-150.
- [Kahng00] A. B. Kahng, S. Muddu, and E. Sarto, "On switch factor based analysis of coupled RC interconnects," in *Proc. DAC*, June 2000, pp. 79-84.
- [Kahng99] A. B. Kahng, S. Muddu, and D. Vidhani, "Noise and delay uncertainty studies for coupled RC interconnects," in *Proc. ASIC/SOC*, 1999, pp. 3-8.

- [Kahng97] A. B. Kahng and S. Muddu, "An analytic delay model for RLC interconnects," *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, vol. 16, no. 12, pp. 1507-1514, Dec. 1997.
- [Kahng95] A. B. Kahng and S. Muddu, "Two-pole analysis of interconnection trees," in *Proc. MCMC*, 1995, pp. 95-104.
- [Kamon94] S. Kamon, M. J. Tsuk, J.K. White, "FASTHENRY: a multipole-accelerated 3-D inductance extraction program," *IEEE Trans. Microwave Theory and Techniques*, vol. 42, no. 9, part 1-2, pp. 1750-1758, Oct. 1994.
- [Kang97] M. Z. W. Kang, W. W. M. Dai, T. Dillinger, and D. P. LaPotin, "Delay bounded buffer tree construction for timing driven floor-planning," in *Proc. ICCAD*, 1997, pp. 707-12.
- [Karush39] W. Karush, "Minima of functions of several variables with inequalities as side conditions," M.S. thesis, Department of Mathematics, University of Chicago, Chicago, 1939.
- [Kawaguchi98] H. Kawaguchi and T. Sakurai, "Delay and Noise Formulas for Capacitively Coupled Distributed RC Lines", in *Proc. Asian and South Pacific Design Automation Conference*, June 1998, pp. 35-43.
- [Kay98] R. Kay and L. T. Pillage, "PRIMO: Probability interpretation of moments for delay calculation," in *Proc. DAC*, 1998, pp. 463-468.
- [Krauter99] B. Krauter, S. Mehrotra, and V. Chandramouli, "Including inductive effects in interconnect timing analysis," in *Proc. CICC*, 1999, pp. 445-452.
- [Krauter98] B. Krauter and S. Mehrotra, "Layout based frequency dependent inductance and resistance extraction for on-chip interconnect timing analysis," in *Proc. DAC*, Jun. 1998, pp. 303-308.

-
- [Kuhn51] W. W. Kuhn and A. W. Tucker, "Nonlinear Programming," in *Proc. 2nd Berkeley Symposium on Mathematical Statistics and Probability*, 1951, pp. 481-492.
- [Landman71] B. S. Landman and R. L. Russo, "On a pin versus block relationship of logic graphs," *IEEE Trans. Computers*, vol. C-20, pp. 1469-79, Dec. 1971.
- [Lee98] T. H. Lee, *The Design of CMOS Radio Frequency Integrated Circuits*, New York, NY: CUP, 1998, pp. 114-131.
- [Lee97] M. Lee, "A multilevel parasitic interconnect capacitance modeling and extraction for reliable VLSI on-chip clock delay evaluation," *IEEE J. Solid-State Circuits*, vol. 33, no. 4, pp. 657-661, Apr. 1998.
- [Leiserson85] C. E. Leiserson, "Fat-Trees: Universal networks for hardware-efficient supercomputing," *IEEE Trans. Computers*, vol. C-34, no. 10, pp. 892-901, Oct. 1985.
- [Lewis84] E. T. Lewis, "An analysis of interconnect line capacitance and coupling for VLSI circuits," *Solid-State Electronics*, vol. 27, pp. 741-749, Aug. 1984.
- [Lillis96] J. Lillis, C. K. Cheng and T. T. Y. Lin, "Optimal wire sizing and buffer insertion for low power and a generalized delay model," *IEEE J. Solid-State Circuits*, vol 31, no. 3, Mar. 1996.
- [Lin00] S. Lin, N. Chang, and S. Nakagawa, "Quick on-chip self- and mutual-inductance screen," in *Proc. ISQED*, Mar. 2000, pp. 513-520.
- [Lin98] T. Lin, E. Acar, and L. T. Pillage, " h -gamma: An RC delay metric based on a gamma distribution approximation to the homogeneous response," in *Proc. ICCAD*, 1998, pp. 19-25.
- [Lin92] S. Lin and E. S. Kuh, "Transient simulation of lossy interconnect," in *Proc. DAC*, Jun. 1992, pp. 81-86.

