

TURUN YLIOPISTON JULKAISUJA
ANNALES UNIVERSITATIS TURKUENSIS

SARJA - SER. A I OSA - TOM. 410

ASTRONOMICA - CHEMICA - PHYSICA - MATHEMATICA

Exploration and Design of High Performance Variation Tolerant On-Chip Interconnects

by

Ethiopia Enideg Nigussie

TURUN YLIOPISTO
UNIVERSITY OF TURKU
Turku 2010

Supervisors

Adjunct Professor Juha Plosila
Professor Jouni Isoaho
Professor Hannu Tenhunen
Department of Information Technology
University of Turku
Turku, Finland

Reviewers

Ph.D. Dinesh Pamunuwa
Center for Microsystems Engineering
Department of Engineering
Lancaster University
Lancaster, UK

Professor Olli Vainio
Department of Computer Systems
Tampere University of Technology
Tampere, Finland

Opponent

Professor Mohammed Ismail El-Naggar
Department of Electrical and Computer Engineering
The Ohio State University
2015 Neil Avenue, Columbus, OH 43210
USA

ISBN 978-951-29-4363-0 (PRINT)

ISBN 978-951-29-4364-7 (PDF)

ISSN 0082-7002

Painosalama Oy – Turku, Finland 2010

Abstract

Continuous technology scaling enables implementation of complex application on a single chip. As a result, there is a design paradigm shift from single-core to multicore systems and from core-centric to interconnect-centric designs, emphasizing the importance of high performance, power efficient, and reliable on-chip interconnects. In sub-90nm technologies, parameter variations are increasing and cause considerable delay variations, which creates difficulty in guaranteeing the reliability of on-chip interconnects.

In this thesis, high performance, variation tolerant and power efficient global on-chip interconnects are presented. The design and implementation of these interconnects are based on formulation and integration of different circuit level techniques. Since delay variations are inevitable, the thesis focuses on self-timed delay-insensitive communication. In this regard, design and optimization of delay-insensitive data encoding/decoding schemes as well as formulation of efficient communication protocols are performed. To compensate the delay overhead of delay-insensitive communication, high speed signaling techniques are developed and implemented. In addition, a novel high speed completion detection technique is devised and implemented to solve the performance bottleneck caused by conventional completion detection methods. A high-throughput and power efficient serial interconnect is also designed to be used as a long-range on-chip communication link.

Furthermore, an interconnect calibration technique after every power start-up of a system is developed and implemented to ensure signal integrity of the interconnects despite process, wearout and aging caused variations. A runtime supply voltage and temperature variation tolerance technique is also devised and implemented for the interconnects. These PVT variation tolerance schemes make the interconnects adaptive to the effect of variations, enabling continuous and reliable operation of the interconnect.

Acknowledgements

It gives me great pleasure to be able to express my gratitude to the people and institutions that have helped me to accomplish this research work. First and foremost, I would like to thank my supervisors Adj. Prof. Juha Plosila, Prof. Jouni Isoaho, and Prof. Hannu Tenhunen for their inspiration, guidance and support. I owe a huge debt of thanks to Adj. Prof. Juha Plosila for his invaluable and continuous guidance. His comments and criticism had a significant impact on the research presented here. I am very grateful to Prof. Jouni Isoaho for his guidance and encouragement throughout the period of my research. I am also indebted to Prof. Hannu Tenhunen for inspiring me to pursue my research in on-chip interconnects and backed me up in the final phase of the thesis. I wish to thank Prof. Olli Vainio and Ph.D. Dinesh Pamunuwa for reviewing the thesis.

The Graduate School in Electronics, Telecommunications and Automation (GETA) is gratefully acknowledged for funding my doctoral studies. This research work was financially supported by the Nokia Foundation, the Ulla Tuominen Foundation, and the Otto A. Malm Foundation.

I would like to acknowledge all my colleagues at the IT department. I am grateful to everyone who have co-authored papers with me. Special thanks go to Sampo Tuuna for co-authoring papers with me which are included in this thesis. I am thankful to D.Sc. (Tech) Johanna Tuominen and Kameswar Rao Vaddina for proof-reading part of this thesis. D.Sc. (Tech) Teijo Lehtonen deserves special thanks, he always took time to discuss any issue I brought to him and came up with practical suggestions each time.

I would like to express my deepest gratitude to my sister Konjit for her constant support, encouragement and most importantly taking care of my responsibilities at home. I am also indebted to my grand mother Abeba and my brothers Girma and Yehwalashete, who always encouraged me to follow

my ambitions, irrespective of how much they missed me. I am grateful to my finnish family Maija, Pekka and Iida who make my stay in Turku much more pleasant.

My final and most heartfelt thanks go to my husband Woubishet for his continuous support, understanding and companionship throughout my doctoral studies.

Turku, August 2010

Ethiopia Enideg Nigussie

Contents

Abstract	ii
Acknowledgments	ii
List of Publications	viii
List of Figures	xii
List of Tables	xvi
List of Abbreviations	xvii
1 Introduction	1
1.1 Emergence of Interconnect-Centric Design	1
1.1.1 Device and Interconnect Scaling	2
1.1.2 System-on-Chip and Multicore Systems	3
1.1.3 Network-on-Chip	4
1.2 Challenges of Global On-Chip Interconnect	5
1.2.1 Performance and Power Consumption	6
1.2.2 Variability and Reliability	7
1.3 Global On-Chip Communication Techniques	8
1.3.1 GALS Communication	9
1.3.2 Self-timed Delay-Insensitive Communication	10
1.4 Scope of Thesis and Contributions	10
1.4.1 Delay-Insensitive Current Sensing Interconnects	11
1.4.2 Completion Detection Technique to Enhance Performance	12
1.4.3 High Throughput Energy Efficient Semi-Serial Link	13
1.4.4 Circuit Techniques for PVT Variation Tolerance	14
1.5 Related Work	15

1.6	Thesis Organization	18
2	Interconnect Design Techniques	19
2.1	Handshaking Protocols	19
2.2	Data Encoding Techniques	22
2.3	Data Decoding Techniques	25
2.4	Completion Detection Techniques	26
2.5	Self-timed Components	27
2.6	On-Chip Signaling Schemes	28
2.6.1	Current-Mode and Current Sensing Signaling	29
2.6.2	Voltage-Mode Signaling: Reference	33
2.7	On-Chip Wire Modeling	33
2.7.1	Wire Parasitic Estimation and Extraction	34
2.7.2	Electrical Level Wire Modeling	39
2.8	Chapter Summary	42
3	Design of Delay-Insensitive Current Sensing Interconnects	43
3.1	Level-Encoded Dual-Rail Current Sensing Interconnect	44
3.1.1	Data Encoder and Driver	46
3.1.2	Receiver, Decoder and Completion Detector	47
3.1.3	Acknowledgment Transmission	48
3.1.4	Simulation Results and Analysis	49
3.1.5	Effect of Crosstalk on Timing	52
3.2	1-of-4 Encoded Current Sensing Interconnect	54
3.2.1	Encoder and Driver	56
3.2.2	Receiver	57
3.2.3	Decoder and Completion Detector	59
3.2.4	Acknowledgment Transmission	61
3.2.5	Reference Voltage-Mode Interconnects	62
3.2.6	Simulation Results and Analysis	63
3.3	Dual-Rail Encoded Differential Current Sensing Interconnect	69
3.3.1	Encoding and Its Implementation	71
3.3.2	Driver, Receiver and Completion Detector	72
3.3.3	Acknowledgment Transmission	75
3.3.4	Simulation Results and Analysis	75
3.4	Chapter Summary	79

4	Enhancing Completion Detection Performance	81
4.1	Delay-Insensitive Bit Parallel Transmission	82
4.2	High-Speed Completion Detection Technique	85
4.3	Case Studies	87
4.3.1	1-of-4 Encoded Current Sensing Interconnect	88
4.3.2	Dual-rail Encoded Differential Current Sensing Interconnect	90
4.3.3	Acknowledgment Transmission	90
4.4	Reference Cases	92
4.5	Simulation Results and Analysis	93
4.5.1	Wire Model	93
4.5.2	Simulations Setup	93
4.5.3	Performance Analysis	94
4.5.4	Power Analysis	97
4.5.5	Area Comparison	99
4.6	Chapter Summary	100
5	Energy Efficient Semi-Serial Interconnect	101
5.1	Long-Range Link in NoC	103
5.2	High-Throughput Serial On-Chip Interconnect	106
5.2.1	Communication Protocol	106
5.2.2	Serializer and Pulse Dual-Rail Encoding	108
5.2.3	High-Speed Differential Pulse Current-Mode Signaling	110
5.2.4	Deserializer	114
5.2.5	Acknowledgment Transmission	115
5.3	Simulation Results and Analysis	116
5.3.1	Wire Model and Simulation Waveforms	116
5.3.2	Performance	117
5.3.3	Power and Energy Consumption	117
5.4	Fully Bit-Parallel vs Serial Links	120
5.5	Chapter Summary	123
6	Comparison of the Designed Interconnects	124
6.1	Summary of the Interconnects	124
6.2	Comparison of the Interconnects	126
6.2.1	Performance	126
6.2.2	Power Efficiency	127

6.2.3	Area	130
6.3	Chapter Summary	130
7	Circuit Techniques for PVT Variation Tolerance	131
7.1	Signal Integrity of Current Sensing Interconnect	132
7.1.1	Effects of Process Variation	132
7.1.2	Runtime Supply Voltage and Temperature Variations .	139
7.2	Post-Manufacture Variation Adaptation	141
7.3	Calibration for Process Variation Tolerance	142
7.3.1	Algorithm and Methodology	142
7.3.2	Reconfiguration Control and Communication Circuits .	148
7.4	Runtime Management of Voltage and Temperature Variations .	152
7.4.1	Sensing Effects of Voltage and Temperature Variation .	153
7.4.2	Sensor Circuit Implementation	155
7.4.3	Reconfiguration and Retransmission	156
7.5	Simulation Results and Analysis	157
7.6	Chapter Summary	163
8	Conclusions	164
8.1	Future Work	166
	References	167

List of Publications

The work discussed in this thesis is based on and extended from the publications listed below:

1. E. Nigussie, J. Plosila and J. Isoaho. *Monitoring and Reconfiguration Techniques for Power Supply Variation Tolerant on-Chip Links*. 2010 IEEE International Symposium on Circuits and Systems (ISCAS 2010), 4 pages, May 2010, Paris, France.
2. E. Nigussie, J. Plosila and J. Isoaho. *Process Variation Tolerant On-Chip Communication Using Receiver and Driver Reconfiguration*. 11th IEEE International Symposium on Quality Electronic Design (ISQED 2010), pages 453-460, Mar 2010, San Jose, CA, USA.
3. E. Nigussie, J. Plosila, S. Tuuna, J. Isoaho and H. Tenhunen. *Energy Efficient Semi-Serial On-Chip Link Through Circuit Optimizations and Integration of Signaling Techniques*. IEEE Transactions on Very Large Scale Integration (VLSI) Systems, under review.
4. E. Nigussie, J. Plosila, S. Tuuna, J. Isoaho and H. Tenhunen. *Boosting Performance of Self-Timed Delay-Insensitive Bit Parallel On-Chip Interconnects*. IET Circuits, Devices and Systems Journal, under review.
5. E. Nigussie, J. Plosila and J. Isoaho. *High-speed Completion Detection for Current Sensing On-Chip Interconnects*. In Electronics Letters, 45(11), pages 547-548, May 2009.
6. E. Nigussie, J. Plosila and J. Isoaho. *Area Efficient Delay-Insensitive and Differential Current Sensing On-Chip Interconnect*. In 21st IEEE International SoC Conference (SOCC 2008), Pages 143-146, Sept. 2008, Newport Beach, USA.

7. E. Nigussie, J. Plosila and J. Isoaho. *Current-Mode On-Chip Interconnect using Level-Encoded Two-Phase Dual-Rail Encoding*. In 2007 IEEE International Symposium on Circuits and Systems (ISCAS 2007), pages 649-652, May 2007, New Orleans, USA.
8. E. Nigussie, T. Lehtonen, S. Tuuna, J. Plosila and J. Isoaho. *High-Performance Long NoC Link Using Delay-Insensitive Current-Mode Signaling*. In VLSI Design (Hindawi), vol. 2007, Article ID 46514, 13 pages, 2007.

The following research papers were also authored/co-authored during the course of this PhD project:

1. E. Nigussie, S. Tuuna, J. Plosila, and J. Isoaho. *Analysis of Crosstalk and Process Variations Effects on On-Chip Interconnects*. In IEEE 2006 International Symposium on System-on-Chip, Nov. 2006, Tampere, Finland.
2. E. Nigussie, J. Plosila and J. Isoaho. *Full-Duplex Link Implementation using Dual-Rail Encoding and Multiple-Valued Current-Mode Logic*. In IEEE International Symposium on Circuits and Systems 2006 (ISCAS 2006), May 2006, Kos, Greece.
3. E. Nigussie, J. Plosila and J. Isoaho. *Delay-Insensitive On-Chip Communication Link using Low-Swing Simultaneous Bidirectional Signaling*. In IEEE Computer Society Annual Symposium on VLSI 2006 (ISVLSI 2006), March 2006, Karlsruhe, Germany.
4. E. Nigussie, J. Plosila and J. Isoaho. *On Asynchronous Full-Duplex Dual-Rail Link with Multiple-Valued Current-Mode Signaling*. In IEEE Circuits and Systems Society 23rd NORCHIP Conference, Nov. 2005, Oulu, Finland.
5. E. Nigussie, J. Plosila and J. Isoaho. *Reliable Asynchronous Links for SoC*. In IEEE 2005 International Symposium on System-on-Chip, Nov. 2005, Tampere, Finland.
6. L. Guang, E. Nigussie, J. Isoaho, P. Rantala and H. Tenhunen. *Interconnection Alternatives for Hierarchical Monitoring Communication in Parallel System-on-Chip*. In Microprocessors and Microsystems:

- Embedded Hardware Design Journal, Elsevier, Vol. 34, No. 5, Aug. 2010.
7. L. Guang, E. Nigussie, P. Rantala, J. Isoaho and H. Tenhunen. *Hierarchical Agent Monitoring Design Approach towards Self-Aware Systems*. In ACM Transactions in Embedded Computing Systems (TECS), 9(3), pages 25:1-25:24, Feb. 2010.
 8. L. Guang, E. Nigussie and H. Tenhunen. *Run-time Communication Bypassing for Energy-Efficient, Low-Latency Per-Core DVFS on Network-on-Chip*. To appear in 23rd IEEE International SoC Conference (SOCC 2010), Nevada, USA.
 9. L. Guang, E. Nigussie and H. Tenhunen. *System-Level Exploration of Run-Time Clusterization for Energy-Efficient On-Chip Communication*. In 2nd international workshop on Network-on-chip architectures, held in conjunction with 42nd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO-42), Dec 2009, New York, USA.
 10. L. Guang, E. Nigussie, L. Koskinen, and H. Tenhunen. *Autonomous DVFS on Supply Islands for Energy-constrained NoC Communication*. In International Conference on Architecture of Computing Systems 2009, Springer Lecture Notes in Computer Science (LNCS), 5545, Mar 2009.
 11. K. R. Vaddina, E. Nigussie, P. Liljeberg and J. Plosila. *Self-Timed Thermal Sensing and Monitoring of Multicore Systems*. In 12th IEEE Symposium on Design and Diagnostics of Electronic Systems, April 2009, Czech.
 12. S. Tuuna, E. Nigussie, J. Isoaho, and H. Tenhunen. *Analysis of Delay Variation in Encoded On-Chip Bus Signaling under Process Variation*. In 21st International Conference on VLSI Design, Jan. 2008, Hyderabad, India.
 13. L. Guang, P. Liljeberg, E. Nigussie and H. Tenhunen. *A Review of Dynamic Power Management Methods in NoC Under Emerging Design Considerations*. In IEEE Norchip 2009, Nov 2009, Trondheim, Norway.
 14. W. Yin, L. Guang, E. Nigussie, P. Liljeberg, J. Isoaho and H. Tenhunen. *Architecture Exploration of Per-Core DVFS for Energy-Constrained On-*

- Chip Networks.* In 12th Euromicro Conference on Digital System Design (DSD 2009), Aug 2009.
15. K. R. Vaddina, L. Guang, E. Nigussie, P. Liljeberg and J. Plosila. *Online Distributed Thermal Sensing and Monitoring of Multicore Systems.* In IEEE 26th NORCHIP Conference, Nov. 2008, Tallinn, Estonia.
 16. L. Guang, P. Rantala, E. Nigussie, J. Isoaho, and H. Tenhunen. *Low-latency and Energy-efficient Monitoring Interconnect for Hierarchical-agent-monitored NoCs.* In IEEE 26th NORCHIP Conference, Nov. 2008, Tallinn, Estonia.
 17. L. Guang, A. W. Yin, P. Rantala, E. Nigussie, P. Liljeberg, J. Isoaho and H. Tenhunen. *Hierarchical Power Monitoring for On-chip Networks.* In Proceedings of Work in Progress Session in Euromicro PDP 2009 Conference, Feb 2009.
 18. A. W. Yin, L. Guang, P. Liljeberg, P. Rantala, E. Nigussie, J. Isoaho and H. Tenhunen. *Hierarchical Agent Based NoC with Dynamic Online Services.* In The 4th IEEE Conference on Industrial Electronics and Applications (ICIEA 2009), May 2009.
 19. A. W. Yin, L. Guang, P. Liljeberg, P. Rantala, E. Nigussie, J. Isoaho and H. Tenhunen. *Hierarchical Agent Architecture for Scalable NoC Design with Online Monitoring Services.* 1st International Workshop on Network on Chip Architectures (NoCArc) held in conjunction with 41st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO-41), Nov. 2008.

List of Figures

1.1	Intel’s 45nm 8-core Xeon-EX processor [5].	4
1.2	Delay comparison [105]	7
2.1	Push channel.	20
2.2	Pull channel.	20
2.3	Four-phase handshaking protocol.	21
2.4	Two-phase handshaking protocol.	22
2.5	Delay-insensitive push channel.	23
2.6	Four-phase dual-rail encoded transmission	24
2.7	Two-phase dual-rail encoded transmission	24
2.8	Four-phase 1-of-4 encoded transmission	25
2.9	Two-phase 1-of-4 encoded transmission	25
2.10	Transistor level implementation of a 2-input C-element	28
2.11	Active-low resettable C-element	28
2.12	3-input upper asymmetric C-element	29
2.13	Single microstrip wire	35
2.14	Distributed RLC wire model with coupling	42
3.1	Bundled-data \Leftrightarrow delay-insensitive conversion	44
3.2	Conventional two-phase dual-rail encoder.	45
3.3	Conventional two-phase dual-rail decoder and completion detector.	45
3.4	Transitions between LEDR code.	46
3.5	Protocol conversion.	47
3.6	LEDR encoded current sensing on-chip interconnect.	48
3.7	Forward latency of <i>LEDRC_m</i> , <i>LEDRV_m</i> and <i>TPDRV_m</i>	50
3.8	Backward latency of <i>LEDRC_m</i> , <i>LEDRV_m</i> and <i>TPDRV_m</i>	50
3.9	Throughput of <i>LEDRC_m</i> , <i>LEDRV_m</i> and <i>TPDRV_m</i>	51

3.10	Power consumption of <i>LEDR_{Cm}</i> , <i>LEDR_{Vm}</i> and <i>TPDR_{Vm}</i>	52
3.11	Energy per bit dissipation of <i>LEDR_{Cm}</i> , <i>LEDR_{Vm}</i> and <i>TPDR_{Vm}</i>	52
3.12	Simulation waveforms of <i>LEDR_{Cm}</i> interconnect	53
3.13	Communication protocol of <i>PMC_m</i>	56
3.14	Encoder and driver of <i>PMC_m</i>	58
3.15	Decoder and completion detector of <i>PMC_m</i>	60
3.16	Acknowledgment transmission of <i>PMC_m</i>	61
3.17	Encoder of <i>TPV_m</i>	63
3.18	Pipeline stage, acknowledgment driver and receiver of <i>TPV_{mP}</i>	64
3.19	Decoder and completion detector of <i>TPV_m</i>	64
3.20	Forward latency of 1-of-4 encoded interconnects.	65
3.21	Throughput of 1-of-4 encoded interconnects.	66
3.22	Power consumption of 1-of-4 encoded interconnects.	67
3.23	Energy per bit dissipation of 1-of-4 encoded interconnects.	68
3.24	Crosstalk effect in latency of <i>PMC_m</i> and <i>TPV_{mP}</i>	69
3.25	Crosstalk effect in throughput of <i>PMC_m</i> and <i>TPV_{mP}</i>	69
3.26	Simulation waveforms of <i>PMC_m</i>	70
3.27	Communication protocol of <i>Dualdiff</i> interconnect	72
3.28	Encoder of <i>Dualdiff</i> interconnect.	73
3.29	Driver of <i>Dualdiff</i> interconnect.	73
3.30	Receiver and completion detector of <i>Dualdiff</i> interconnect	75
3.31	Acknowledgment transmission of <i>Dualdiff</i> interconnect	76
3.32	Simulation waveforms of <i>Dualdiff</i>	77
3.33	LEDR encoded differential interconnect	78
3.34	Forward latency of <i>Dualdiff</i> and <i>LEDR_{diff}</i>	79
3.35	Throughput of <i>Dualdiff</i> and <i>LEDR_{diff}</i>	79
3.36	Power consumption of <i>Dualdiff</i> and <i>LEDR_{diff}</i>	80
3.37	Energy per bit dissipation of <i>Dualdiff</i> and <i>LEDR_{diff}</i>	80
4.1	Completion detector of 32-bit Two-Phase 1-of-4 Transmission	83
4.2	Completion detector of 32-bit Two-Phase Dual-Rail Transmission	84
4.3	Completion detection in a pipelined voltage-mode link	84
4.4	High-speed completion detection circuit.	87
4.5	Bit-width versus SNR of a wire for a reliable detection.	87
4.6	Completion detection circuit of 4-bit <i>PMC_mFCD</i> link.	89
4.7	Completion detection circuit of 4-bit <i>DualdiffFCD</i> link.	91

4.8	Acknowledgment signal transmission.	92
4.9	Latency of links with high-speed completion detection.	95
4.10	Throughput of links with high-speed completion detection.	95
4.11	Throughput of two-phase 1-of-4 encoded 2mm interconnects.	96
4.12	Throughput of two-phase dual-rail encoded 2mm interconnects.	97
4.13	Power consumption of 1-of-4 encoded 2mm interconnects.	98
4.14	Power consumption of dual-rail 2mm encoded interconnects.	98
4.15	Energy per bit dissipation of 1-of-4 encoded 2mm Links	99
4.16	Energy per bit dissipation of dual-Rail encoded 2mm Links	99
5.1	64 nodes flattened butterfly topology [67]	104
5.2	64 nodes concentrated mesh with express channel [72]	105
5.3	16 nodes NoC topologies requiring long-range links	105
5.4	High-throughput serial on-chip communication link.	106
5.5	Serializer communication protocol.	108
5.6	Deserializer communication protocol.	109
5.7	Serializer and pulse dual-rail encoder.	110
5.8	Shift register's TSPC flip-flop with parallel data loading.	111
5.9	Counter's resettable TSPC flip-flop.	112
5.10	Pulse dual-rail encoder input and output signals.	112
5.11	Driver, receiver and data validity decoder of serial link.	113
5.12	Deserializer circuit.	115
5.13	Double edge-triggered TSPC flip-flop.	115
5.14	Pulse generator for acknowledgment signal transmission.	116
5.15	Simulation waveforms of serial link.	118
5.16	Energy per bit of 64-bits word bit-serial and semi-serial links.	120
5.17	Throughput of 64-bit word serial and parallel links.	122
5.18	Energy per bit of 64-bit word serial and parallel links.	122
6.1	Throughput versus bit width of links	128
6.2	Energy per bit versus communication distance of interconnects	129
6.3	Energy per bit versus transmission bit width of links	129
7.1	Variable parameters of current sensing interconnect	132
7.2	<i>LEDRCm</i> interconnect I_{win} and $I_{rec,in}$ variations	134
7.3	<i>PMCMFCD</i> interconnect I_{win} and $I_{rec,in}$ variations	136
7.4	<i>DualdiffFCD</i> interconnect I_{win} and $I_{rec,in}$ variations	137

7.5	Interconnect model for analysis	138
7.6	$I_{rec,in}$ versus supply voltage and temperature	141
7.7	I_{ref} versus supply voltage and temperature	141
7.8	Interconnect calibration flow chart	143
7.9	Interconnect with calibration wires	144
7.10	Reconfigurable driver and receiver	145
7.11	Best case : driver reconfiguration is not required	148
7.12	Average case calibrations	149
7.13	Worst case calibrations	150
7.14	Calibration failure cases	151
7.15	Calibration control and communication signals	151
7.16	Receiver reconfiguration control circuit	153
7.17	Driver side calibration control	154
7.18	Transition-to-pulse and pulse-to-transition converters	154
7.19	Interconnect with VT runtime variations detector	155
7.20	VT variation sensor and reconfiguration circuit	156
7.21	Simulation waveforms of average-case calibration	160
7.22	Simulation waveforms of VT variation tolerance	162

List of Tables

2.1	The truth table of 2-input C-element.	27
3.1	Comparing <i>LEDRCm</i> and <i>LEDRVm</i> interconnects.	53
3.2	Comparing <i>LEDRCm</i> and <i>TPDRVm</i> Interconnects.	53
3.3	Effect of crosstalk in <i>LEDRCm</i> and <i>BundledVm</i>	54
3.4	Encoding and wire current of <i>PMCM</i>	56
3.5	Encoding protocol of <i>Dualdiff</i> interconnect.	72
3.6	Area comparison between <i>Dualdiff</i> and <i>LED Rdiff</i>	78
4.1	Area comparison of interconnects.	100
5.1	Power and energy dissipations of bit-serial communication. . .	119
5.2	Performance and energy dissipation of semi-serial links. . . .	120
5.3	Area comparison between serial and parallel links.	122
6.1	Summary of interconnects	125
6.2	Throughput of interconnects.	127
6.3	Throughput of 64-bit word 5mm long links.	128
6.4	Area comparison of 64-bit 2mm long transmission links.	130
7.1	Calibration delay and power consumption.	158
7.2	Calibration area overhead.	159
7.3	Power overhead of VT variation management.	161
7.4	VT variation management area overhead.	161

List of Abbreviations

AC	Alternating Current
BundledVm	Bundled-data encoded Voltage-mode interconnect
CMOS	Complementary Metal Oxide Semiconductor
Dualdiff	Two-phase Dual-rail encoded differential current sensing interconnect
DualdiffFCD	Dualdiff interconnect using fast completion detection
DualVmP	Two-phase dual-rail encoded pipelined voltage-mode interconnect
DVD	Dynamic Voltage Drop
GALS	Globally Asynchronous Locally Synchronous
IC	Integrated Circuit
IP	Intellectual Property
ITRS	International Technology Roadmap for Semiconductors
LEDR	Two-phase Level-Encoded Dual-Rail
LEDRcm	Level-Encoded Dual-Rail Current Sensing interconnect
LEDRdiff	LEDR Encoded Differential Current Sensing interconnect
LEDRVm	Level-Encoded Dual-Rail Voltage-Mode interconnect
NoC	Network-on-Chip

PMCM	Two-phase 1-of-4 encoded multilevel current sensing interconnect
PMCMFCD	PMCM interconnect using fast completion detection
PVT	Process, Voltage and Temperature
SoC	System-on-Chip
TPDRVm	Two-Phase Dual-Rail Voltage-Mode interconnect
TPVm	Two-Phase 1-of-4 encoded Voltage-mode interconnect
TPVmP	Two-Phase 1-of-4 encoded Voltage-mode Pipelined interconnect
TPVmRep	Two-Phase 1-of-4 encoded Voltage-mode Repeated interconnect
TSPC	True Single Phased Clocked
VT	Voltage and Temperature

Chapter 1

Introduction

The continuous development of semiconductor technology over the last five decades has been the enabling factor that has driven many huge changes in our everyday life. Personal computing, mobile communications, Internet, broadband technology and automobile industry, are obvious examples. This remarkable development is the result of technology scaling that led to fabrication of Integrated Circuit (IC) with smaller feature sizes, higher levels of integration and faster operating frequencies. The process of device scaling evolved from few micrometers to nanometers today, and the circuit complexity has advanced from Small-Scale Integration (SSI) in 1960s to Giga-Scale Integration (GSI) in 2000s. It is predicted that this integration continues at a faster speed towards a trillion transistors per chip, Tera-Scale Integration (TSI) era, in 2020s. Today, not only digital devices and memories, but also analog/mixed-signal blocks, MEMS based sensors, and other functional blocks are being integrated on the same die to build a complete system. However, the benefits of system integration are significantly reduced without efficient communication between these blocks. Thus, this thesis addresses the problems of global on-chip interconnects using novel circuit level techniques.

1.1 Emergence of Interconnect-Centric Design

The performance of transistors is continually improved through scaling. However, the impact of technology scaling on long wires is reverse. In order to tackle this problem, there is a shift in system design approach from computation-centric to interconnect-centric.

1.1.1 Device and Interconnect Scaling

The primary goals of technology scaling are decreasing gate delay, increasing gate density, and reducing energy per storing/operating [1]. At the moment, the feature size is scaling down at a rate of 0.7 per year [105] in compliance with Moore's Law [2]. This decreases the gate delay by 30%, doubles the gate density and reduces the energy per switching by 65%. As the history of integration density reveals, in 1960 one transistor consisted of 10^{20} atoms in a volume of 0.1 cm^3 and in 2000 these number were 10^7 atoms in 0.01^3 , leading to a higher capacity of integration. Similarly, the energy for storing/operating 1 bit is reduced because the energy required for charging and discharging capacitors is lowered due to the reduction in capacitors area from 1 cm^2 to 0.01^2 and furthermore the supply voltages are scaled down from 10 V to 1 V [3].

Scaling down of transistor's dimensions leads to improvements in both transistor cost and performance. However, scaling down of interconnect cross-sectional dimensions degrades performance. The ideal scaling of interconnect assumes that the width and height of the wire are reduced with the same scaling factor as gates' dimensions, leading to taller and narrower wires. As a result, the resistance of a unit length wire increases at the rate of 104% per year. The length of local wires scales the same way as the logic, whereas global wires tend to track the chip dimensions. In general, die area should decrease by 50% per year in successive technology, but new designs integrate more transistors and functionality per chip, resulting in a need for die area increment. Die area has been increasing 13% per year. Consequently global interconnect length increases at a rate of 6% per year, and its RC time constant increases by approximately 130% per year.

To mitigate the increase in wire delay, various techniques have been developed from both geometric structure and materials perspective, such as high aspect ratio, multiple-layer metalization, copper technology and low-k dielectrics. A higher aspect ratio along with smaller wire pitch leads to a reduction in RC delay but this approach has two problems. First, manufacturing of lines and vias with aspect ratio larger than 4 becomes unreliable due to the difficulty in filling a deep and narrow trench completely with metal. Second, the increase in line thickness due to a higher aspect ratio results in a larger coupling capacitance to neighboring wires, which increases both the RC delay component and the signal coupling noise. These two undesirable

effects limit the practicality of this technique for nanometer technologies. In multiple-layer metallization approach, different scaling factors for local and global interconnects are used to satisfy the need for higher density, reduced RC delay, and smaller resistive loss. In order to match the device density on the substrate and maintain the RC delay, the pitches and length of local wires scales at a much faster rate than the vertical dimensions. For global wires, the scaling is determined by the length of the chip edge, and as a result signal delay on global wires increases continuously from one technology generation to the next. This reverse scaling of global interconnect is an undesirable consequence of technology scaling. From materials perspective two major advances have been done, the change in metal material from aluminum to copper and the introduction of low-k dielectrics to replace silicon oxide (SiO_2). However, in nanometer technology, all these techniques are insufficient to achieve the needed high-speed global on-chip communication. Hence, requiring additional techniques and other approaches.

1.1.2 System-on-Chip and Multicore Systems

The continued scaling of the semiconductor technology creates the potential of System-on-Chip (SoC) integration, that is, the integration of a complete electronic system including all its periphery and its interfaces to the outside world on a single die. SoC consists of several heterogeneous components with different implementation styles such as programmable processors, dedicated hardware to perform specific tasks, on-chip memories, input-output interfaces, and on-chip communication architecture that serves as the interconnection fabric for communication between these components. For instance, Intel's 45nm Xeon® EX processor (Nehalem-EX) is a SoC which has eight 64-bit cores and a 24 MB shared L3 cache (Figure 1.1). At the top stripe it has four Quick Path Interconnect (QPI) links, while the bottom stripe houses the Scalable Memory Interconnect (SMI) links. It also has a system interface that includes two memory controllers, two hub interfaces to the last level cache, an 8-port router, the power control unit (PCU) and the DFX control box. Usually such building blocks can be shared and also re-used as Intellectual Property (IP) blocks, which further improves the productivity and reduces time-to-market.

Today, SoC combines a diverse set of components using adaptive circuits, integrated sensors, sophisticated power-management techniques, and increased




Figure 1.1: Intel’s 45nm 8-core Xeon-EX processor [5].

parallelism to build products that are multicore and multi-function. Examples of such prototypes are Intel’s 80 core TeraFLOPS [64], 167-processor computational platform [59], and FAUST chip (a reconfigurable baseband platform consisting of 23 computing units that can be configured to support the functions of specific baseband processing) [83]. This trend will continue and it will open up the feasibility of a wide range of applications, such as data mining, visual computing, and recognition, making use of massive parallel processing and tightly interdependent processes which brings into front the underlying interconnection capability. The interconnection between SoC components should provide reliable routing of data from the source to destination. It must also be able to guarantee latency or bandwidth to ensure that the application performance constraints are met.

1.1.3 Network-on-Chip

The increasing number of IP cores that can be integrated on a single chip enables implementation of complex applications using the SoC approach. The huge communication demands of these applications and the abundant computation power available on-chip put tremendous pressure on the communication architecture. Consequently, scalable communication architectures are needed for efficient implementation of SoC. Simple on-chip communication solutions do not scale up when the number of processing and storage arrays on a chip increases. For example, on-chip buses can serve a limited number of units, and beyond that performance degrades due to the bus parasitic capacitance and

the complexity of arbitration. Network-on-Chip (NoC) is a communication infrastructure targeted for SoC consisting of tens or hundreds of resources. NoCs are an attempt to scale down the concepts of large-scale networks, and apply them to SoC domain. It separates the concerns of communication from computation by building on-chip communication structure. Each component of a SoC is viewed as a node of the on-chip communication network. NoC use packets to route data from the source to the destination component, via a network fabric that consists of switches(routers) and point-to-point links, which connect the resources to routers as well as the routers to each other to form a network. NoC provides better scalability than on-chip buses because as more resources are introduced to a system, also more routers and links are introduced to connect them to the network. The additional links and routers provide the communication capacity needed for the new resources. Many NoC realizations inherently contain some redundancy in the communication media, which can be used to provide a higher reliability and traffic balancing. This is in contrast to bus structures, which rely on a single communication medium.

NoC has been initially proposed as a design paradigm for on-chip communication in the beginning of this millennium [15, 16, 17, 44, 74]. Today, there are NoCs in commercial use such as *ArterisTM* [134] and *STNoCTM* [135, 136] as well as industrial products which use NoC as a communication backbone. The *TILE64TM* 64-core processor from Tiler [18, 61] and the 80-core Intel's TeraFLOPs processor [64] are recent examples of industrial products which proves the feasibility and potential of NoC. The interconnects designed and presented in this thesis can be used as a link between two NoC routers.

1.2 Challenges of Global On-Chip Interconnect

On-chip interconnect has become a primary challenge for high performance high complexity SoCs. Transmitting clock, data, and communication signals over large die areas requires long interconnections among the various circuit modules. As technology scales, the interconnect cross section decreases while operating frequencies increase. The impact of these trends on high performance systems is significant. Long interconnects with smaller cross sections exhibit increased capacitance and resistance, resulting in larger power consumption, and higher latency. Furthermore, wire inductance can no longer be ignored due to high signal frequencies and long wire lengths. The in-

creasing number of cores per chip also places a premium on high-bandwidth, low-latency and low-power links between cores.

1.2.1 Performance and Power Consumption

The higher wire resistance, the increase in wire length and reduced wire spacing cause the global wire delay to increase considerably compared to the gate delay [105]. The gap between global interconnect delay and gate delay increases with technology scaling as can be seen from Figure 1.2. Furthermore, the increase in die size due to increasing chip functionality makes it more difficult to deliver signals across the chip in one clock cycle [14]. According to the prediction of International Technology Roadmap for Semiconductors (ITRS), at 45nm technology node, the RC delay is 542 *ps* for a 1 *mm* long minimum pitch copper global wire, whereas the clock frequency will reach 10 *GHz* (equivalent to 100 *ps* cycle time). The conventional approach to deal with this is to use pipelining in global signals, which increases the latency and power consumption when routing signals across functional blocks. Also, the total on-chip wire length will increase linearly with technology, reaching about 2.22 *km/cm*² by the year 2010 [105]. This trend supports the assumption that long interconnects will be significant in future technologies. To combat these phenomena, traditional repeater insertion methods have been widely developed and adopted. Unfortunately, as interconnect lengths increase, the required number of repeaters increases tremendously. This results in significant power dissipation, increased delay, and larger area.

Interconnects, especially the global wires have also become a major source of power consumption. It has been reported that wire capacitance can take up to 70% of the total chip capacitance in contemporary designs [6]. Moreover, a rapid increase in chip operating frequencies further exacerbates the amount of dynamic power dissipated in the interconnect. Magen et al. found that interconnection power accounts for half the total dynamic power of a 130nm microprocessor, and nearly 50% of the interconnect power is consumed by global wires [7]. With a projection that without changes in design philosophy, in the next five years up to 80% of microprocessor power will be consumed by interconnect.

As discussed in Section 1.1.3, NoCs have emerged as the seemingly best candidate to connect the cores on present and future SoCs. Latency and power consumption of NoC are among the critical challenges and need to be




Figure 1.2: Delay comparison [105]

addressed at all abstraction levels [156]. The latency of networks is too large, leading to performance degradation when they are used in high performance systems. The power consumption of NoC implemented with current techniques is too high, by a factor of 10, to meet the expected needs of future SoCs. In NoC the network interconnects consume a significant part of the total power budget. For example, in TeraFLOPS [64] the network consumes up to 39% of the total chip power (76 W when operating at 5.1 GHz) [8]. 17% of the network power is consumed in the links (13 W at 5.1 GHz). Therefore, more emphasis should be put on circuit techniques that increase signal velocity on channels and reduce the power consumption of the interconnects.

1.2.2 Variability and Reliability

Variability has become a major challenge for designs in sub-100nm technology nodes and it is considered one of the primary limiters for technology scaling [9] - [12], [113]. It is affecting device as well as interconnecting wire parameters. The inability to precisely control the manufacturing process leads to unpredictable device and wire characteristics, which in turn cause performance and power variability besides error-prone behavior [13], [113] - [125]. According to ITRS, within a few years delay and power variability reaches 63% and 76%, respectively [110]. In addition, systems performance is also affected by the environment in which it operates such as temperature, power supply voltage

and noise.

The performance of an on-chip interconnect is determined by the electrical characteristics of the signaling circuit's devices and the interconnecting wire parasitics. From a process perspective, almost all manufacturing phases, etches, thin-film deposition, hot processes, and even wafer clean processes, influence device parameters and thus contribute to variabilities. Increased process complexity related to subwavelength lithography, chemical-mechanical polishing, and the implementation of low-k dielectrics leads to higher variability of wire resistance and capacitance. Variations in operating environment, spatial as well as temporal effects, can also have a similar impact. For example, the effective supply voltage of a transistor may vary across the chip due to changes in the voltage drop along the power grid. The local operating temperature of a transistor is affected by local variations in power dissipation. Crosstalk, resulting from capacitive and inductive coupling, could severely affect the timing and signal integrity of an interconnect. Because each victim wire experiences a different capacitive coupling length or a different inductive-coupling return path, the interconnect exhibits varying signal propagation delays under different switching patterns. All these variations cause the signal propagation delay of the interconnect to be uncertain which in turn affects the performance and reliability of the communication significantly.

Traditionally corner based analysis has been used to guard against yield loss resulting from these variations; however, with increasing number of sources of variation, corner based methods are becoming overly pessimistic and computationally expensive. Self-timed design methodologies can make the communication resilient to delay variations. More specifically, self-timed delay-insensitive links can operate correctly in the presence of delay variations in gates and interconnecting wires.