- [Lin83] S Lin et al., “Error Control Coding: Fundamentals and Applications”, Prentice-Hall, 1983.
- [Lin75] H. C. Lin and L. W. Linholm, “An optimized output stage for MOS integrated circuits,” *IEEE J. Solid-State Circuits*, vol. sc-10, no. 2, pp. 106-9, Apr. 1975.
- [Liu03] J. Liu, D. Pamunuwa, L.-R. Zheng, and H. Tenhunen, “A global wire planning scheme for network-on-chip,” in *Proc. ISCAS*, May 2003, vol. 4, pp. 892-5.
- [Massoud98] Y. Massoud, S. Majors, and G. V. Surovic, “Layout techniques for minimising on-chip interconnect self inductance,” in *Proc. DAC*, Jun. 1998, pp. 566-571.
- [McCormick90] S. P. McCormick and J. Allen, “Waveform moment methods for improved interconnection analysis,” in *Proc. DAC*, June 1990, pp. 406-412.
- [Mead80] C. Mead and L. Conway, “Introduction to VLSI systems,” Reading MA: Addison-Wesley, 1980, Chapter 1.
- [Meijs84] N. v.d. Meijs and J. T. Fokkema, “VLSI circuit reconstruction from mask topology,” *Integration*, vol. 2, no. 2, pp. 85-119, 1984.
- [Menezes99] N. Menezes and C. P. Chen, “Spec-based repeater insertion and wire sizing for on-chip interconnect”, in *Proc. VLSI Design*, 1999, pp. 476-482.
- [Millberg02] M. Millberg, “The Nostrum protocol stack and suggested services provided by the Nostrum backbone”, Technical Report TRITA-IMIT-LECS R 02:01, Laboratory of Electronics and Computer Systems, Department of Micro-Electronics and Information Technology, Royal Institute of Technology, Stockholm, Sweden, 2003.
- [Moore65] G. E. Moore, “Cramming more components onto integrated circuits,” *Electronics*, vol 38, no. 8, pp. 114-117, 1965.

-
- [Nekili93] M. Nekili and Y. Savariya, "Parallel regeneration of interconnections in VLSI and ULSI circuits," in *Proc. ISCAS*, 1993, pp. 2023-2026.
- [Nekili92] M. Nekili and Y. Savariya, "Optimal methods of driving interconnections in VLSI circuits," in *Proc. ISCAS*, 1992, pp. 21-24.
- [Nemes84] M. Nemes, "Driving large capacitances in MOS LSI systems," *IEEE J. Solid-State Circuits*, vol. sc-19, no. 1, pp. 159-161, Feb. 1984.
- [Nilsson02] E. Nilsson, "Design and implementation of a hot-potato switch in a network on chip," MsC Thesis, Royal Institute of Technology, Department of Micro-Electronics and Information Technology, Laboratory of Electronics and Computer Systems, Stockholm, Sweden, Jun. 2002.
- [Odabasioglu98] A. Odabasioglu, M. Celik and L. T. Pillegi, "PRIMA: Passive reduced-order interconnect macromodeling algorithm," *IEEE Trans. Comp.-Aided Design of ICs and Sys.*, vol. 17, no. 8, pp. 645-654, Aug. 1998.
- [Pamunuwa03a] D. Pamunuwa, S. Elassaad and H. Tenhunen, "Modelling noise and delay in VLSI circuits," *Electronics Letters*, Vol. 39 Issue 3, pp. 269-271, Feb. 2003.
- [Pamunuwa03b] D. Pamunuwa and S. Elassaad, "Closed form metrics to accurately model the response in arbitrarily-coupled RC trees," in *Proc. IEEE International Symposium on Circuits and Systems (ISCAS 2003)*, Bangkok, Thailand, May 2003, vol. 4, pp. 892-5.
- [Pamunuwa03c] D. Pamunuwa, S. Elassaad and H. Tenhunen, "Analytic Modeling of Interconnects for Deep Sub-Micron Circuits", in *Proc. International Conference on Computer-Aided Design (ICCAD 2003)*, (in press), Nov. 2003.
- [Pamunuwa03d] D. Pamunuwa, J. Öberg, L. R. Zheng, M. Millberg, A Jantsch and H. Tenhunen, "Layout, performance and power trade-offs for in mesh-based network-on-chip architectures," in *Proc. IFIP Inter-*