1.3 Global On-Chip Communication Techniques

SoC consists of many Intellectual Property (IP) blocks. The different functions among different SoC blocks naturally cause them to work in different clock rates for optimal performance. Hence, coordination and communication between these components become challenging. Globally Asynchronous Locally Synchronous (GALS) scheme has been proposed as a solution. The idea of GALS is to partition a system into separate clock domains, which run at

different clock rates, and the separated domains communicate with each other in an asynchronous manner.

1.3.1 GALS Communication

Globally synchronous communication is a thing of past because it is difficult to design with growing chip sizes, clock rates, relative wire delays and parameter variations. Moreover, high speed global clocks consume a significant portion of system power budgets and lack the flexibility to independently control the clock frequencies of submodules to achieve high energy efficiency. GALS facilitate the integration of independently designed blocks operating at different frequencies as well as fast block reuse by providing wrapper circuits to handle the inter-block communication.

NoCs with GALS clocking styles have been used in many proposed network designs and are expected to be an attractive approach to overcome many of the timing problems [74]. GALS simplifies clock tree design and results in easily scalable clocking systems. It also allows better energy savings since each functional unit can easily have its own independent clock and voltage [75]. Furthermore, it enables easy implementation of distributed power management system for the entire chip [76].

Generally there are two different implementation of GALS NoC: fully asynchronous (self-timed) and multi-synchronous. In self-timed GALS NoC, IP blocks use locally generated clock and there is a synchronous \Leftrightarrow asynchronous interface between the network and the synchronous IP. Clockless networks such as MANGO [81], ANoC [82], ALPIN [76], FAUST chip [83], and QNoC [84] are examples of self-timed NoC. A systematic comparison between these two implementations shows that the self-timed network gives better saturation threshold, smaller average power consumption, slightly higher maximal bandwidth and 2.5 times smaller packet latency than the multi-synchronous implementation [53]. Furthermore, the risk of metastability introduced by the multiple bi-synchronous FIFOs used in the multi-synchronous implementation can be a critical issue. This risk is much lower in the self-timed approach since the metastability is entirely confined in the synchronous \Leftrightarrow asynchronous interface. Due to these advantages self-timed delay-insensitive interconnects are designed in this thesis which can be used in clockless NoC.

1.3.2 Self-timed Delay-Insensitive Communication

Delay-insensitive codes have been used in many applications for error detection and delay-insensitive communication. Their main feature is the ability of allowing the correct interpretation of the code word independently of the delay of individual bits. Hence, self-timed delay-insensitive data transfer is one of the most promising approaches to deal with delay uncertainties in on-chip interconnects. A self-timed delay-insensitive communication link assumes nothing about the delays in the wires and devices except that they are finite and positive, and therefore the reliability of communication is unaffected by the delay variations. Several delay-insensitive coding schemes have been proposed, but effective Complementary Metal Oxide Semiconductor (CMOS) implementations are needed in order to make feasible self-timed on-chip interconnect. Dual-rail and 1-of-4 codes are well known and mostly used for on-chip delay-insensitive communication [54].

The delay insensitivity feature does not come free of cost, it has delay and area overhead. The data encoding at the transmitting side and decoding at the receiver side, as well as completion detection cause additional delay to the communication. In this thesis, different performance enhancement techniques for delay-insensitive interconnects are developed and implemented in order to compensate the delay overhead and achieve high performance communication.

1.4 Scope of Thesis and Contributions

This thesis addresses the challenges of global on-chip interconnects, more specifically performance, power consumption, delay variability and signal integrity. The author examines, develops, and integrates different circuit level techniques and implements them in order to mitigate or withstand these problems. Since signal propagation delay variations are unavoidable in nanometer CMOS technologies, the focus of this thesis is on design and optimization of delay-insensitive on-chip interconnects. Therefore, a self-timed delay-insensitive data transmission scheme is adopted for all interconnects designed in this thesis. The presented interconnects use custom designed circuits as in most of self-timed designs and can be integrated to an automated IC design flow as an IP.

The main contributions of this thesis are fourfold. The first one is circuit and signaling solutions to achieve high performance and power efficient delay-

insensitive on-chip communication. Second, high speed completion detection technique and its circuit implementation in order to overcome the performance penalty of conventional detection methods. Third, high throughput and energy efficient semi-serial link for long-range on-chip communication. The last contribution is circuit level techniques to ensure signal integrity of current sensing interconnects despite PVT variations. A brief summary of each contribution along with the related publications are discussed in the following subsections.

1.4.1 Delay-Insensitive Current Sensing Interconnects

Self-timed delay-insensitive data transmission has delay overhead due to data encoding/decoding and completion detection. To minimize this overhead in power efficient manner, formulation, customization and integration of several techniques are done. This encompasses data encoding/decoding, communication protocols and signaling schemes along with design of their data encoder/decoder, completion detector, driver and receiver circuits. Based on the developed techniques and circuits, three low latency and high throughput delay-insensitive current sensing interconnects are designed and verified through transistor-level simulations. Their performance and power consumption are analyzed and compared with conventionally implemented delay-insensitive on-chip interconnects as well as each other.

Related Publications

1. Ethiopia Nigussie, Teijo Lehtonen, Sampo Tuuna, Juha Plosila and Jouni Isoaho. High-Performance Long NoC Link Using Delay-Insensitive Current-Mode Signaling. VLSI Design, Vol. 2007, Article ID 46514, 13 pages, Mar. 2007.

Technical Contributions: Implementation of high-performance interconnect based on multilevel current sensing signaling and delay-insensitive two-phase 1-of-4 encoding is presented. A communication protocol which enables power saving by using multilevel current signaling is developed. Data encoder/decoder circuits which convert between two-phase bundled-data and two-phase 1-of-4 encoding formats and completion detection circuit are designed. In addition, multilevel current-mode driver and current sensing receiver circuits are designed.

2. Ethiopia Nigussie, Juha Plosila and Jouni Isoaho. Current-Mode On-Chip Interconnect using Level-Encoded Two-Phase Dual-Rail Encoding. In 2007 IEEE International Symposium on Circuits and Systems (ISCAS 2007), pages 649-652, May 2007.

Technical Contributions: As the exploration for efficient delay-insensitive communication continues, another on-chip interconnect which uses LEDR encoding and current sensing signaling is implemented and verified. LEDR encoder and decoder circuits which convert between two-phase bundled-data and LEDR encoding along with completion detection circuit is designed. Furthermore, current-mode driver and current sensing receiver is designed.

3. Ethiopia Nigussie, Juha Plosila and Jouni Isoaho. Area Efficient Delay-Insensitive and Differential Current Sensing On-Chip Interconnect. In IEEE International SoC Conference (SOCC 2008), Pages 143-146, Sept. 2008.

Technical Contributions: Implementation of area and power efficient on-chip interconnect based on integration of two-phase dual-rail encoding and differential current sensing signaling is presented. The contribution includes developing a communication protocol which integrates two-phase dual-rail encoding and differential current sensing signaling in area and power efficient manner. Moreover, a signaling technique which senses both the direction of current flow and current levels is formulated and implemented. Data encoder, differential current-mode driver, current direction sensor, current comparator, differential amplifier and completion detection circuits are also designed.

1.4.2 Completion Detection Technique to Enhance Performance

Conventional completion detection requires logic circuitry, the delay of which increases drastically when the channel bit width increases, making delay insensitive interconnects problematic for high performance systems. In order to solve this performance bottleneck, a novel completion detection technique for delay insensitive current sensing on-chip interconnects is presented. The proposed completion detection technique directly uses the current on each data wire and carries out completion detection in the current mode, enabling high

speed operation. Furthermore, its speed does not degrade when increasing the channel bit width.

Related Publications

1. Ethiopia Nigussie, Juha Plosila and Jouni Isoaho. High-speed completion detection for current sensing on-chip interconnects. In *Electronics Letters*, 45(11), pages 547-548, May 2009.

Technical Contributions: High-speed and bit width insensitive completion detection technique along with its CMOS implementation is presented. Its performance and power consumption are also analyzed and compared with conventional detection scheme.

2. Ethiopia Nigussie, Juha Plosila, Sampo Tuuna, Jouni Isoaho and Hannu Tenhunen. Boosting Performance of Self-Timed Delay-Insensitive Bit Parallel On-Chip Interconnects. *IET Circuits, Devices and Systems Journal*, under review.

Technical Contributions: Implementations of two delay-insensitive current sensing interconnects which use the proposed completion detection method are presented. The benefit of the proposed completion detection is verified through simulations by varying wire lengths and channel bit widths. In addition, performance, power efficiency and area comparisons with conventionally implemented interconnects are performed.

1.4.3 High Throughput Energy Efficient Semi-Serial Link

The throughput of the interconnects which are designed so far decrease considerably in long-range communications. This is due to in part the need to acknowledge every transmission, the backward latency. So, the goal here is to achieve high-throughput and maintain it regardless of the communication distance. Furthermore, to dissipate low energy per bit. The presented serial link implementation is based on the design of high speed serializer/deserializer circuits, formulation of pulse dual-rail encoding, and integration of wave-pipelining, pulse signaling and differential current-mode signaling. The formulated pulse dual-rail encoding provides an opportunity to implement pulse signaling at no cost and allows per word acknowledgment without losing the delay insensitivity feature. The integration of wave-pipelining, pulse signaling

and differential current-mode signaling schemes leads to significant power savings while maintaining the throughput. This integration also helps to achieve fast decoding of data validity indicator signal.

Related Publication

1. Ethiopia Nigussie, Juha Plosila, Sampo Tuuna, Jouni Isoaho and Hannu Tenhunen. Energy Efficient Semi-Serial On-Chip Link Through Circuit Optimizations and Integration of Signaling Techniques . IEEE Transactions on Very Large Scale Integration (VLSI) Systems, under review.

Technical Contributions: Design of high throughput and energy efficient semi-serial link is presented. A novel communication protocol and delay-insensitive encoding are formulated. Integration and implementation of pulse signaling, differential current mode signaling and wave-pipelining are done. Serializer/deserializer, pulse dual-rail encoder, differential current mode driver, differential amplifier receiver, and data validity decoder circuits are also designed.

1.4.4 Circuit Techniques for PVT Variation Tolerance

It has been verified that the current sensing interconnects which are designed in this thesis have achieved higher performance, better power efficiency and smaller area compared to optimally pipelined voltage-mode interconnects. These interconnects use a current sensing receiver and the signal integrity of this type of receiver can be affected by process, voltage and temperature (PVT) variations. The conventional approach of allocating large current margin by assuming worst-case variation has high power consumption cost. Hence, here two techniques to guarantee signal integrity of a current sensing interconnect with low power overhead are developed. The first one deals with process variation and the other one with voltage and temperature variations. To tolerate the effect of process variation, interconnect calibration technique which adjusts the current by detecting the existing amount of variation at every power start-up of the system is developed and implemented. Online monitoring of receiver's currents variation is devised and implemented in order to withstand the effect of voltage and temperature variations. If the variation violates the margin, receiver reconfiguration and request for retransmission will be carried out. These two PVT variation tolerance schemes allow to maintain the signal integrity of an interconnect with low power overhead.

Related Publications

1. Ethiopia Nigussie, Juha Plosila and Jouni Isoaho. Process Variation Tolerant On-Chip Communication Using Receiver and Driver Reconfiguration. In 11th IEEE International Symposium on Quality Electronic Design (ISQED 2010), Pages 453-460, Mar 2010.

Technical Contributions: Process variation tolerance technique for current sensing interconnects is developed and implemented. An error detection scheme as well as a reconfiguration algorithm and methodology are developed. Furthermore, reconfiguration control and communication circuits in addition to reconfigurable driver and receiver circuits are designed. The technique is implemented in two current sensing interconnects and simulated for different bit width transmissions. In addition, analysis of the calibration delays for best-, average- and worst-case variations, power and area overheads are performed.

2. Ethiopia Nigussie, Juha Plosila and Jouni Isoaho. Monitoring and Reconfiguration Techniques for Power Supply Variation Tolerant on-Chip Links. In 2010 IEEE International Symposium on Circuits and Systems (ISCAS 2010), 4 pages, May 2010.

Technical Contributions: A runtime power supply and temperature variation tolerance technique for current sensing interconnects is designed. Algorithm and methodology, which consists of sensing the effect of variations, reconfiguration control, retransmission request and reporting error to higher level error controlling system, are developed. The technique is implemented in two current sensing interconnects which includes the design of variation sensor, reconfiguration control, retransmission request and error reporting circuits. The two interconnects which use the proposed technique are simulated. The power and area overheads are examined for 2- to 64-bits transmissions.

1.5 Related Work

Increasing attention is placed on the design of on-chip interconnects due to the dominant limitation of global interconnect signal delays, power dissipation and delay uncertainty on overall system performance and reliability. It is imperative that future on-chip interconnect designs overcome these challenges.

The conventional technique to improve the interconnect delay bottleneck is to insert repeaters by breaking the wire into several sections [106, 149]. Usually these wire sections are highly capacitive and high strength repeaters are needed. The adverse effect of this is increased power consumption; it has been estimated that over 50% of the power in a high performance microprocessor is dissipated by repeaters charging and discharging interconnects [6, 7]. The other approach is inserting register pipelines [151]- [154] to increase the communication throughput. This approach increases the latency of the communication, and furthermore, the number of registers needed increases with the size and the complexity of the system. This in turn increases the power consumption. These show that the conventional solutions are inadequate to meet the overall performance requirements of high performance electronic systems in nanometer regime.

High Performance Interconnect

In order to solve the delay and power problems of global interconnects, several alternatives have been proposed by ITRS [58]. Using different signaling methods is among the proposed alternatives. This approach utilizes available technology with innovative approaches to signaling and circuit operation to implement high speed global interconnects. In [40], high speed on-chip signaling method that relies on differential current-mode sensing to improve both delay and energy dissipation has been proposed and implemented. High speed and power efficient on-chip interconnect has been demonstrated using current mode signaling along with circuit techniques in [33]. In [39], a 10Gbps/channel on-chip signaling system has been fabricated in 90nm technology. It consists of current mode logic driver and receiver, and differential transmission line. By using impedance-unmatched driver it saves the energy per bit by 21% compared with a conventional impedance-matched driver. Energy-aware differential current sensing signaling through the use of differential leakage-aware amplifier has been proposed in [137]. A method to propagate signal near the speed of light has been demonstrated in [34], though the interconnect has high power consumption. Wave-pipelining has also been proposed for on-chip interconnects as a means to increase throughput [77, 103, 138, 150]. There are researchers whose focus is on using signal conditioning, and high-speed transceivers in order to improve interconnect throughput [139]-[143]. All of these works concentrate in achieving high performance communication with-

out dealing properly the reliability problem due to delay variations.

Variation Tolerant Interconnect

In nanometer scale technologies sources of variability are increasing and unavoidable, which creates several challenges in building reliable systems. Variability causes signal propagation delay uncertainty in interconnects which in turn causes error. Due to this, variation tolerant on-chip interconnects are needed. Self-timed design methodologies can make the interconnect robust to delay variations. Different self-timed delay-insensitive interconnects have been proposed in [37, 55, 56, 107, 145, 146, 147, 148, 155, 108]. Most of these works concentrate only in the delay insensitivity feature and ignore the need for high performance interconnects. For example, the work in [37, 55, 56, 107, 145, 155, 148] use four-phase handshaking, which requires four traversals of the long wire per each data transfer. This decreases the performance of such interconnects significantly. In [145], delay-insensitive encoding, which minimizes the wiring overhead has been proposed but the encoder, decoder and completion detection logic complexity increases which increases the delay overhead and consequently reduces the throughput of the communication. An asynchronous DI interconnect which uses two-phase dual-rail encoding has been implemented and compared with synchronous interconnect in [108]. To improve the throughput of this interconnect locally clocked pipeline stages have been inserted, which increase both power consumption and the required area.

High Performance and Variation Tolerant Interconnect

The aim of this thesis is to achieve both high performance and variation tolerant on-chip communication. The designed interconnects use two-phase handshaking and self-timed delay insensitive data transfer. To compensate the delay overhead due to delay-insensitive encoding, decoding and completion detection, different high-speed signaling techniques have been implemented. In addition, to minimize the delay overhead due to completion detection of wide bit transmissions, bit width insensitive high-speed completion detection technique has been developed and implemented. Furthermore, self-calibration, monitoring and reconfiguration techniques have been developed to guarantee the signal integrity of the interconnects despite Process, Voltage and Temperature (PVT)

variations.

1.6 Thesis Organization

The rest of the thesis is organized as follows. In Chapter 2, the design techniques used to implement the presented high performance delay-insensitive interconnects are discussed. Modeling of on-chip wires is also discussed in the last section of this chapter. Design and analysis of the three delay-insensitive current sensing on-chip interconnects are presented in Chapter 3. In addition, analysis of their performance and power consumption as well as comparison with conventional delay-insensitive on-chip interconnects are presented. In Chapter 4, a high speed completion detection technique as well as its design is presented in order to enhance the performance of the delay-insensitive interconnects. Furthermore, two of the interconnects presented in Chapter 3 are redesigned and presented as case studies to demonstrate the advantage of the presented completion detection. Analysis of their performance, energy dissipation and area besides comparison with the reference cases are also presented. In Chapter 5, implementation and analysis of high-throughput serial on-chip interconnect targeted for long-range communication is presented. Also, comparison of throughput, energy and area between fully bit-parallel, bit-serial and semi-serial links are performed. All the interconnects which are presented in Chapter 3 and 5 are redesigned using 65nm CMOS technology and their performance, energy dissipation, and area are compared in Chapter 6. In Chapter 7, circuit techniques as well as implementations to tolerate process, supply voltage and temperature variation effects on the signal integrity of the interconnects are presented. Finally, Chapter 8 concludes the thesis and presents future direction of this research.

Chapter 2

Interconnect Design Techniques

In this chapter, power efficient design techniques for the delay-insensitive global on-chip interconnects are presented. It is a foundation work for the rest of the chapters. The chapter starts with the handshaking protocols and is followed by discussion of data encoding, decoding, and completion detection techniques. Furthermore, customized and advanced signaling techniques that have been designed during the course of this thesis are also explained in detail. The last section describes the electrical and magnetic view of interconnecting wires and the way they are modeled.

2.1 Handshaking Protocols

In self-timed on-chip communication, handshaking protocol is used to transmit data between a sender and a receiver. The sender delivers data onto the channel and the receiver accepts data from the channel. The communication parties can be further classified as follows: active part initiates the data transfer and passive part responds to the active. The transfer direction is determined by the communication protocol.

The exchange of data through a channel is negotiated between the sender and the receiver using a handshaking protocol. For every data transfer, a request is transmitted by the active module which indicates the validity of the data on the channel. An acknowledgment transmission from the passive module indicates data acceptance and readiness of the receiver to accept the




Figure 2.1: Push channel.




Figure 2.2: Pull channel.

next data. The transmission of request and acknowledgment signals may occur on dedicated signaling wires or may be implicit in the data depending on the data encoding technique used but in either case, one event indicates data validity and the other data acceptance. The flow of data relative to the request event determines whether the channel is a push or pull channel. In a push channel data flows in the same direction as the request whereas in the pull channel data flows in the same direction as acknowledgment signal. These two types of channels are illustrated in Figure 2.1 and 2.2. The push channel is assumed in all of the interconnects designed in this thesis.

In general, the request and acknowledgment signals may be transmitted using one of the two protocols described below; a two-phase transition-based handshaking also called a non return-to-zero protocol or a four-phase level-based protocol (a return-to-zero scheme). There are other customized protocols such as single-track [19, 20], and one-phase [21, 22, 23]. The advantage of single-track handshaking is that it requires just two transitions per data transfer as opposed to four in four-phase protocol and avoids the requirement for event-triggered logic circuits of the two-phase protocol. However, the implemented circuit will run correctly only if it is not exposed to heavy ambient noise because single-track protocol relies momentarily on high impedance states on wires. The one-phase protocol requires only one communication action between the sender and the receiver which makes it faster than both two and four-phase handshaking. It uses a data coloring scheme to indicate data validity and acceptance. The transmitted symbol consists of both bit value




Figure 2.3: Four-phase handshaking protocol.

and color information. There is a color detector circuit at both the transmitter and the receiver. The detector detects the signal in the wire and extracts the color information. The receiver accepts the data and changes its color when the color in the wire is the same as its own (data is valid). Also the transmitter sends the next data with new color after its color is the same as the wire color (data acceptance). Due to color detection at both sides, there is no need to transmit either the data validity or acceptance to one another. Even if it is attractive in saving communication time, it requires complex circuits and incurs additional power consumption.

The four-phase handshaking protocol, shown in Figure 2.3, uses signal levels to indicate the validity of data and its acceptance by the receiver. That is the sender issues data and sets the request high, the receiver absorbs the data and sets the acknowledgment high. The sender responds by setting the request low (at this point data is no longer valid) and the sender acknowledges this by setting the acknowledgment low. This requires two transitions per data transfer on both request and acknowledgment signals. With increasing wire delay due to the reverse scaling effect of long wires, accommodating the return-to-zero phase leads to a significant reduction in throughput which makes the four-phase handshaking unattractive for global on-chip communication.

Unlike the four-phase protocol the signal levels are unimportant in the two-phase handshaking protocol where the information is carried by the transition. Both rising and falling transitions are equivalent, each being interpreted as a




Figure 2.4: Two-phase handshaking protocol.

handshaking event. A push channel that uses the two-phase protocol passes data using a request signal transition, and acknowledges data reception with an acknowledge signal transition. Two-phase handshaking is preferred for long on-chip communication since it reduces the required number of transitions by half and avoids the requirement of spacer compared to four-phase signaling [25]. This saves communication time and energy of the system significantly. Figure 2.4 illustrates the two-phase protocol.

Some may argue that the two-phase communication requires edge sensitive control logic circuits, which leads to considerable delay overhead. But the delays of edge-sensitive logic circuits are much smaller than the global wire delay in nanometer CMOS technologies, which makes the use of two-phase handshaking advantageous over the four-phase. For example, in 65nm CMOS technology, global wire delay is nine times the gate delay [105].

2.2 Data Encoding Techniques

So far the two-phase and four-phase handshaking protocols are presented. Another dimension of self-timed communication is the use of data encoding. A communication can be carried out using control wires separately of data. This approach is known as bundled-data encoding where it is assumed that by the time request arrives, data is already arrived and is stable. In other words,




Figure 2.5: Delay-insensitive push channel.

the delay in data validity indicator (request) wire must be larger than the delay in the data wire. To remove this timing constraints, the data validity indicator signal is included in the data itself resulting a delay-insensitive communication. As already explained, delay-insensitive communication is a necessity for global on-chip communication due to the unavoidable PVT variations and the resulting delay uncertainties in nanometer regime. Most of the existing GALS communication wrappers however have bundled-data encoding interface [26, 27, 28, 29]. Hence, there is a need to convert the single-rail data representation to delay-insensitive encoding for global communication. There are many types of delay-insensitive encodings, but the most commonly used in on-chip implementations are dual-rail (1-of-2) and quad-rail (1-of-4) encodings [54]. In a delay-insensitive channel, to transmit N -bit data in parallel it requires $2N + 1$ wires. Only one handshake wire is required since data itself acts as data validity or acceptance indicator depending on the channel type (only acknowledgment wire for push channel and request wire for pull channel). A delay-insensitive channel is shown in Figure 2.5.

Dual-rail encoding uses two signals to represent each bit of information, and therefore, to transmit N bits of data $2N$ wires are required. Each bit transfer will involve activity in only one of the two wires. As in all delay-insensitive codes, timing information is implicit in the code, that is, it is possible to determine when the entire data word is valid. This is done, for instance, by detecting a level in a four-phase dual-rail data transfer or by detecting a transition in a two-phase transmission on one of the two wires for every bit in the word. A separate wire to convey data readiness is thus not necessary. Transmission of four consecutive bits, 1001 using four and two-phase dual-rail encoding is shown in Figures 2.6 and 2.7, respectively.

There are other customized dual-rail encodings such as Two-phase Level-Encoded Dual-Rail (LEDR) encoding, one-phase dual-rail encoding and pulse dual-rail encoding aimed at either minimizing the timing overhead or circuit

complexity. LEDR encoding is used in one of the on-chip interconnects presented in this thesis (See Section 3.1). Pulse dual-rail encoding has been formulated and used along with wave-pipelining in the serial on-chip link, presented in Chapter 5. It encodes each bit into Pulse and No Pulse (P, NP) pair depending on bundled-data and request signal from the transmitter.

In 1-of-4 data encoding, a group of four wires is used to transmit two bits of information per symbol. A symbol is one of the two-bit codes 00, 01, 10, or 11 and it is transmitted through activity on one of the four wires. Since it is possible to detect the arrival of each symbol at the receiver, 1-of-4 encoding is delay-insensitive. Besides being delay-insensitive, 1-of-4 encoding has more immunity against crosstalk effects as compared to single-rail (bundled-data) encoding, because the likelihood of two adjacent wires switching at the same time is much smaller. In dual-rail encoding, representation of a valid N -bit value requires $2N$ transitions, whereas in 1-of-4 it requires only N transitions. The reduction of transitions in 1-of-4 encoding decreases the dynamic power consumption due to the lower wire capacitance. Transmission of three consecutive symbols using four-phase and two-phase 1-of-4 encoding is illustrated in Figure 2.8 and Figure 2.9, respectively. Four-phase 1-of-4 encoding with voltage-mode pipelining signaling has been used in [56]. However, in nanometer era where the wire delay dominates over the gate delay the use




Figure 2.6: Four-phase dual-rail encoded transmission




Figure 2.7: Two-phase dual-rail encoded transmission




Figure 2.8: Four-phase 1-of-4 encoded transmission




Figure 2.9: Two-phase 1-of-4 encoded transmission

of four-phase handshaking has significant delay overhead as it requires four communications per transfer. Therefore, the two-phase 1-of-4 data encoding is used in one of the high-performance current sensing on-chip interconnects presented in Chapter 3.

2.3 Data Decoding Techniques

In delay-insensitive data transmission, the receiver has to decode the transmitted encoded data. The complexity of the decoding circuit depends on the chosen encoding. The simpler the decoding logic is, the more attractive is the encoding.

The decoder of four-phase dual-rail encoded channel detects whether one of the dual-rail wires is set to high state or not. The high state indicates a valid data. This detection can be done using an OR gate. The two-phase dual-rail transmission requires more complex decoding logic because it has to detect current transitions on both wires. Furthermore, the decoder has to compare the current transition with the previous transition. Due to this

it is not suitable for high-performance communication. In LEDR encoded transmission, data is decoded directly from the state wire using an inverter or buffer to make it full swing. This technique does not require complex decoding logic and the required number of communication actions is two (similar with two-phase dual-rail). Therefore, LEDR is a better alternative over two-phase dual-rail. The pulse dual-rail encoded transmission together with differential pulse signaling, requires no decoding logic. That is, the receiver output provided by the differential amplifier is the transmitted data. This increases the throughput and it is one of the reasons for formulating this type of encoding for the serial link presented in Chapter 5.

The decoding of 1-of-4 encoded data has similar problems as the dual-rail encoding. That is, the decoder needs to sense the voltage levels of the wires, which requires two 2-input OR gates per one 1-of-4 group. The decoding of voltage-mode two-phase 1-of-4 encoded transmission is complex. The decoder consists of XNOR gates which detect the transitions on the wires, NAND gates and a SR latch to decode the data back into the single-rail form. The data decoding in a current sensing 1-of-4 encoded interconnect becomes simpler and faster than voltage-mode one because it does not need to detect transitions and compare with the previous transitions. It consists of current comparators and OR gates (section 3.2).

2.4 Completion Detection Techniques

In synchronous interconnects, the role of the clock is to define points in time where signals are stable and valid. In a self-timed communication, the absence of the clock means that there must be another way to detect when signals are stable and valid. In a delay-insensitive channel, the validity of data is encoded within the data by the transmitter and data validity test (completion detection) is performed by the receiver. The validity test is used to determine that the arrived data is a valid value for the chosen delay-insensitive encoding. In practice, it is also necessary to perform data neutrality test. The implementations of validity and neutrality tests play an important role in the efficiency of a delay-insensitive communication channel.

The completion detection for a four-phase dual-rail (1-of-4) encoded channel is carried out by sensing voltage levels on each pair (one 1-of-4 group) of wires. In a two-phase channel sensing voltage transitions of each pair (group)

of wires is required, in this case it requires XOR gates instead of OR gates. Completion detection logic of two-phase dual-rail and 1-of-4 encoded 32-bits channel is shown in Section 4.1 Figure 4.2 and Figure 4.1, respectively. This way of detecting data validity requires logic circuitry whose delay increases drastically when the channel bit width increases, making delay-insensitive interconnects problematic for high performance systems. A fast completion detection technique, where its delay does not increase with transmission bit width, is proposed in Chapter 4 for current sensing interconnects.

2.5 Self-timed Components

In this section, design of self-timed components which are used in the interconnects are discussed briefly. A C-element is a basic building block of self-timed logic. It is a state-holding element, a special kind of latch. When all of its inputs are 0 or 1 the output is set to 0 or 1, respectively. For other input combinations, it preserves its state. Its truth table is shown in Table 2.1 where t and $t - 1$ indicate the current and previous values, respectively. Transistor-level implementation of a C-element is shown in Figure 2.10.

Table 2.1: The truth table of 2-input C-element.

a_t	b_t	c_t
0	0	0
0	1	c_{t-1}
1	0	c_{t-1}
1	1	1

A resettable C-element is a variant of C-element which has a reset input. Its output can be enforced to 0 using the reset input, independently of its other inputs. Its circuit is shown in Figure 2.11. An active-low resettable C-element has been used in one of the interconnects designed in this thesis, see Section 3.2.3. An upper asymmetric C-element is also a variant of C-element where one of its inputs acts like an active-low reset signal. When all inputs are 1 its output is set to 1 and if the input that acts as active-low reset is *low* the output is set to *low* regardless of the other inputs value. For other input combinations, the C-element preserves its state. A 3-input upper asymmetric C-element has been used in the serializer circuit presented in Section 5.2.2. Its CMOS implementation is shown in Figure 2.12.




Figure 2.10: Transistor level implementation of a 2-input C-element




Figure 2.11: Active-low resettable C-element

2.6 On-Chip Signaling Schemes

The signal transmission systems used in CMOS circuits can be broadly classified into two categories: voltage-mode and current-mode signaling. The important difference between these two transmissions systems lies in the type of the transmitted signal. That is, the signal can be transmitted using voltage or current. Several design options for interconnect signaling exists, for example, single-ended or differential signaling, pulse signaling, and wave-pipelining. A designer has to choose the optimal signaling scheme and possibly customize




Figure 2.12: 3-input upper asymmetric C-element

it. To do so, there is a tradeoff among latency, throughput, power and area that should be considered. In this Section, different signaling techniques that have been designed in order to improve the performance of delay-insensitive on-chip interconnects are discussed. The conventional voltage-mode signaling with repeater insertion and pipelining is also discussed since it has been used as a reference case.

2.6.1 Current-Mode and Current Sensing Signaling

The key to current-mode and current sensing signaling is the low-impedance termination at the receiver which results in reduced signal swings without the need of separate voltage references and increased bandwidth performance. Also this low-impedance termination shifts the dominant pole of the system and leads to a smaller time constant and thus less delay. It is typically implemented by terminating the line with a diode connected transistor. This signaling can operate at a much lower noise margin than the voltage-mode network, and at a much lower swing as well due to its immunity to power supply noise. All these translate into increased bandwidth performance [30], decreased delay and reduction in dynamic power dissipation and higher noise immunity. The other important feature of current-mode signaling is its reduced delay sensitivity due to process induced variations [31]. For these reasons, current-mode signaling technique becomes a better alternative than voltage-mode for contemporary and future high-speed noise-prone single-chip systems. Current-mode and current sensing signaling have already been proven to provide drastic speed enhancements for on-chip signaling [32, 33, 34]. It is also shown theoretically in [32] that current sensing signaling can be three times faster than voltage-mode signaling.

There are three sources of power dissipation in current-mode circuits: static, dynamic, and short-circuit power dissipation. In current-mode signaling static power dissipation is the major component of the total power dissipation that arises from the constant current path from VDD to ground via the termination. Static power dissipation can be minimized using different circuit techniques which reduce leakage currents. Dynamic power is dissipated when the parasitic capacitance of the wire is charged and discharged. Since current-mode signaling operates at low voltage swing dynamic power consumption is not as significant source of power dissipation as in voltage-mode signaling. The third source of power dissipation arises from the finite input signal edge rates that result in short-circuit current. Generally, careful control of input edge rates can minimize the short circuit current component to within 20% of the total dynamic power dissipation [35].

Inspired by the advantages explained above, different signaling techniques based on customization of current-mode or current sensing signaling have been designed and used in all of the presented on-chip interconnects (Chapters 3 and 5).

Current-mode and current sensing signaling refers to sensing a signal with a low impedance termination at the receive-end which results in a shift or extension in dominant pole position thereby increasing the bandwidth of the line. The difference between these two is their receiver type. That is, in the current-mode signaling the receiver senses the voltage at the end of the wire, compares it with a reference voltage and then amplifies the result. On the other hand, in the current sensing signaling, the receiver senses the current at the end of the wire, compares it with a reference current and finally outputs the result in voltage levels. The current sensing signaling makes the implementation of a delay-insensitive interconnect circuits simpler, especially the data decoding and completion detection circuits as will be discussed in Chapter 3.

Binary and Multilevel Current Sensing Signaling

In binary current sensing signaling either there is current I through the wires or there is no current. The receiver compares the wire current with a reference current in order to decode out the transmitted data and also to perform the completion detection test. The LEDR encoded current sensing interconnect, presented in Section 3.1, uses a binary current sensing signaling. It uses a diode connected NMOS transistor both as termination load and to mirror the

wire current to a current comparator. The current comparator compares the wire current with a reference current.

Using a current comparator with more than one reference current, it is possible to detect more than one current level in the wire. Multilevel current sensing signaling has been proposed for both synchronous and self-timed on-chip interconnects [36, 37, 38]. Multilevel current sensing signaling is very attractive for delay-insensitive interconnects because it opens up the possibility to represent each code with current level. This simplifies the encoding, decoding and completion detection circuits implementation complexity besides minimizing the delay incurred due to decoding and completion detection.

In a delay-insensitive transmission the data validity indicator is the transmitted data itself. The transmission of every new bit needs to be seen in the wire and detected in the receiver. Since two-phase handshake is preferred for long on-chip interconnects either transition in voltage (in case of current-mode signaling) or different current values can be used as data validity indicator. Using transition in current-mode signaling may cause unnecessary power consumption due to the constant current flow in some of the wires which previously made a transition to a high state. In order to save this power, the interconnect presented in 3.2 allows current flow in the wires only during the respective symbol transmission. If binary current mode signaling is used with this type of power saving transmission scheme, the data validity indicator cannot be seen in the wires when there is consecutive transmission of the same symbol. Thus, in two-phase 1-of-4 encoded current sensing interconnect implementation it becomes possible to differentiate between the consecutive transmission of the same symbols using multilevel currents. The transmitted multilevel current is first detected at the receiver by a detecting circuit based on a current comparator. Then, the encoded voltages are estimated using decoding circuitry.

Differential Multilevel Current Sensing Signaling

Differential current sensing signaling has better noise robustness than single-ended signaling. It has been demonstrated that high speed and energy efficient on-chip communication has been achieved using differential current sensing signaling [39, 40, 41, 42]. In [43] comparisons between differential current sensing signaling and voltage-mode signaling with optimal repeaters insertion have been performed using 250nm, 130nm, 65nm and 45nm technologies. Besides

its superiority in speed for longer wires, differential current sensing signaling consumes less power than optimal repeaters insertion for activity of 50% and higher and length 4mm and longer for 130nm, 65nm and 45nm technologies.

In order to get both noise and delay variations robustness, four wires per bit are required if binary current sensing signaling is used. Two wires per bit for the delay-insensitive encoding and two wires per each encoded wire to support differential signaling. This has a much larger area overhead and higher power dissipation, as it requires four wires per bit transmission. By using current directions and current values simultaneously both the delay-insensitivity and differential signaling has been achieved with only two wires per bit transmission instead of four. This technique has been implemented in the on-chip interconnect presented in Section 3.3. A change in the current level on the wire indicates arrival of new data (delay-insensitivity), while the direction of the current flow reveals the logical value of the transmitted bit. This way of integration leads to more power and area efficient robust communication. The data transfer needs three current levels.

Wave-pipelined Differential Pulse Current-Mode Signaling

In pulse signaling only a small portion of the wire is charged during pulse propagation, significantly reducing the amount of capacitance needed to be charged and hence, saving a considerable amount of power over level-based signaling. It has been shown that the use of pulse signaling can save up to 50% of energy compared to level-based signaling with repeater insertion [88]. Furthermore, it has been demonstrated through analytical models that more than 70% power saving could be achieved by combining pulse signaling with wave-pipelining technique without penalties of data throughput [89].

In [34], a prototype 8Gbps serial link employing pulsed current-mode signaling was manufactured and measured. Sharp current-pulse data transmission was used to modulate transmitter energy to higher frequencies, where the effect of wire inductance is maximized, allowing the on-chip wires to function as transmission lines. In addition to power saving, pulse current-mode signaling mitigates the effect of dispersion due to its return-to-zero signaling scheme in which receiver termination is employed.

The serial link, presented in Chapter 5, employs differential current-mode pulse signaling along with wave-pipelining since this helps to achieve both high-throughput and low-power consumption for global communication. It

has combined pulse dual-rail encoding with wave-pipelined differential pulse current-mode signaling, enabling both delay variation and noise robustness.

2.6.2 Voltage-Mode Signaling: Reference

In voltage-mode signaling the voltage has to swing from rail-to-rail over the entire length of the wire. This leads to large dynamic power consumption, larger delay and it also generates power-supply noise [86]. The optimal repeater insertion technique [48, 106] used in voltage-mode signaling, was developed to reduce the wire delay and improve performance of lengthy global interconnections. However, with the increase in number and density of interconnects, the number of repeaters would increase manifold, presenting significant overhead in terms of power and area. Furthermore, as the optimal repeater insertion distance decreases with each technology node due to increased resistive effects of interconnect, the overall improvement in delay can be undermined by the exponential increase in the number of repeaters and associated driver/repeater power dissipation. A higher throughput can be obtained by using pipeline latches instead of repeaters to both amplify the signal and spread the link delay over multiple pipeline stages. This further increases power consumption and area costs compared to the simple repeater approach. Since most of high performance delay-insensitive links use either voltage-mode signaling with optimal repeater insertion or pipelining [56, 107, 108, 109], both signaling techniques are employed for the delay-insensitive interconnect presented in Chapters 3 and 5. Comparison between the conventionally implemented voltage-mode with repeaters/pipelining latches and current sensing delay-insensitive interconnects helps designers to make appropriate decision on which signaling techniques to use for specific circumstances.

2.7 On-Chip Wire Modeling

A chip is non-functional without wires that connect devices each other. Wires carry signals from one place to another. On-chip wires constitute the lowest level in a hierarchy that spans chip to package-level connections. On-chip wire is not an ideal conductor with zero resistance, capacitance and inductance, but rather it is an unwanted parasitic circuit element. With the increase in circuit performance, complexity, density and levels of integration in nanometer technologies, it is essential to include all parasitic effects during the optimization

process. However, this is not a feasible approach due to the large amount of design variables in the optimization process and the overall complexity of the chip. Furthermore, this approach has the disadvantage of not seeing the exact problem, because at a given circuit node, only few dominant parameters affect the overall performance. Thus, designers need to have a clear insight into the parasitic wiring effects, their relative importance and their reduced-order models. Wire parasitics estimation is required to compare different interconnect schemes because interconnect figures of merits (performance, power consumption and noise coupling) [93, 94] are functions of wire parasitics. In this thesis a wire refers to just the metal that interconnects different blocks and the interconnect refers to a wire with its driver and data encoder, load (receiver input impedance) and receiver along with data decoder and completion detector. This section discusses briefly methods and basis for estimating wire parasitics and the electrical level modeling of wires.

2.7.1 Wire Parasitic Estimation and Extraction

Wire parasitic extraction is usually done by representing complex structures as a collection of simple geometric elements and then each parasitic value is combined using superposition or introducing scale factors to obtain the parasitics of the complex structure. There are many commonly used tools which extract the wire parasitics by assuming that the electromagnetic field through interconnects are quasi-static; they ignore the displacement current in maxwell equations. With such simplification, electrical fields remain static outside conductors, but magnetic fields retain frequency dependency inside conductors so that the skin effect can be accounted properly. Capacitance and conductance of a structure are determined by electrical fields while resistance and inductance are determined by magnetic fields. In other words, by ignoring the displacement current, magnetic and electrical fields are decoupled in the quasi-static theory. Because of this decoupling, a quasi-static field solver is quicker and can solve much bigger problems than a full-wave solver. For example, FastHenry [45] and FastCap [24] are among the quasi-static field solvers.