national Conference on VLSI Systems-on-Chip, Darmstadt, Germany, Dec. 2003 (in press).

- [Pamunuwa03e] D. Pamunuwa, S. Elassaad and H. Tenhunen, "Modeling delay and noise in arbitrarily-coupled RC trees," *Submitted to IEEE Trans. Computer-Aided of ICs and Sys.*, Nov. 2003.
- [Pamunuwa02a] D. Pamunuwa, L. R. Zheng and H. Tenhunen, "Maximising Throughput over Parallel Wire Structures in the Deep Submicrometer Regime," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 11, no. 2, pp. 224-243, April, 2003.
- [Pamunuwa02b] D. Pamunuwa, L. R. Zheng and H. Tenhunen, "Optimising Bandwidth Over Deep Sub-micron Interconnect," in *Proc. ISCAS*, vol. 4, pp. 193-6, 2002.
- [Pamunuwa02c] D. Pamunuwa and H. Tenhunen, "On Dynamic Delay and Repeater Insertion," in *Proc. ISCAS*, vol. 1, pp. 97-100, 2002.
- [Pamunuwa02d] D. Pamunuwa and H. Tenhunen, "On Dynamic Delay and Repeater Insertion in Distributed Capacitively Coupled Interconnects", in *Proc. ISQED*, pp. 240-5, 2002.
- [Pamunuwa02e] "Signalling Techniques for DSM Designs," in *Workshop notes, European Solid-State Circuits Conference*, Florence, Italy, September, 2002.
- [Pamunuwa01] D. Pamunuwa and H. Tenhunen, "Repeater Insertion to Minimise Delay in Parallel Coupled Interconnects," in *Proc. VLSI Design*, pp. 513-17, 2001.
- [Pamunuwa00] D. Pamunuwa, L. R. Zheng and H. Tenhunen, "Combating Digital Noise in High Speed ULSI Circuits Using Binary BCH Encoding," in *Proc. ISCAS*, vol. 4, pp. 13-16, 2000.
- [Pamunuwa99] D. Pamunuwa, L. R. Zheng and H. Tenhunen, "Error-Control Coding to Combat Digital Noise in Interconnects for ULSI Circuits," in *Proc. Norchip*, pp. 275-82, 1999.