The interconnects implemented throughout this thesis are assumed having a microstrip configuration. A microstrip is a strip of metal over a return ground plane, as shown in Figure 2.13, where w , h , and d are the wire width, wire height, and wire length, respectively. The t_{ox} is the distance to the underlying




Figure 2.13: Single microstrip wire

ground plane. An electric and magnetic field is created around the microstrip if a driving circuit injects a voltage and current signal, respectively, onto it.

Resistance

The resistance of a wire is the ratio of potential difference of the two ends of a wire to the total current flowing through it:

$$R \equiv \frac{\Phi_{12}}{I} \quad (2.1)$$

where Φ_{12} is the potential difference between the two ends of the wire and I is the current flowing through the wire.

Resistance is dominated by the cross sectional area and the resistivity (inverse of conductivity) of the signal conductor. The DC-resistance, r_{dc} of a microstrip structure, shown in Figure 2.13, is given by:

$$r_{dc} = \frac{\rho d}{h w} = R_{square} \frac{d}{w} \quad (2.2)$$

where w , h , ρ are width, thickness and resistivity of the wire, respectively.

Since the thickness is usually a constant for a given technology, it is customary to incorporate it with the resistivity and form a single constant called sheet resistance of the material (R_{square}). At low signal frequencies, Equation 2.2 is sufficient since the entire cross section of the wire carries the current. As the frequency increases, the current density inside is not uniform, but drops away exponentially with depth into the conductor. This phenomenon is called the skin effect since most of the current is now flowing through the skin of the conductor. This leads to current crowding primarily on the surface and the effective cross-section where current flows reduces. As a consequence, wire resistance increases with the frequency. Skin effect is defined as the depth below the surface of the conductor at which the current density decays to $1/e$

(about 0.37) of the current density at the surface and it is given by:

$$\delta_e = \sqrt{\frac{\rho}{\mu\pi f}} \quad (2.3)$$

Skin effect starts to occur close to the cutoff frequency, f_s where $\delta_e \leq 0.3h$ and is fully developed when $\delta_e \ll h$ (as a guideline $\delta_e \leq 0.1h$) [96]. The obvious and generally accepted term is to get the minimum of width and thickness to obtain the cutoff frequency. For typical on-chip wires, δ_e is found to be equal to $1.5hw/(h+w)$ with relative error less than 5% for $0.25 < h/w < 10$ [95]. There is a widely used empirical formula which describes the frequency dependent behavior of a wire over a ground plane (microstrip structure).

$$R(f) = \begin{cases} r_{dc} & f \leq f_0 \\ r_{dc}\sqrt{\frac{f}{f_0}} & f \geq f_0 \end{cases} \quad (2.4)$$

where

$$f_0 = \frac{\rho}{\mu\pi\delta_e^2}$$

is referred as the break frequency at which this phenomenon begins to dominate. Skin effect decreases the effective cross sectional area that carries the current, which causes resistance to increase. The accurate frequency dependent modeling of wire parameters is usually done considering both resistance and inductance.

Besides the skin effect there are other causes that increase the resistivity of a metal such as metal barrier and surface scattering. The purpose of the barrier is to prevent the diffusion of copper into the surrounding dielectric. Since it is fabricated from a higher resistivity metal, it is safe to assume that the copper carries all the current but the effective area through which current conducts reduces. Surface scattering has a significant contribution to the resistivity when the minimum dimension of the metal line becomes comparable to the mean free path of the electrons. Furthermore, temperature also affects the resistivity of the copper, that is, its resistivity increases with temperature. In all of the interconnects designed in this thesis, the resistance values of the wire were extracted using FastHenry [45].

Capacitance

When two conducting objects are charged to different electric potentials, an electric field is created between them and a capacitance arises. It always takes some time to build up a voltage between two objects. The capacitance can be seen as the reluctance of voltage to instantaneously increase or decrease in response to an input signal. The capacitance for the single isolated microstrip wire shown in Figure 2.13, can be approximated by:

$$C = C_{parallel} + C_{fringe} = \frac{w\epsilon_{ox}}{t_{ox}}d + \frac{2\pi\epsilon_{ox}}{\ln(2 + 4t_{ox}/h)}d \quad (2.5)$$

where $C_{parallel}$ is the parallel-plate (bottom area-to-substrate) capacitance, C_{fringe} is the fringing (side-wall-to-substrate) capacitance, and ϵ_{ox} is the insulator dielectric constant. This simplification is only useful for estimating rough capacitance values. In reality, a wire is surrounded by a large number of other wires on the same layer and adjacent layers in case of the multilevel structure. Each wire is coupled not only to the grounded substrate, but also to neighboring wires. To model the capacitance in such a complex environment is a non-trivial task and the above equation is not a good model for the capacitance of a wire in such a complicated structure. In practice, field solver extraction tools are utilized to numerically calculate the parasitic capacitance values. Hence, capacitance values of the wires are extracted using a Linpar [46] field solver for all interconnects designed in this thesis.

Inductance

Inductance is a measure of the distribution of the magnetic field near and inside a current-carrying conductor. This measure is a property of physical layout of the conductor and is also a measure of the ability of the conductor to link magnetic flux or to store magnetic energy. The fundamental equation for inductance is as follows:

$$L = \frac{\oint \vec{B} \cdot d\vec{A}}{I} \quad (2.6)$$

where I is the current, B the magnetic field induced from I , and A is the integration loop. The definition of inductance follows a loop property, the current return path should be known to determine the inductance value. In contemporary interconnect structures the return current is spread all over the

range and the exact return path of a current is not known. In these cases, the possible current return paths are the power distribution network and the adjacent wires [97]. The loop formed by the wire and its 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 a particular wire, but also on several thousands of other wires. Therefore, in order to find the inductance, the induced current is assumed to return at the infinity. This method was first proposed in [98] and was further introduced for circuit analysis in [99].

A simple approach for inductive parasitic extraction is to use a free space relationship, which relates loop inductance (L) of a wire to its capacitance ($C_{\epsilon_r=1}$) by assuming no dielectric in the medium, the inductance is given by:

$$L = \frac{\epsilon_0 \mu_0}{C_{\epsilon_r=1}} \quad (2.7)$$

This formulation is used in the tool Raphael RC2 [100], which is a two dimensional parasitic extraction tool. Considering the middle conductor in a three parallel conductor system, the self and mutual inductance equations become:

$$L_S = \frac{\epsilon_0 \mu_0}{2} \left(\frac{1}{C_S} + \frac{1}{C_S + 2C_c} \right) \quad (2.8)$$

$$L_m = \frac{\epsilon_0 \mu_0}{2} \left(\frac{1}{C_S} - \frac{1}{C_S + 2C_c} \right) \quad (2.9)$$

where C_s and C_c are self and coupling capacitances, respectively. Unfortunately in an IC, this assumption does not hold and more detailed methods need to be used. For a wire with a finite conductivity, the magnetic flux exists both inside and outside the conductor, subdividing the wire inductance into internal and external components. The internal inductance of a wire is due to the magnetic flux inside the wire and the external inductance is due to a magnetic flux outside the wire (loop or partial inductance is external to the wire). When modeling the internal inductance, high frequency effect of the current distribution has to be considered because of the skin effect. The current distribution inside a conductor also changes with frequency due to the proximity effect, the current tends to concentrate closer to the current return path in or-

der to minimize the inductance. Another effect of frequency on the inductance is due to multi-path current re-distribution. In an integrated circuit, there are many possible current return paths, e.g., the power/ground network, nearby signal lines, and the substrate. The distribution of the return current among these possible paths is determined by the impedance of the individual paths. At different frequencies, the relationship among the impedances of different paths will change, as well as the distribution of the return current. The return current is distributed in a way that the total impedance is minimized at a specific frequency. If the frequency dependent effects are very important to consider in a desired frequency range, the cross-sections are subdivided into sections smaller than the skin depth at the maximum frequency of interest. Then, the current distribution in each filament can be regarded as uniform. To calculate the partial inductances of rectangular cross-sectional wires, closed-form equations proposed in [98] are used. In this manner, an inductively coupled RL circuit can be formed for the conductor. By solving currents in this circuit at several points in the frequency domain, the frequency dependent resistance and inductance can be obtained [101]. This technique, which is known as partial element equivalent circuit (PEEC) is the foundation for frequency dependent parasitic extraction tools such as FastHenry [45]. For all the interconnects presented in this thesis, the inductance values of the wires were extracted using FastHenry.

2.7.2 Electrical Level Wire Modeling

In order to analyze the performance and signal integrity of an interconnect, it is necessary to translate the wire layout and technology information such as the width and length of the wire, neighboring line conditions and related dielectrics into electrical parameters. Then these parameters can be combined with other circuit components to evaluate performance. This is achieved through parasitics extraction. Based on the design and technology specifications, a physical line is usually converted into a netlist composed of resistors, capacitors and inductors (if necessary). Due to the technology scaling and increasing operating speeds, accurate modeling of wires has become a necessity. Wires have traditionally been modeled as lumped RC segments but with circuit operation frequency on the rise, this model lacks the required accuracy to model a high-performance interconnect.

The fundamental electrical behavior of a metal wire can be fully determined

using Maxwell's equations (Equation 2.10 - 2.14) in conjunction with the rule of charge conservation (Equation 2.14).

$$\nabla \cdot D = \rho \quad (2.10)$$

$$\nabla \times E = \frac{\partial B}{\partial t} \quad (2.11)$$

$$\nabla \cdot B = 0 \quad (2.12)$$

$$\nabla \times H = J + \frac{\partial D}{\partial t} \quad (2.13)$$

$$\nabla \cdot J + \frac{\partial \rho}{\partial t} = 0 \quad (2.14)$$

Since solving these equations requires a huge amount of computation, they are usually simplified depending on the range of frequencies and wire lengths of interest. As already discussed in the above subsection, the behavior of a wire is frequency dependent. At DC it behaves as a resistor, causing both losses in the voltage supply (IR drop) and static power consumption (IR^2). Wire activities are also affected by the interaction between electric and magnetic fields when operating in AC range. In current IC designs, quasi-static assumption is usually applicable since the signal frequency is relatively low and the wire length is much shorter than the wavelength of the signal. For instance at 10GHz, the wavelength is about 17cm for $k=3.0$ dielectrics. As explained in Section 2.7.1, under the quasi-static assumption the electric and magnetic field can be decoupled. Thus, wire capacitance and inductance can be defined and extracted independently and the resulting wire is represented by an RC or RLC equivalent circuit. To solve the electrical response, the wire is assumed to be uniform, and therefore the Maxwell's equations can be reduced to the telegraph equation (transmission line theory) as follows:

$$\frac{\partial^2 V}{\partial x^2} = RC \frac{\partial V}{\partial t} + LC \frac{\partial^2 V}{\partial t^2} \quad (2.15)$$

where x is the length dimension, t is the time and V is the voltage. A dimensionless ratio of the physical length of a wire to the signal wavelength, $\frac{\ell}{\lambda}$, is referred as the electrical length. This ratio is used to determine whether to model the wire using a lumped or distributed model. A wire is considered to be electrically short if the electrical length is less than unity. These electrically short wires belong to the classical circuit analysis and it is quite safe to approximate the entire line as a lumped RC or RLC segment because the

signal level along the entire length of the wire is almost constant. A rule of thumb to determine whether a wire can be represented by a lumped circuit or not is to test its length against the following criterion [102]:

$$length \leq \lambda/20 \tag{2.16}$$

where λ is the signal wavelength. Since the frequency spectrum that a digital signal contains is more closely related to its rise time ($\frac{1}{3.14t_r}$) than to the signal frequency itself, λ should be estimated from the rise time of the signal. Transmission line modeling needs to be applied when the time of flight (time required for a signal to travel round trip from the driver to the end of a line) across the wire becomes comparable to the signal rise time. A transmission line can be thought as a large number of lumped segments in series so that they represent the distributed nature of the wire.

The importance of modeling inductive effects in wires is increasing because of faster rise times and longer wires. Wide wires used in upper metal layers can be especially susceptible to inductive effects due to their low resistance [91]. Wires should be modeled as RLC lines if they satisfy the following two conditions [91, 102]: input signal rise time is smaller than the time of flight and the time of flight is greater than the Elmore delay of an RC line. The latter criteria describes a situation where wire resistance is considerably smaller than line characteristic impedance. By combining these two criteria, the following condition is obtained:

$$\frac{t_r}{2\sqrt{LC}} < length < \frac{2}{R}\sqrt{\frac{L}{C}} \tag{2.17}$$

where R, L, and C are the resistance, inductance and capacitance per unit length, respectively. In case the constraint on the left-hand side of Equation 2.17 is larger than the right-hand side, $t_r > 4L/R$, the input signal is not fast enough and the inductance effect can be ignored regardless of the wire length.

Since the interconnects designed in this thesis are targeted for high-performance signaling over global wires, all wires are modeled using a distributed RLC model by considering the inductance effect, as shown in Fig.2.14. In order to accurately consider crosstalk noise effects, both capacitive and inductive coupling between all wires was also included.




Figure 2.14: Distributed RLC wire model with coupling

2.8 Chapter Summary

In this chapter self-timed delay-insensitive communication techniques, high speed on-chip signaling schemes and wire modeling has been discussed. These are the foundation topics for the next chapters. The techniques presented in this chapter are the enabling factor to achieve delay-insensitivity, higher performance and lower power consumption on-chip communication and they have been implemented in the interconnects designed during the period of this PhD work.

Chapter 3

Design of Delay-Insensitive Current Sensing Interconnects

Unlike synchronous design style which uses a globally distributed clock signal to indicate moments of stability of the data, asynchronous circuits exchange information using handshakes to explicitly indicate the validity and acceptance of data. Depending on the type of handshaking, data encoding, channel type, and data-validity schemes there are a number of alternative communication protocols. As already discussed in Chapter 2, two-phase handshaking is preferred for global on-chip communication since it reduces the number of transitions and avoids the requirement of a spacer between consecutive data symbols. This saves communication time and energy of the system. The most common asynchronous data encoding in GALS design is bundled-data (single-rail) encoding which uses N lines to represent N -bit information and two additional handshake lines indicating data validity and acceptance. Since this encoding has a timing constraint between control (data validity) and data lines, communication through a long on-chip interconnect becomes sensitive to delay variations. Therefore, converting bundled-data encoding to delay variation insensitive encoding is necessary for global on-chip interconnects where delay variations are unavoidable. The general block diagram of conversion between bundled-data and delay-insensitive encoding is shown in Figure 3.1. The conversion between the two encodings requires a data encoder at the transmitter side and a data decoder as well as a completion detector at the

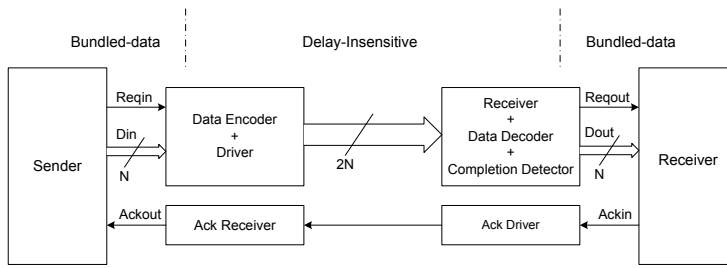


Figure 3.1: Bundled-data \leftrightarrow delay-insensitive conversion

receiver side.

In this chapter design and analysis of three delay-insensitive current sensing on-chip interconnects are presented. Their performance and power consumption are analyzed and compared with conventional delay-insensitive on-chip interconnects. The design of these interconnects is among the author’s contributions. The design and analysis of each interconnect are presented in separate sections. The performance, energy and area of these three interconnects are compared to each other using the same technology and wiring models in Chapter 6.

This chapter is organized as follows. In Section 3.1, design of an on-chip interconnect which uses LEDR encoding and current sensing signaling is presented. Analysis of its performance and power consumption along with two dual-rail encoded reference interconnects are also discussed. The design and simulation results of a 1-of-4 encoded multilevel current sensing interconnect is presented in Section 3.2. Its performance and power efficiency has been compared with two 1-of-4 encoded voltage-mode interconnects. In Section 3.3, area and power efficient two-phase dual-rail encoded differential current sensing interconnect is presented. The summary of this Chapter is presented in the last section.

3.1 Level-Encoded Dual-Rail Current Sensing Interconnect

LEDR encoding is among the preferred encoding schemes for global on-chip communication, because it needs no resetting transitions that consume time and power. Its completion detection and decoding circuitry is faster and much simpler than that of two-phase dual-rail encoding since detection is level based

rather than transition based. The conventional two-phase dual-rail protocol has more complex and slower decoding and completion detection circuitry compared to LEDR. In the two-phase protocol, if the transmitted data has the value 0 there is a transition on one wire and a transition on the other wire if 1 is transmitted. To detect completion and decode the data, the current and previous state on both wires need to be detected. This makes the circuit relatively complex and slow. The gate level implementations of the encoder, decoder and completion detector of two-phase dual-rail encoded transmission are shown in Figure 3.2 and 3.3.




Figure 3.2: Conventional two-phase dual-rail encoder.




Figure 3.3: Conventional two-phase dual-rail decoder and completion detector.

In LEDR one data bit is encoded into a 2-bit codeword as follows. A data sequence $D(i)$ of bits is encoded into a sequence $DS(i)$ and $DP(i)$ of state and

phase bits, respectively.

$$DS(i) = D(i), \quad \forall i$$

Given $DP(0)$,

if $(DS(i + 1) = DS(i))$ **then**

$$DP(i + 1) = \neg DP(i),$$

else

$$DP(i + 1) = DP(i)$$

As each codeword has a phase, even or odd, it is possible to differentiate between two consecutive same bit transmissions. Figure 3.4 shows the four possible codewords organized into overlapping groups of value and phase. The arrows illustrate the allowed transitions between the codewords.




Figure 3.4: Transitions between LEDR code.

3.1.1 Data Encoder and Driver

The encoder takes the request and data bit in the voltage-mode bundled-data form and converts this information into current-mode LEDR signaling. The conversion from the two-phase voltage mode to the LEDR current mode is shown in Figure 3.5 at the protocol level. As shown in Figure 3.6, the outputs of the double-edge-triggered flipflops 1 and 2 ($DFF1$ and $DFF2$) control the current flow through the phase and state wires (DP and DS) by serving as

gate voltages of the transistors $Mn2$ and $Mn4$. Considering the data phase wire $W(DP)$, the transistors $Mp1$ and $Mn1$ generate the source current. The transistor $Mp2$ is used to mirror the generated current I from the current source and drive this current through $W(DP)$ when $Mn2$ is 'ON'. The same principle applies to the state wire $W(DS)$.




Figure 3.5: Protocol conversion.

3.1.2 Receiver, Decoder and Completion Detector

At the receiver side, the current comparator circuit, as depicted in Figure 3.6, is composed of the diode-connected input NMOS transistor $Mn6$, the NMOS transistor $Mn7$ connected to replicate this input current, the threshold current generating pair of transistors $Mn5$ and $Mp3$, and the PMOS transistor $Mp4$ that replicates the threshold current. In addition to serving as an input transistor, $Mn6$ acts also as a termination load. The drains of the PMOS replicating transistor $Mp4$ and NMOS replicating transistor $Mn7$ are connected to generate the comparator circuit's output voltage $V(DP)$. The comparator provides a logical *high* output voltage when the input current $I(DP)$ is less than the threshold current and a logical *low* output voltage when the input current $I(DP)$ is greater than the threshold current.