-
- [Paul94] C. R. Paul, *Analysis of Multi-conductor Transmission Lines*, New York, NY: John Wiley and Sons, 1994.
- [Paul92] C. R. Paul, *Introduction to Electromagnetic Compatibility*, New York, NY: John Wiley and Sons, 1994.
- [Pillage90] L. T. Pillage and R. A. Rohrer, "Asymptotic waveform evaluation for timing analysis," *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, vol. 9, pp. 352-366, Apr. 1990.
- [Priore93] D. A. Priore, "Inductance on silicon for sub-micron CMOS VLSI," in *Proc. IEEE Symp. on VLSI Circuits*, May 1993, pp. 17-18.
- [Rabaey96] J. M. Rabaey, *Digital Integrated Circuits*, Upper Saddle River, NJ: Prentice Hall, 1996.
- [Raghavan92] V. Raghavan, J. E. Bracken, and R. A. Rohrer, "AWESpice: a general tool for the accurate and efficient simulation of interconnect problems," in *Proc. DAC*, June 1992, pp. 740-745.
- [Ratzlaff94] C. L. Ratzlaff and L. T. Pillage, "RICE: Rapid interconnect circuit evaluation using AWE," *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, vol. 13, pp. 763-776, Jun. 1994.
- [RIPE97] RIPE: Rensselaer Interconnect Performance Estimator. [Online]. Available: <http://latte.cie.rpi.edu/ripe.html>
- [Rorabaugh96] C B Rorabaugh, "Error Coding Cookbook", McGraw-Hill, 1996.
- [Rosa08] E. B. Rosa, "The self and mutual inductance of linear conductors," *Bulletin of the National Bureau of Standards*, vol. 4, pp. 301-344, 1908
- [Roychowdhury91] J. S. Roychowdhury and D. O. Pederson, "Efficient transient simulation of lossy interconnect," in *Proc. DAC*, June 1991, pp. 406-412.

- [Rubinstein83] J. Rubinstein, P. Penfield, and M. Horowitz, "Signal delay in RC tree networks," *IEEE Trans. Computer Aided Design*, vol CAD-2, no. 3, pp. 202-211, July 1983.
- [Sai-Halasz95] G. A. Sai-Halasz, "Performance trends in high-end processors," in *Proc IEEE*, vol. 83, pp. 20-36, Jan. 1995.
- [Sakurai90] T. Sakurai and A. R. Newton, "Alpha-power law MOSFET model and its applications to CMOS inverter delay and other formulas," *IEEE J. Solid-State Circuits*, vol. 25, no. 2, pp. 584-594, Apr. 1990.
- [Sakurai83a] T. Sakurai and K. Tamaru, "Simple formulas for two- and three-dimensional capacitances," *IEEE Trans. Electron Devices*, vol. ED-30, no. 2, pp-183-5, Feb. 1983.
- [Sakurai83b] T. Sakurai, "Approximation of wiring delay in MOSFET LSI," *IEEE J. Solid-State Circuits*, vol. 18, no. 4, pp. 418-426, Aug. 1983.
- [Schaller97] R. E. Schaller, "Moore's law: past, present and future," *IEEE Spectrum*, vol. 34, no. 6, pp. 53-59, Jun. 1997.
- [Schockley49] W. Schockley, "The theory of pn junctions in semiconductors and pn-junction transistors," *Bell System Technical Journal (BSTJ)*, vol. 48, p. 435, 1949.
- [Schockley48] W. Schockley, "The transistor - a crystal diode," *Electronics*, pp. 68-71, Sep. 1948.
- [Sgroi01] M. Sgroi, M. Sheets, A. Mihal, K. Keutzer, S. Malik, J. Rabaey, and A. Sangiovanni-Vincentelli, "Addressing the system on-a-chip interconnect woes through communication based design," in *Proc. DAC*, 2001, pp. 667-672.
- [Shepard00] K. L. Shepard, D. Sitaram, and Y. Zheng, "Full-chip, three-dimensional, shapes-based RLC extraction," in *Proc. ICCAD*, Nov. 2000, pp. 142-149.