As shown in Figure 3.6, the decoder takes as input the output of the state wire's current comparator, $V(DS)$, and reconstructs the data sent by the transmitter. Unlike conventional two-phase transmission which needs to

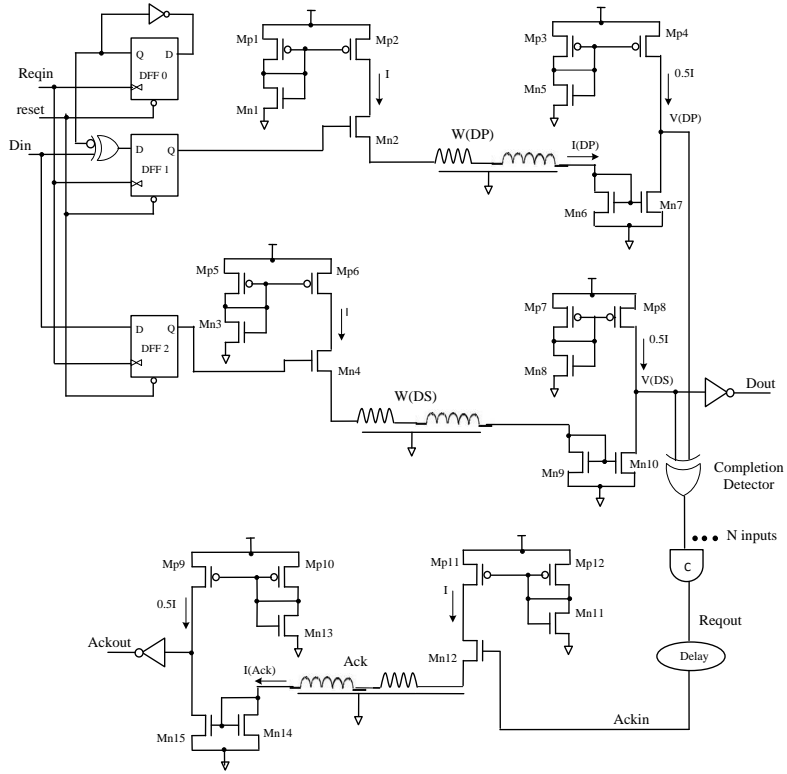


Figure 3.6: LEDR encoded current sensing on-chip interconnect.

detect both wires' current and previous states (Figure ??), here the decoder needs sensing voltage level of the state ($W(DS)$) wire using only an inverter. The completion detection is carried out using a 2-input XOR gate, the outputs of the two current comparators are the inputs for the completion detector. The completion detection circuit is also simpler and faster, just one XOR gate per each LEDR group is needed. For N bit transmission, completion detection is carried out using N 2-input XOR gates connected to an N -input C-element. The output of the C-element acts as the bundled-data request signal ($Reqout$) passed to the receiving module.

3.1.3 Acknowledgment Transmission

The acknowledgment signal transmission circuitry is also shown in Figure 3.6. The voltage-mode bundled-data acknowledge signal ($Ackin$), sent by the receiving module, is converted into a current-mode signal Ack during transmission and back into a voltage-mode signal ($Ackout$) at the transmitter side.

As can be seen from Figure 3.6, the signaling circuits are basically equivalent to the ones used for data transmission. When *Ackin* is *high*, current I flows through the *Ack* wire causing an up-going transition on the signal *Ackout* at the transmitter side. When *Ackin* goes *low*, the current is switched off and *Ackout* eventually returns to zero.

3.1.4 Simulation Results and Analysis

Latency, throughput and average total power consumption are considered as main parameters to evaluate the presented LEDR on-chip interconnect (*LEDRCm*). Also the performance and power consumption of two reference interconnects are analyzed. One of the reference interconnects uses LEDR encoding along with voltage-mode signaling with repeaters (*LEDRVm*). The other one uses two-phase dual-rail encoding and voltage-mode signaling with repeaters (*TPDRVm*). This helps to determine the performance improvement and power overhead due to the use of current-mode signaling along with LEDR encoding over *LEDRVm* and *TPDRVm*. During simulations a transmission line model of the wires was assumed by using 20 distributed RLC sections. Metal 4 of a 130nm CMOS technology with minimum metal width, spacing and pitch was used to model the transmission line. The resistance and inductance matrices of the interconnect structure were extracted using FastHenry [45], while the capacitance matrices were extracted using Linpar [46]. The interconnect circuitry was designed and simulated using Cadence Analog Spectre with 130nm CMOS technology from STMicroelectronics. The supply voltage was 1.2V.

Here forward latency is defined as the delay from a transition on the bundled-data request signal (*Reqin*) at the transmitter side to the corresponding transition on the bundled-data request signal (*Reqout*) at the receiver side (see Figure 3.6). Reverse latency is defined as the delay from a transition on the bundled-data acknowledgment signal (*Ackin*) at the receiver side to the corresponding transition on the bundled-data acknowledgment signal (*Ackout*) at the sender side. The change in forward and reverse latency when the wire length is varied from 1 to 11mm is shown in Figures 3.7 and 3.8 for *LEDRCm*, *LEDRVm* and *TPDRVm* interconnects. The *LEDRCm* interconnect latencies are much smaller than *LEDRVm* and *TPDRVm* interconnects for longer wires. For example, at 7mm communication distance, the forward latency of *LEDRCm* is about two-thirds of *LEDRVm* and half of the *TPDRVm* latencies.

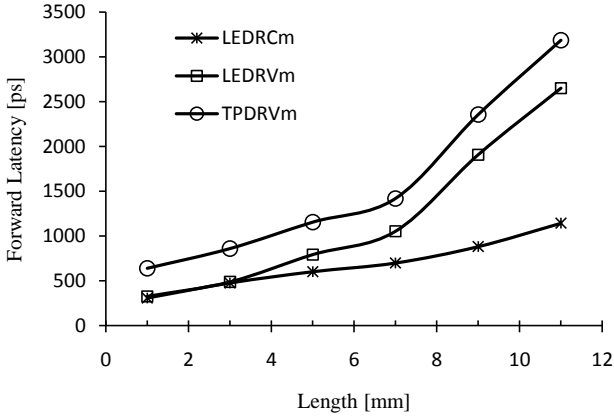


Figure 3.7: Forward latency of *LEDRCm*, *LEDRVm* and *TPDRVm*.

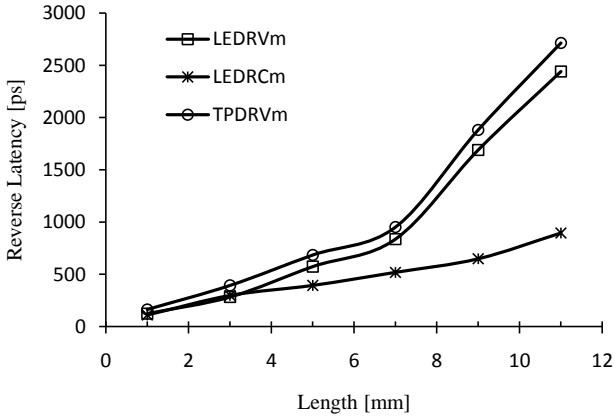


Figure 3.8: Backward latency of *LEDRCm*, *LEDRVm* and *TPDRVm*.

The forward latency of *TPDRVm* interconnect is higher than the two LEDR encoded interconnects, showing the impact of its complex encoding/decoding and completion detection logics. The latency difference between *LEDRCm* and *LEDRVm* interconnects shows the use of current sensing signaling in enhancing the performance of delay-insensitive interconnects especially at global wire lengths. The forward latency of 9mm long *LEDRCm* interconnect is only 46% of that of the *LEDRVm* interconnect's latency with the same communication distance.

As in the latency, the *LEDRCm* interconnect throughput is higher than that of the *LEDRVm* and *TPDRVm* interconnects. At 7mm long communi-




Figure 3.9: Throughput of *LEDRCm*, *LEDRVm* and *TPDRVm*.

cation distance, throughput of *LEDRCm* interconnect is 1.55 and 1.94 times higher than that of *LEDRVm* and *TPDRVm*, respectively. Its throughput is not dropping as fast as the other two interconnects with the increase in the communication distance, this can be seen from Figure 3.9. The *LEDRCm* interconnect achieves 1.005Gbps throughput per one dual-rail group (two data transmission wires + 1 acknowledgment wire) at 5mm wire length without using repeaters or pipelining. The main reason for the higher throughput is the use of current sensing signaling. If a number of these dual-rail groups concatenate in parallel, throughput increases linearly albeit the completion detection (data arrival and stability check) for all groups deviate the linear increase slightly (see Chapter 4).

The average total power consumption of the one-bit *LEDRCm* interconnect varies from $195\mu\text{W}$ to $1242\mu\text{W}$ when wire length is varied from 1 to 11mm. Its power consumption is higher than that of the two reference interconnects especially with long wires as shown in Figure 3.10. However, it is more power efficient for longer communication distances as can be seen from the energy per bit diagrams in Figure 3.11. At global wire lengths, that is, starting from 5mm, it dissipates least energy per bit compared to the other two, and *TPDRVm* dissipates the most.

To summarize the throughput improvement and energy savings of *LEDRCm* interconnect over the reference interconnects (*LEDRVm* and *TPDRVm*), Tables 3.1 and 3.2 are presented. Table 3.1 shows the advantage of current sensing signaling by comparing *LEDRCm* and *LEDRVm* interconnects. The

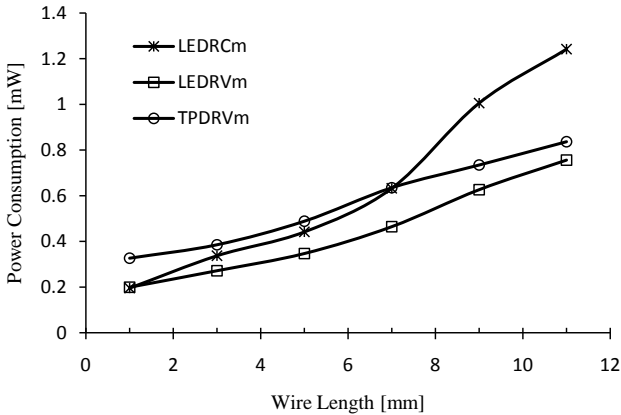


Figure 3.10: Power consumption of *LEDRCm*, *LEDRVm* and *TPDRVm*.




Figure 3.11: Energy per bit dissipation of *LEDRCm*, *LEDRVm* and *TPDRVm*.

benefit of LEDR encoding along with current sensing signaling over conventional two-phase dual-rail encoded voltage-mode interconnect is demonstrated in Table 3.2. *LEDRCm* interconnect gains almost double throughput and 50% energy savings compared to *TPDRVm*.

The simulation waveforms of the one-bit *LEDRCm* interconnect are shown in Figure 3.12. As can be seen, there is a change in current only in one wire per data transfer either on $W(DS)$ or $W(DP)$.

3.1.5 Effect of Crosstalk on Timing

Since *LEDRCm* is a delay-insensitive interconnect, crosstalk induced signal propagation delay variations can cause only performance penalty, do not affect its reliability. In this subsection the performance penalty caused by crosstalk

Table 3.1: Comparing *LEDRCm* and *LEDRVm* interconnects.

Wire Length [mm]	Gain in Throughput	Energy Savings
5	1.37X	7.24%
7	1.82X	25.36%
9	2.34X	31.65%
11	2.5X	34.28%

Table 3.2: Comparing *LEDRCm* and *TPDRVm* Interconnects.

Wire Length [mm]	Gain in Throughput	Energy Savings
5	1.85X	51.14%
7	1.95X	48.91%
9	2.76X	50.51%
11	2.89X	48.78%




Figure 3.12: Simulation waveforms of *LEDRCm* interconnect

is examined. In this analysis, 4-bit parallel data transmission is considered. This requires 8 parallel physical wires since delay-insensitive encoding is used. The wires are modeled as transmission lines with both capacitive and inductive coupling between each other. Minimum wire separation distance of 210nm is used with minimum global pitch specified in 130nm technology and 1.2V supply voltage. The worst-case switching pattern was defined by assuming that capacitive coupling dominates inductive coupling, which is the most usual case in on-chip parallel wires. The effect of crosstalk on performance of *LEDRCm* interconnect is compared with bundled-data voltage mode (*BundledVm*) interconnect. Because the worst-case switching pattern of *LEDRCm* is the same as *LEDRCm*. Furthermore, from crosstalk effect on timing perspective, bundled-data encoded interconnect can represent synchronous transmission. The delay variation percentage of the *LEDRCm* interconnect due to worst-case crosstalk is less than one-third of that of a bundled-data voltage mode (*BundledVm*) one, as shown in Table 3.3. This is because in LEDR encoded data transmission only the state wires ($W(DS)$) make transitions when there is a switching in the input data. The phase wires ($W(DP)$) are quiet. At most the victim wire has one nearest aggressor.

Table 3.3: Effect of crosstalk in *LEDRCm* and *BundledVm*.

Interconnect	Worst-Case Switching	% Delay Variations
BundledVm	$\uparrow\uparrow\downarrow\uparrow - - - -$	+141
LEDRCm	$- \uparrow - \uparrow - \downarrow - \uparrow$	+42

3.2 1-of-4 Encoded Current Sensing Interconnect

In 1-of-4 data encoding, a group of four wires is used to transmit two bits of information per symbol. A symbol is one of the two-bit codes 00, 01, 10, or 11 and it is transmitted through activity on one of the four wires. Since it is possible to detect the arrival of each symbol at the receiver, 1-of-4 encoding is delay-insensitive, as are all the 1-of-N codes [47]. Besides being delay-insensitive, 1-of-4 encoding has more immunity against crosstalk effects when compared to bundled-data encoding, because the likelihood of two adjacent wires switching at the same time is one-eighth times smaller. Furthermore, dynamic power consumption due to wire capacitance is smaller for the 1-of-4 code than for the simpler 1-of-2 (dual-rail) code. This is because the 1-of-4

code conveys two bits of information using only a single transition, while the 1-of-2 code requires two transitions for two bits of information.


In this section, implementation of a novel high-performance link based on multilevel current sensing signaling and delay-insensitive two-phase 1-of-4 encoding is presented. Current sensing signaling reduces communication latency of global wires significantly compared to voltage-mode signaling, making it possible to achieve high throughput without pipelining and/or using repeaters. Performance of the proposed multilevel current-mode interconnect is analyzed and compared with two reference voltage-mode interconnects.

The 1-of-4 Encoded Multilevel Current Sensing (*PMCM*) scheme converts two-phase bundled-data voltage-mode signaling into pulsed 1-of-4 multilevel current sensing signaling at the transmitter side. At the receiver side, delay-insensitive current sensing signaling is turned back into bundled-data voltage-mode communication. The *PMCM* scheme is logically equivalent to a 1-of-4 encoded voltage-mode scheme, the difference is that information is presented as current pulse rather than voltage transitions, as shown in Table 3.4. Hence, one of the four data wires draws current to indicate the presence of a new two-bit data symbol. Similarly, an acknowledgment is signaled as current on the acknowledgment wire. As explained in Chapter 2, such a current sensing implementation is inherently much faster and more immune against power supply noise and delay variations compared to a voltage-mode implementation. The communication protocol is shown in Figure 3.13 (from the receiver's perspective) and the signaling circuits are depicted in Figures 3.14 and 3.15. The advantage of this interconnect implementation is that high throughput and low latency can be achieved without using area and power hungry pipelining or repeaters.

The multilevel and pulsed nature of the *PMCM* scheme can be seen in Figure 3.13. The current detected at the receiver has three different values: 0, I_1 and I_2 . The values I_1 and I_2 are used when the voltage-mode request signal *Reqin* at the transmitter side is *low* and *high*, respectively, reflecting the adopted two-phase communication protocol. The value 0, in turn, means that there is no symbol on a wire. It is used as the initial value of the data wires and for switching off current on a wire when the 2-bit symbol to be transmitted changes, making current on a wire pulse shaped. This feature reduces the overall power consumption of the current-mode interconnect. The values of I_1 and I_2 are determined by considering the speed, power consumption, and

Table 3.4: Encoding and wire current of *PMCM*.

Bundled-data		PMCM			
D_1D_0	Reqin	$I(WQ_3)$	$I(WQ_2)$	$I(WQ_1)$	$I(WQ_0)$
00	$0 \rightarrow 1$	0	0	0	I_2
	$1 \rightarrow 0$	0	0	0	I_1
01	$0 \rightarrow 1$	0	0	I_2	0
	$1 \rightarrow 0$	0	0	I_1	0
10	$0 \rightarrow 1$	0	I_2	0	0
	$1 \rightarrow 0$	0	I_1	0	0
11	$0 \rightarrow 1$	I_2	0	0	0
	$1 \rightarrow 0$	I_1	0	0	0


 Figure 3.13: Communication protocol of *PMCM*.

noise margin of the interconnect. In the following consecutive sections, the implementations of the encoder, decoder and completion detector are separately discussed.

3.2.1 Encoder and Driver

The encoder takes the request and two data bits in the voltage-mode bundled-data form and converts this information into multilevel current sensing 1-of-4 signaling. The double-edge triggered flip-flops shown in Figure 3.14 are used to sample the value of the 2-bit data symbol at each transition of the two-phase request signal *Reqin*. For instance, consider the encoder circuit of the

wire $Q3$. Depending on the value of the signal $Reqin$, either transistor $Mn1$ or $Mn2$ conducts making either current I_1 or I_2 to flow through the wire $Q3$ when the symbol 11 has arrived from the sender module. To prevent the line from drawing current continuously, the transistor $Mn4$ is used to ground the line when other than the symbol 11 is sent. The reset signal rst is controlled by the transmitting module. When a data burst is about to begin, rst is set to *high* enabling the sampling flip-flops. When the burst has been completed, rst is initialized back to *low*, meaning that all the data wires become grounded. This is necessary to prevent the data wires of the link from drawing current (consuming power) during possibly long idle periods between bursts. In nanometer scale technologies process variation effects are one of the major concerns. The driver output currents may vary from their expected values due to process variation effects. In order to minimize this variation, transistors $Mp1$ and $Mp2$ which operate in the linear region form a resistive path from the supply voltage to $Mn1$ and $Mn2$ which in turn keeps the switching threshold of $Mn1$ and $Mn2$ transistors constant.

3.2.2 Receiver

At the receiver side, consider the current comparator circuit of $Q3$, as depicted in Figure 3.15. It is composed of the diode-connected input NMOS transistor $Mn2$, the NMOS transistors $Mn3$ and $Mn4$ connected to replicate this input current, the reference or threshold current generating pair of transistors $Mn1$ and $Mp1$, and the PMOS transistors $Mp2$ and $Mp3$ that replicate the threshold current. In addition to serving as an input transistor, $Mn2$ acts also as a termination load. The drains of the PMOS reference current replicating transistors and line current replicating NMOS transistors are connected together to generate the comparator circuit's output voltages, $V(30)$ and $V(31)$. This comparator provides a logical high output voltage when its input current $I(Q3)$ is less than the threshold current and a logical low output voltage when the input current $I(Q3)$ is greater than the threshold current. Here the current comparator compares current on the wire $Q3$ with two different threshold currents, I_{ref1} and I_{ref2} , in order to distinguish the three current levels. To

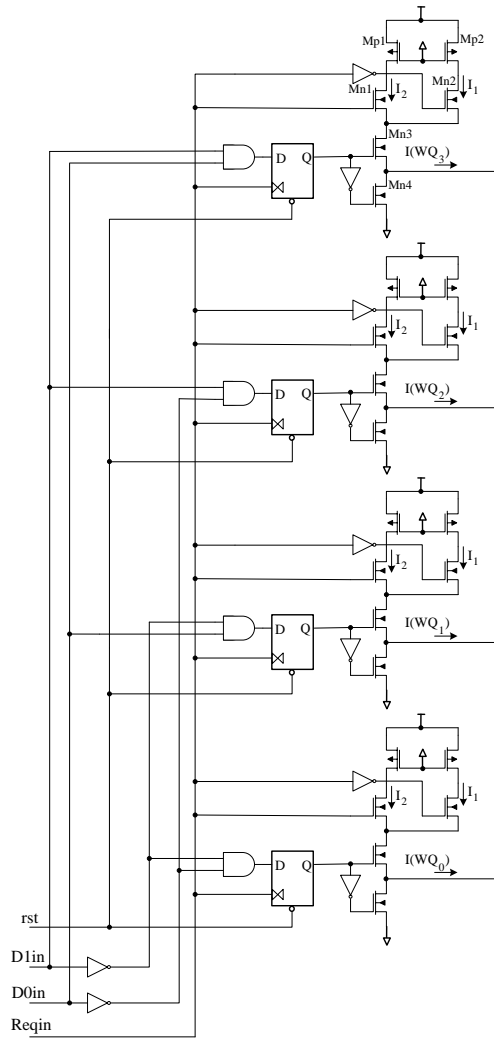


Figure 3.14: Encoder and driver of $PMCM$.

be more specific,

$$\begin{aligned}
 & \text{If } (I(Q3) < I_{ref1}) \text{ then} \\
 & \quad (V(30) = 1) \wedge (V(31) = 1) \quad //(\text{initial state}) \\
 & \text{If } (I_{ref1} < I(Q3) < I_{ref2}) \text{ then} \\
 & \quad (V(30) = 0) \wedge (V(31) = 1) \\
 & \text{If } (I(Q3) > I_{ref2}) \text{ then} \\
 & \quad (V(30) = 0) \wedge (V(31) = 0)
 \end{aligned}$$

In nanometer scale technologies, the line and reference currents at the input of the receiver may vary from the nominal value due to supply voltage, process and temperature variation effects. In Chapter 7 different techniques are developed to ensure reliability of the communication by restoring the current levels within the desired margins at power start-up.

3.2.3 Decoder and Completion Detector

As shown in Figure 3.15, the data decoder, composed of three inverters and two OR gates, needs as inputs the outputs of the current comparators of the wires $Q3$, $Q2$, and $Q1$ to reconstruct the two bits ($D1out$, $D0out$) sent from the transmitter module. Only the comparator outputs of the threshold current I_{ref1} (i.e., $V(10)$, $V(20)$, and $V(30)$) are needed for this purpose. Formally, the logic is as follows:

$$\begin{aligned}
 & \textit{If} (V(30) = 0) \wedge (V(20) = 1) \wedge (V(10) = 1) \textit{ then} \\
 & \quad (D1out = 1) \wedge (D0out = 1) \\
 & \textit{If} (V(30) = 1) \wedge (V(20) = 0) \wedge (V(10) = 1) \textit{ then} \\
 & \quad (D1out = 1) \wedge (D0out = 0) \\
 & \textit{If} (V(30) = 1) \wedge (V(20) = 1) \wedge (V(10) = 0) \textit{ then} \\
 & \quad (D1out = 0) \wedge (D0out = 1) \\
 & \textit{If} (V(30) = 1) \wedge (V(20) = 1) \wedge (V(10) = 1) \textit{ then} \\
 & \quad (D1out = 0) \wedge (D0out = 0)
 \end{aligned}$$

The completion detector reads all current comparator outputs as illustrated in Figure 3.15. For each 4-wire block, the completion detection circuit includes two 4-input NAND gates ($N0$ and $N1$), a 2-input NAND gate ($N2$), and a resettable 2-input C-element ($C1$). To produce the receiver-side request signal $Reqout$, the completion signals of the $N/2$ 4-wire blocks are combined with an $N/2$ -input C-element, where N is the bit-width of the transmitted data. The completion detection process is started by sensing the current values on the four wires. In this pulsed implementation of 1-of-4 encoding, current flows only in one of the four wires. Current through the wire becomes I_1 or I_2 when the transmitter-side request signal $Reqin$ is *low* or *high*, respectively. Hence, if the input current of the comparator is greater than the threshold I_{ref2} , then the output of the C-element $C1$ and subsequently the receiver-

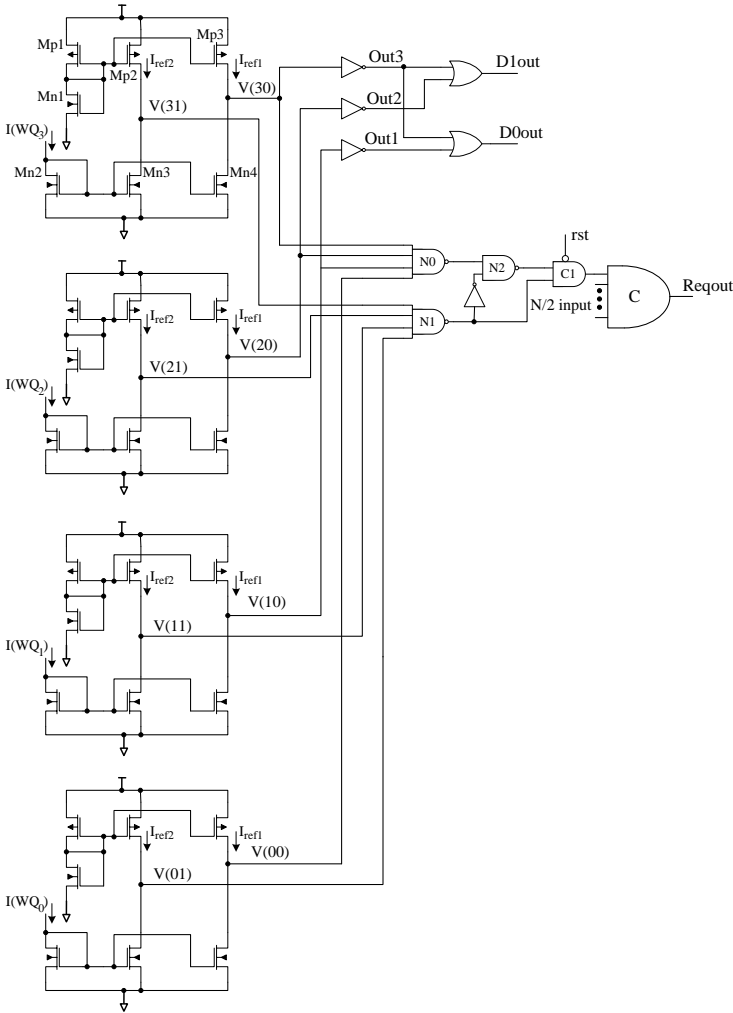


Figure 3.15: Decoder and completion detector of $PMCM$.

side request signal $Reqout$ go *high*. Correspondingly, if the comparator input current is between the thresholds I_{ref1} and I_{ref2} , the output of $C1$ and the signal $Reqout$ go *low*. The completion detection logic uses as inputs the current comparator outputs $V(30)$ and $V(31)$ of $Q3$, $V(20)$ and $V(21)$ of $Q2$, $V(10)$ and $V(11)$ of $Q1$, and $V(00)$ and $V(01)$ of $Q0$. For instance, consider again the receipt of the symbol 11 through the wire $Q3$. Assuming that the transmitter-side request signal $Reqin$ is *high*, the current on the wire $Q3$ is I_2 . Consequently, the comparator outputs $V(30)$ and $V(31)$ become *low*, and all the other comparator outputs remain *high* since no current flows through the wires $Q2$, $Q1$, and $Q0$. This makes the outputs of the NAND gates $N1$ and




Figure 3.16: Acknowledgment transmission of *PMCM*.

$N2$ high, causing an up-going transition on the output of the C-element $C1$. Formally, the completion detection logic for the symbol 11 is as follows (The output of a gate X is denoted by $O(X)$):

$$\begin{aligned}
 &(V(30) = 0) \wedge (V(31) = 0) \quad (\text{current is } I_2) \\
 &\Rightarrow (O(N0) = 1) \wedge (O(N1) = 1) \\
 &\Rightarrow (O(N2) = 1) \\
 &\Rightarrow (O(C1) = 1) \\
 &(V(30) = 0) \wedge (V(31) = 1) \quad (\text{current is } I_1) \\
 &\Rightarrow (O(N0) = 1) \wedge (O(N1) = 0) \\
 &\Rightarrow (O(N2) = 0) \\
 &\Rightarrow (O(C1) = 0)
 \end{aligned}$$

3.2.4 Acknowledgment Transmission

The voltage-mode bundled-data acknowledge signal (*Ackin*), sent by the receiver module, is converted into a current-mode signal during transmission and back into a voltage-mode signal (*Ackout*) at the transmitter side. In this interconnect design, transmission of the acknowledgment signal also uses multilevel current sensing signaling. The driver and receiver circuits of this transmission along with distributed RLC model of the Acknowledgment wire is shown in Figure 3.16. The current through the acknowledgment wire becomes I_1 or I_2 when the acknowledgment signal from the receiving module is *low* or *high*, respectively. The receiver uses a current comparator circuit to detect the value of the current through the acknowledgment wire and output the result in voltage form. An inverter is used to amplify the comparator's output to full-swing.

3.2.5 Reference Voltage-Mode Interconnects

This reference interconnect also uses a two-phase protocol and 1-of-4 encoding, the difference being that it is implemented using voltage-mode signaling. In the $TPVm$ scheme one of the four wires makes a transition to indicate the presence of a new two-bit symbol. When this new symbol arrives to the receiving module, the receiver accepts the symbol and sends an acknowledgment to the sender module by changing the state of the acknowledge signal. Since voltage-mode signaling is used, the voltage on the interconnect swings from rail-to-rail over its entire length. This leads to large dynamic power consumption, large delay, and generation of power-supply noise. The usual approach to improve the performance of a voltage-mode interconnect is to insert repeaters or pipeline latches. Inserting repeaters decreases the signal propagation delay at the cost of increasing power consumption and chip area. A higher throughput can be obtained by using pipeline latches instead of repeaters to both amplify the signal and spread the link delay over multiple pipeline stages. This further increases power consumption and area costs compared to the simple repeater approach. Here both schemes are considered for the reference $TPVm$ interconnect. The pipelined and repeater-based implementations are called $TPVmP$ and $TPVmRep$, respectively. In the $TPVmP$ implementation pipeline stages are inserted in every 2mm along the link wire. This is based on the assumption that the typical distance between two neighboring (adjacent) routers in the on-chip mesh structure is 2mm [44] and that the local link length can be considered an upper limit for pipeline-free signal transmission as in [73]. In the $TPVmRep$ implementation optimal repeater insertion is used for both data and acknowledgment transmission. The required optimal number of repeaters and optimal size of the repeater are calculated using equation (36) of [48]. Using this equation the required number of optimal repeaters becomes $2.22 * L$ and the optimum size of the repeater becomes $76.5 * \text{minimum size inverter}$, where L is the wire length in mm.

The straightforward gate level implementations of the encoder which converts the two-phase bundled-data input to the delay-insensitive two-phase 1-of-4 protocol, the pipeline stage, and the decoder and completion detector which converts the delay-insensitive code back to the two-phase bundled-data form at the receiver side are shown in Figures 3.17, 3.18, and 3.19, respectively. The encoder consists of NOR gates which generate the select inputs for the multiplexers depending on the two-bit input codes, double-edge triggered flip-




Figure 3.17: Encoder of TPVm.

flops which are used to sample the symbol value at both edges of the request signal, and multiplexers each of which allows transition on the corresponding flip-flop output only when the appropriate input symbol is present. The decoder and completion detector circuit consists of XNOR gates which detect the transitions on the wires, NAND gates and a SR latch to decode the data back into the bundled-data form, and a four-input XOR gate together with an $N/2$ -input C-element for detecting completion. An inverter is used as both driver and receiver for the transmission of the two-phase acknowledgment signal between the pipeline stages in the $TPVmP$ implementation, as shown in Figure 3.18.

3.2.6 Simulation Results and Analysis

Simulation of $PM Cm$ and the two reference voltage-mode interconnects ($TPVmP$ and $TPVmRep$) was carried out in Cadence Analog Spectre and Hspice using

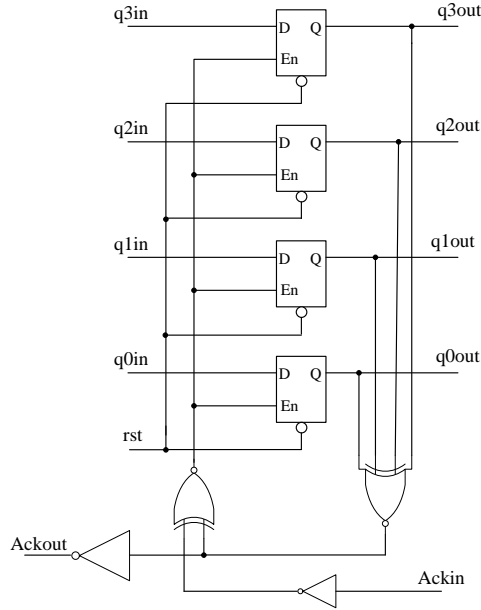


Figure 3.18: Pipeline stage, acknowledgment driver and receiver of *TPVmP*.




Figure 3.19: Decoder and completion detector of *TPVm*.

130nm technology from STMicroelectronics, and the supply voltage was set to 1.2V. Since high-performance signaling over long wires is considered, the wires were modeled using a distributed RLC model of metal 4. In order to accurately model crosstalk noise, both capacitive and inductive coupling between all wires was included. The bus consisted of eight parallel wires. The RLC values of the wires were extracted using FastHenry [45] and Linpar [46]




Figure 3.20: Forward latency of 1-of-4 encoded interconnects.

field solvers as in LEDR encoded interconnect in Section 3.1. The wire length was varied in the simulations from 2mm to 12mm.

Performance Analysis

Latency and throughput are considered the main parameters to analyze the performance of the multilevel current sensing on-chip interconnect along with the two reference voltage-mode interconnects. In the first reference interconnect, *TPVmP*, pipeline stages are inserted every 2mm assuming that the local wire length (between neighbor routers in a network) is 2mm. This improves the throughput at the expense of increased forward latency, power consumption and chip area. In the second reference interconnect, *TPVmRep*, optimal size repeaters are inserted at optimal distances.

Here forward latency is defined as the delay from a transition on the bundled-data request signal (*Reqin*) at the transmitter side to the corresponding transition on the bundled-data request signal (*Reqout*) at the receiver side (see Figure 3.1). In other words, the time required for one flit to traverse from the sending router to the receiving router. The change in the forward latency of the three interconnects when wire length is varied from 2mm to 12mm is shown in Figure 3.20. Since the *PMCm* interconnect uses current sensing signaling, its forward latency is much smaller than the latency of the two reference interconnects. The *PMCm*'s forward latency was less than one third of *TPVmP*'s latency for all simulated wire lengths. The forward latency of the pipelined voltage-mode interconnect was greater than 1.5 times *TPVmRep*'s latency for 4mm and longer communication distances.




Figure 3.21: Throughput of 1-of-4 encoded interconnects.

The throughput of $PMCM$, along with the two reference interconnects, is shown in Figure 3.21. The throughput of $PMCM$ varied from $5.102Gbps$ to $1.602Gbps$ when the wire length was varied from 2 to 12mm. At global communication distance of 8mm, the throughput of $PMCM$ was 1.53 and 1.88 times the throughput of $TPVmP$ and $TPVmRep$, respectively. In the case of the reference interconnects, $TPVmP$ achieved a throughput of $1.597Gbps$ while the throughput of $TPVmRep$ varied from $2.534Gbps$ to $1.041Gbps$ when the wire length was varied from 2 to 12mm. The reported latency and throughput values are for one group of 1-of-4 encoding.

The $PMCM$ interconnect is a better alternative than $TPVmP$ or $TPVmRep$ to realize high-performance long-range links. In addition to achieving high performance, $PMCM$ circuitry takes a smaller chip area compared to voltage-mode reference interconnects, $TPVm$. This is because the complexity and required chip area of the encoder and decoder of both $TPVm$ and $PMCM$ interconnects are almost the same. However, the number of required pipeline stages and the number of repeaters increase with wire length, which leads to increase in layout complexity and required area.




Figure 3.22: Power consumption of 1-of-4 encoded interconnects.

Power Analysis

The average total power consumption for 2-bit data transfer on the proposed current sensing and the two reference interconnects when communication distance was varied from 2 to 12mm is shown in Figure 3.22. *PMCm* has consumed 38% or more power than that of *TPVmP* at all wire lengths. The power consumption of *TPVmP* increases at a faster rate with wire length compared to *PMCm* due to the increase in the number of pipeline stages. As a result the power consumption difference between these two interconnects decreases at global wire lengths. *PMCm*'s power consumption was 10 to 36% lower than that of *TPVmRep* starting from 6mm wire length, this is because of the increase in the number of repeaters inserted at global lengths of the wire. The power dissipated by the *TPVmRep* interconnect was higher than 2 times *TPVmP*'s consumption for all wire lengths.

The energy per bit of the interconnects is shown in Figure 3.23. The energy per bit of *PMCm* was 26 to 58% less than that of *TPVmRep* and 15 to 37% larger than the *TPVmP*'s energy dissipation. *TPVmP* and *TPVmRep* dissipate least and highest energy, respectively at all wire lengths.

Noise Analysis

The impact of crosstalk noise on latency and throughput was also studied. In this analysis, 4-bit parallel data transfer was assumed. This requires 9 (8 parallel data transmissions + 1 acknowledgment) physical wires since 1-of-4 encoding is used. The acknowledgment wire was designed as having shielding from the parallel data transmission wires, to counteract the coupling effect.

The wires were modeled as transmission lines which have both capacitive and inductive coupling between each other. During this analysis, minimum wire separation distance with minimum global pitch specified in 130 nm technology and 1.2V supply voltage were used. The delay variation due to both capacitive and inductive coupling was simulated by considering the worst-case and best-case switching patterns. These switching patterns depend on the RLC values of the wire. In the simulation setup it is assumed that the capacitive coupling dominates the inductive coupling which is the most usual case in on-chip parallel wires. The effect of crosstalk on latency and throughput when the wire length was varied from 2mm to 12mm is shown in Figures 3.24 and 3.25, respectively.

During best-case and worst-case switching, the latency variation of $TPVmP$ was slightly less than that of $PM Cm$. For example, at a wire length of 8mm, the increase in latency due to best-case switching from the crosstalk free latency of $TPVmP$ and $PM Cm$ was 59.8% and 62.3%, respectively. In worst-case switching, the $TPVmP$ and $PM Cm$ latency variations were 144% and 147%, respectively, at the same wire length. In fact, these percentage values are rather large because in the nominal case shown in Figure 3.20 the considered capacitive loads were only to ground. In other words, the nominal case capacitive loads do not consider the loading effect of the coupling capacitances. The decrease in throughput due to crosstalk was greater for $TPVmP$ than for $PM Cm$, specially at long wire length. For example at 12mm wire length, the throughput of $TPVmP$ was decreased by 38% while the $PM Cm$ was only by 30%.




Figure 3.23: Energy per bit dissipation of 1-of-4 encoded interconnects.




Figure 3.24: Crosstalk effect in latency of *PMCM* and *TPVmP*.




Figure 3.25: Crosstalk effect in throughput of *PMCM* and *TPVmP*.

3.3 Dual-Rail Encoded Differential Current Sensing Interconnect

As already discussed in Chapter 1, global on-chip interconnects get slower with technology scaling and dissipate more power. At the same time signal integrity issues become challenging due to crosstalk, PVT variations and noise. PVT variations cause the signal propagation delay to be uncertain, which in turn affects the performance and reliability of the interconnect significantly. It has been demonstrated that high speed, energy efficient and better noise immunity




Figure 3.26: Simulation waveforms of *PMCM*.

can be achieved using differential current-mode signaling (See Section 2.6.1). In addition, reliable on-chip communication in the presence of delay variations is possible through the use of self-timed delay-insensitive data transfer. Thus, integrating differential signaling with delay-insensitive data transfer enables high-performance as well as robustness towards both noise and delay variations. However, integrating these two techniques has considerable area and power overhead as it requires four wires per bit (two for delay-insensitive encoding and two for differential signaling). In this section, a high-performance on-chip interconnect based on novel area and power efficient integration of delay-insensitive data transfer and differential current sensing signaling is presented.

The proposed Dual-rail encoded differential current sensing interconnect (*Dualdiff*) implements both delay-insensitive and differential signaling schemes with only two wires per bit by using a novel encoding and differential current sensing. This leads to a smaller area and smaller power consumption. Its encoding technique and circuits are discussed in the next subsections. As in the other interconnects presented in this chapter, both its input and output signals are assumed to be in the two-phase bundled-data encoded form.

3.3.1 Encoding and Its Implementation

In conventional delay-insensitive data encoding transmission of N -bit of data requires $2N$ wires. The doubling of the wire count compared to bundled-data encoding has a significant effect on the wiring area and routing complexity. Delay-insensitive data transmission using N instead of $2N$ wires for four-phase [37] and two-phase [36] handshaking has been proposed using single-ended multilevel current-mode signaling. Both of these works use different current levels to encode the data and data validity indicator together.

In the novel encoding technique considered here, current directions and current values are used simultaneously to get both delay-insensitivity and differential signaling features. A change in the current level on the wire indicates arrival of new data (delay-insensitivity), while the direction of the current flow reveals the logical value of the transmitted bit. The encoding protocol is shown in Table 3.5. When the transmitted bit is 1 the driver sources current to one of the wires and sinks current from the other wire, and vice versa for bit 0 transmission. The value of the current on the wires switches between I_1 and I_2 at every new transmission event. The communication protocol of this interconnect is shown in Figure 3.27.

The encoder receives as inputs data (Din) and request ($Reqin$) signals in the bundled-data form. It encodes the data and request together and outputs voltage pulses which serve as inputs to the driver. The encoder circuit is shown in Figure 3.28. The En signal is used for enabling transmission and it is an active-high signal. That is, it is *high* during transmission and *low* during idle periods. The edge sensitive flip-flops sample the data input at the edges of the bundled-data two-phase request input ($Reqin$). Only one of the AND gates ($N1$ to $N4$) outputs voltage pulse at every transmission event. The encoder outputs, $Ind1$, $Ind2$, $Ind3$ and $Ind4$ are inputs to the differential driver.