-
- [Shepard97] K. L. Shepard, V. Narayanan, P. C. Elmendorf, and G. Zheng, "Global Harmony: coupled noise analysis for full-chip RC interconnect networks," in *Proc. ICCAD*, Nov. 1997, pp. 139-146.
- [Shin01] Y. Shin and T. Sakurai, "Coupling-driven bus design for low-power application-specific systems," in *Proc. DAC*, Jun. 2001, pp. 750-753.
- [Shoji96] M. Shoji, *High Speed Digital Circuits*, Addison Wesley, Massachusetts, 1996.
- [Silveira95] L. M. Silveira, M. Kamon and J. White, "Efficient reduced-order modeling of frequency-dependent coupling inductances associated with 3-D interconnect structures," in *Proc. DAC*, 1995, pp. 376-80.
- [Sirichotiyakul01] S. Sirichotiyakul, D. Blaauw, C. Oh, R. Levy, V. Zolotov, and J. Zuo, "Driver modeling and alignment for worst-case delay noise," in *Proc. DAC*, June 2001, pp. 720-725.
- [Stallings94] W. Stallings, *Data and Computer Communications*, Fourth Edition, Prentice-Hall, 1994.
- [Sylvester99a] D. Sylvester and K. Keutzer, "Getting to the bottom of deep sub-micron II: a global wiring paradigm," in *Proc. ISPD*, 1999, pp. 193-200.
- [Sylvester99b] D. Sylvester and K. Keutzer, "System level performance modeling with BACPAC - Berkeley Advanced Chip Performance Calculator," in *Workshop notes, International Workshop on System Level Interconnect Prediction*, 1999, pp. 109-114.
- [Sylvester98] D. Sylvester and K. Keutzer, "Getting to the bottom of deep sub-micron," in *Proc. ICCAD*, 1998, pp. 203-211.
- [Szyperski98] *Component Software: Beyond Object Oriented Software*, Reading, MA: ACM/Addison Wesley, 1998

- [Takahashi01] M. Takahashi, M. Hashimoto, and H. Onodera, "Crosstalk noise estimation for generic RC trees", in *Proc. ICCD*, 2001, pp. 110-116.
- [Tong00] X. Tong and M. Marek-Sadowska, "Efficient delay calculation in presence of crosstalk," in *Proc. ISQED*, 2000, pp. 491-497.
- [Tutuianu96] B. Tutuianu, F. Dartu and L. T. Pillage, "An explicit RC-circuit delay approximation based on the first three moments of the impulse response," in *Proc. DAC*, 1996, pp. 611-616.
- [Vittal99] A. Vittal, L. H. Chen, M. Marek-Sadowska, K. Wang, and S. Yang, "Crosstalk in VLSI Interconnections," *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, vol. 18, no. 12, pp. 1817-1824, Dec. 1999.
- [Wu90a] C. Y. Wu and M. Shiau, "Delay models and speed improvement techniques for RC tree interconnections among small-geometry CMOS inverters," in *IEEE J. Solid-State Circuits*, vol. 25, no. 5, pp. 1247-1256, Oct. 1990.
- [Wu90b] C. Y. Wu and M. Shiau, "Accurate speed improvement techniques for RC line and tree interconnections in CMOS VLSI," in *Proc. ISCAS 1990*, pp. 2.1648-2.1651.
- [Yu99] Q. Yu and E. Kuh, "Passive multipoint moment matching model order reduction algorithm on multiport distributed interconnect networks," *IEEE Trans. Circuits and Systems-I*, Vol. 46, pp. 140-160, Jan. 1999.
- [Yuan82] C. P. Yuan and T. N. Trick, "A simple formula for the estimation of the capacitance of two-dimensional interconnects in VLSI circuits," *IEEE Electron Device Lett.*, vol. EDL-3, pp. 391-393, 1982.
- [Zheng01] L. R. Zheng, "Design, analysis and integration of mixed signal systems for signal and power integrity," Ph.D. dissertation, Royal Institute of Technology, Stockholm, May, 2001.

- [Zheng00] L. R. Zheng, D. Pamunuwa and H. Tenhunen, “Accurate *A Priori* Signal Integrity Estimation Using a Dynamic Interconnect Model for Deep Submicron VLSI Design”, in *Proc. ESSCIRC*, sep. 2000, pp. 324-327.
- [Öberg02] J. Öberg, D. Pamunuwa, L. R. Zheng, M. Millberg, A Jantsch and H. Tenhunen, “A feasibility study on the performance and power distribution of two possible Network-on-Chip architectures,” under review, *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, Oct. 2002.