Table 3.5: Encoding protocol of *Dualdiff* interconnect.

Din	Reqin	En	Ind1	Ind2	Ind3	Ind4
0	1 \rightarrow 0	1	1	0	0	0
0	0 \rightarrow 1	1	0	0	1	0
1	1 \rightarrow 0	1	0	1	0	0
1	0 \rightarrow 1	1	0	0	0	1




Figure 3.27: Communication protocol of *Dualdiff* interconnect

3.3.2 Driver, Receiver and Completion Detector

The designs of driver, receiver and completion detector circuits are presented in this section and shown in Figures 3.29 and 3.30. As shown in Figure 3.29, source coupled CMOS bipolar current-mode drivers are used. Such a driver conveys two currents of the same amplitude but opposite polarity to the wires such that not only the effect of supply voltage fluctuations on the wires is minimized, but also the noise injection from the driver to the substrate is minimal. Two bipolar current-mode drivers are used in order to drive the two current levels I_1 and I_2 . At every transmission event, the driver sources current to one of the wires and sinks the same amount of current from the other wire.

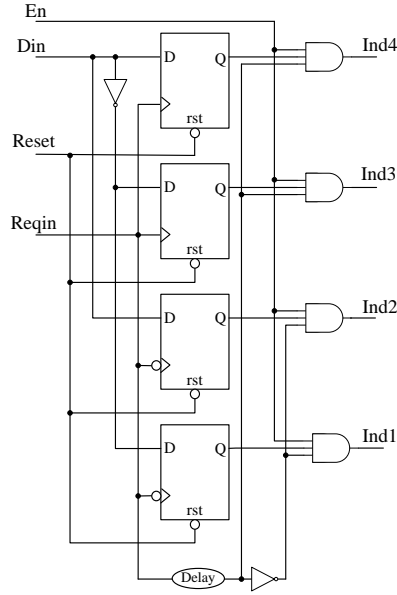


Figure 3.28: Encoder of *Dualdiff* interconnect.

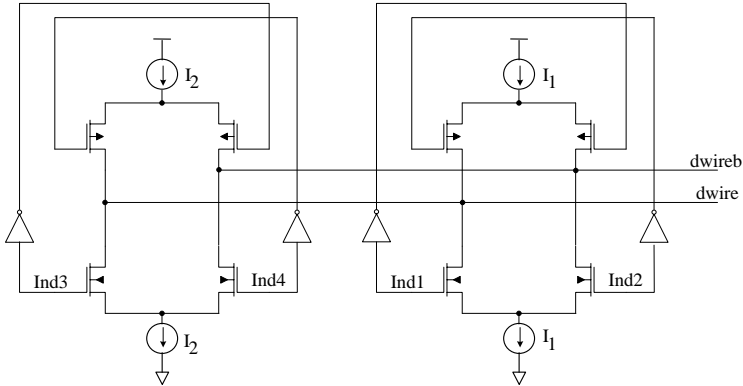


Figure 3.29: Driver of *Dualdiff* interconnect.

The receiver senses the direction of the current flow to retrieve the transmitted data. The receiver of this interconnect has two stages, a current direction sensor and a differential amplifier as shown in Figure 3.30. The current direction sensing circuit is a modified version of the one presented in [49]. In our design, the termination transistor is diode connected to ensure its saturation operation and also to mirror the wire current. Furthermore, unlike in [49] the current sensor output is not connected back to the termination transistor gate to avoid its effect on output switching. Thus, the *Mn3* and *Mn6* transistors provide a low impedance path to ground for current sourced by the driver

and serve as best matched impedance termination since they operate in saturation. Consider the top current sensor in Figure 3.30. The transistor $Mp1$ provides negative feedback to transistor $Mn1$. It turns the gate of $Mn1$ on and off as required and helps in modulating the input impedance. The transistor $Mp2$ provides a constant current bias hence regulating the transconductance of $Mn1$. The source terminal of the transistor $Mn1$ is connected to *dwire*. When current is sunk by the driver $Mn1$ becomes *on* and pulls the output of the current sensor to *low*. When current is sourced by the driver, the source voltage of $Mn1$ rises thus turning it *off*. In this case the current flows through the load transistor $Mn2$ to the output making the output voltage of the current sensor *high*.

The output of the current sensor is not full swing. In addition, the receiver needs to have high common-mode noise rejection capability in order to take full advantage of differential signaling. Due to these, the second stage, a high-speed self-biased differential amplifier is used. The amplifier consists of source coupled NMOS and PMOS transistors ($Mn8$, $Mn9$, $Mp6$, and $Mp7$). It operates at high speed because its output switching currents are significantly greater than its quiescent current. It has also a higher differential-mode gain compared to conventional amplifiers and a large common-mode input range because its bias condition adjusts itself to accommodate the input swing [86, 90]. The bundled-data input to the receiving module, $Dout$, is the output of the amplifier without requiring additional data decoding logic.

Since the receiving side has a bundled-data interface, it requires a separate data validity indicator ($Reqout$ in Figure 3.30). Based on the encoding, it is known that the current on the wire becomes $\pm I_2$ when the request from the sending module is *high* and $\pm I_1$ when the request is *low*. To decode the request signal the output currents of both wires are compared separately with the reference current using a current comparator as shown in Figure 3.30. The currents in *dwire* and *dwireb* are mirrored into transistors $Mn10$ and $Mn11$, respectively. These mirrored wire currents will be compared with the reference current, which is generated by diode connected $Mp8$ and $Mn12$ transistors. The reference current value is $0.5 * (I_1 + I_2)$. If either of the wire output currents is greater than the reference current, $Reqout$ will be *high* and if the current in both wires is less than I_{ref} then $Reqout$ will be *low*.

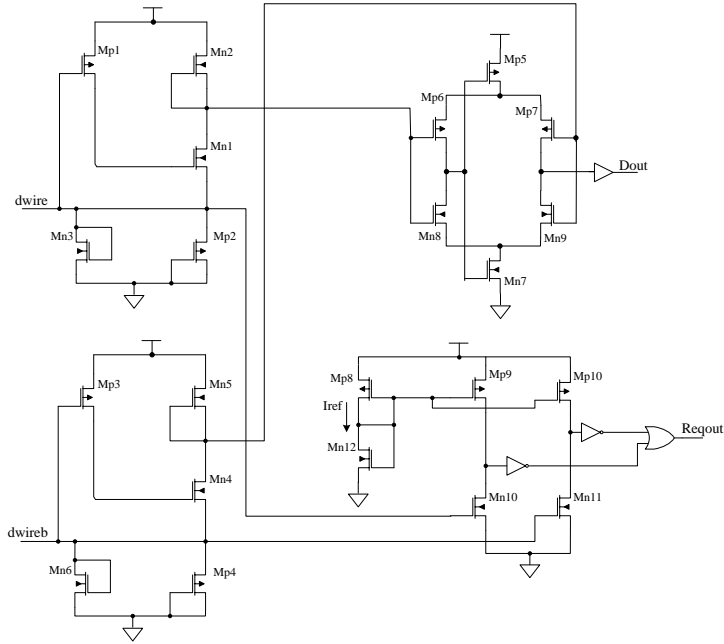


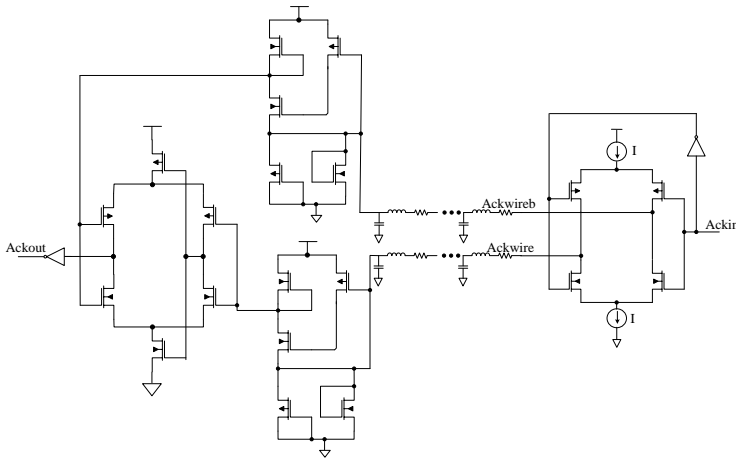
Figure 3.30: Receiver and completion detector of *Dualdiff* interconnect

3.3.3 Acknowledgment Transmission

In this interconnect design, an acknowledgment is sent for each transmitted bit from the receiving module. The voltage-mode bundled-data acknowledgment signal (*Ackin*) is converted into a differential current signal during transmission and back into a voltage-mode signal (*Ackout*) at the transmitter side. The acknowledgment transmission also uses differential current sensing signaling. The driver and receiver circuits of this transmission along with distributed RLC model of the acknowledgment wire is shown in Figure 3.31. The current through the differential acknowledgment wires, *Ackwire* and *Ackwireb* becomes $+I$ and $-I$, respectively when *Ackin* is *high* and vice versa when *Ackin* is *low*. The two current direction sensors, which are the same as in data receiver, detect the direction of current flow. The self-biased differential amplifier retrieves the transmitted acknowledge signal using the output of the current sensors.

3.3.4 Simulation Results and Analysis

The wire properties were set according to the ITRS 65nm technology node for global wiring. RLC matrices of the wire were extracted using FastHenry


 Figure 3.31: Acknowledgment transmission of *Dualdiff* interconnect

and Linpar fieldsolvers. During extraction, both wire width and separation distance were set to 210nm and the wire thickness was set to 242nm. In the interconnect simulation, the wires were modeled as a distributed RLC using the extracted per unit values. The circuits were designed and simulated in Cadence Analog Spectre using 65nm CMOS technology from STMicroelectronics and 1V supply voltage. The simulation waveforms are shown in Figure 3.32. It consists of the input data and request, current and voltage of the wire, the amplifier output, the request output, and acknowledgment input and output.

A reference interconnect is designed and simulated to compare the performance, power consumption and area of the *Dualdiff* interconnect. In order to determine the contribution of the novel integration scheme in *Dualdiff*, the reference interconnect uses conventional integration of delay-insensitive data transfer with differential current sensing signaling. It uses LEDR data encoding due to its simpler data decoding and completion detection schemes. It requires four wires per bit, that is, two for differential phase and the other two for differential state transmissions as in [78]. The same signaling circuits are used as in the *Dualdiff* interconnect to have proper comparison. The reference LEDR encoded differential current sensing interconnect (*LEDRdiff*) is shown in Figure 3.33. It requires two bipolar differential current-mode drivers, four current direction sensors, and two differential amplifiers. Its data encoding and data validity indicator decoding circuits are also shown in Figure 3.33.

The forward latencies of both interconnects when the wire length varies from 1 to 5mm are shown in Figure 3.34. The latency of the proposed interconnect is much smaller than that of the reference one. For example, at




Figure 3.32: Simulation waveforms of *Dualdiff*.

2mm wire length its latency was less than one-half of that of the reference. As shown in Figure 3.35, the throughput of the *Dualdiff* interconnect is 1.92 and 1.54 times that of the reference interconnect for 1 and 5mm long links, respectively. It has a throughput of 1.34Gbps at 5mm wire length. The performance penalty of the reference interconnect comes mainly from its encoding and completion detection circuits.

The power consumption of the proposed interconnect is smaller than that of the reference interconnect as expected (Figure 3.36). At 5mm wire length 24% power savings has been gained compared to the reference interconnect. The average energy dissipated per every transmitted bit is also examined for both interconnects and is shown in Figure 3.37. The energy per bit of *Dualdiff* interconnect is much smaller than the reference energy dissipation. Its energy

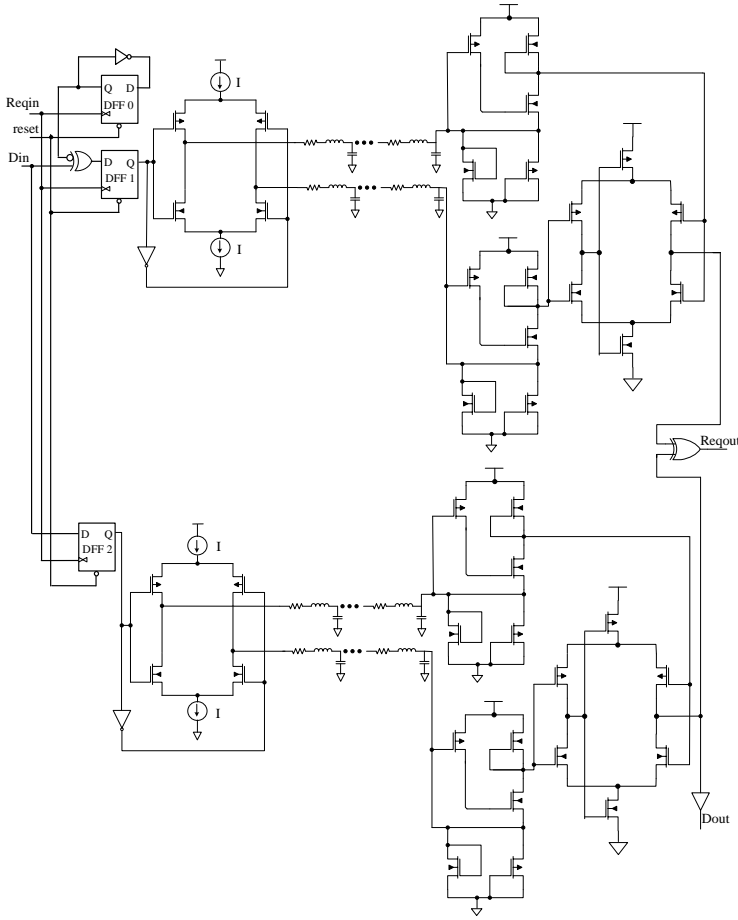


Figure 3.33: LEDR encoded differential interconnect

per bit is less than one-third and one-half of that of the reference at 1 and 5mm wire lengths, respectively. The reference interconnect energy dissipation increased much faster than the proposed one for longer wires. All these analyses show the superiority of the proposed interconnect over the conventional delay-insensitive differential interconnect. Moreover, the proposed interconnect is area efficient since it reduces the required number of wires by half. It requires 15% less active area and 40% less wiring area for 2mm long interconnect (See Table 3.6).

Table 3.6: Area comparison between *Dualdiff* and *LEDRdiff*.

Interconnect	Active Area [μm^2]	Wiring Area [μm^2]
Dualdiff	16	2520
LEDRdiff	19	4200




Figure 3.34: Forward latency of *Dualdiff* and *LEDRdiff*.




Figure 3.35: Throughput of *Dualdiff* and *LEDRdiff*.

3.4 Chapter Summary

The design and analysis of three high-performance and delay-insensitive global on-chip interconnects were presented. The delay-insensitivity makes the communication robust and attains average-case performance rather than worst-case which is the situation in communication based on timing constraints. The first interconnect presented in this chapter, *LEDRCm*, has achieved higher throughput and dissipated lower energy per bit than the conventionally implemented *LEDRVm* and *TPDRVm* interconnects. The second interconnect presented in this chapter, *PMCm*, uses two-phase 1-of-4 encoding and multi-level current sensing signaling. The performance analysis showed that *PMCm* has achieved higher throughput and lower latency than its two reference in-




Figure 3.36: Power consumption of *Dualdiff* and *LEDRdiff*.




Figure 3.37: Energy per bit dissipation of *Dualdiff* and *LEDRdiff*.

terconnects, *TPVmP* and *TPVmRep*. In addition, the energy per bit dissipation of *PMCM* was lower than that of the *TPVmRep*. The last interconnect presented in this chapter is based on a novel integration of delay-insensitive encoding and differential current sensing signaling. Only half number of wires are required compared to conventional integration of the two schemes, making it both area and power efficient. It has achieved higher performance than a reference interconnect, *LEDRdiff*. It has also consumed lower power and dissipated lower energy per bit than the *LEDRdiff*. Therefore, the presented three on-chip interconnects are prominent candidates for high-performance, energy efficient and delay variations tolerant global on-chip communication.

Chapter 4

Enhancing Completion Detection Performance

In the previous chapter, designs of high-performance and delay-insensitive current sensing interconnects have been presented. In delay-insensitive transmission, validity of the data is encoded within the data itself at the transmitter, and the data validity test, i.e., completion detection, as well as data decoding is performed at the receiver. The delay incurred due to completion detection increases with bit width of the transmission channel and affects the performance of the communication significantly. In order to overcome this overhead, a high speed completion detection technique along with its CMOS implementation is designed and presented in this chapter. Unlike the conventional detection circuits, the delay of the presented completion detection circuit is not affected by the bit width of the channel. This optimizes the performance of delay-insensitive current sensing links further since it was already demonstrated in Chapter 3 that the current sensing interconnects achieve higher performance and better power efficiency than the interconnects using voltage-mode signaling with repeaters or pipelines.

The chapter is organized as follows. First delay-insensitive bit parallel transmission and the overhead of completion detection in such transmission are discussed. The novel high speed completion detection technique and its implementation details are presented in Section 4.2. Two delay-insensitive links, which use the proposed completion detection technique are presented as case studies in Sections 4.3.1 and 4.3.2. The design of the acknowledgment interconnect for the the case study links are discussed in Section 4.3.3 and the

design of reference links are explained briefly in Section 4.4. In Section 4.5 simulation details and analysis of performance, power, energy and area of the case studies as well as reference links are presented. The summary presented in Section 4.6 concludes the chapter.

4.1 Delay-Insensitive Bit Parallel Transmission

A GALS communication method is used in almost all proposed NoC designs and is expected to be an attractive approach to overcome many of the timing problems [50]. The GALS approach simplifies clock tree design and results in easily scalable clocking systems. It also enables better energy savings since each functional unit can easily have its own independent clock and voltage [51]. Furthermore, it allows easy implementation of a distributed power management system for the entire chip [52]. A fully self-timed NoC in the GALS clocking scheme gives a better network saturation threshold, smaller average power consumption, slightly higher maximal bandwidth and much smaller packet latency (2.5 times smaller) than the multi-synchronous NoC implementation [53]. A number of fully asynchronous GALS NoCs have been proposed and implemented, such as MANGO [81], ANoC [82], ALPIN [76], FAUST chip [83] and QNoC [84]. Hence, due to the advantages of a fully self-timed GALS NoC and given its ability to work reliably in the presence of variations, a self-timed delay-insensitive link between NoC routers is a natural choice.

In bit parallel transmission, the throughput of a delay-insensitive link decreases when the bit width of a channel increases, because of the increase in the delay of completion detection. Conventionally, completion detection is carried out by sensing either voltage transitions or levels on each data wire. This requires logic circuitry whose delay increases drastically when the channel bit width increases, causing a bottleneck to achieve high performance communication using a delay-insensitive interconnect. The completion detection logic for two-phase dual-rail and 1-of-4 encoded links are shown in Figure 4.1 and 4.2 for 32-bit parallel transmission. These two encodings are the simplest and the most commonly used on-chip delay-insensitive codes [54], requiring two signal wires per each transmitted bit. In 32-bit 1-of-4 encoded transmission, the delay incurred due to the completion detection is the sum of delays of a 4-inputs XOR gate and four 2-input C-elements, see Figure 4.1. The completion detection logic of 32-bit two-phase dual-rail encoded transmission has an

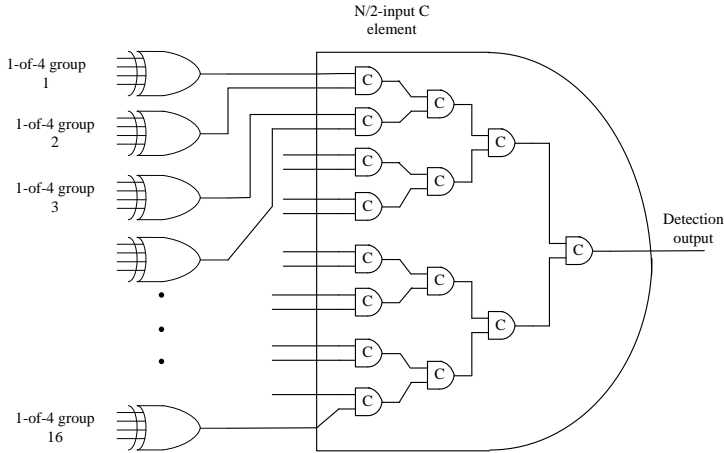


Figure 4.1: Completion detector of 32-bit Two-Phase 1-of-4 Transmission

extra delay of a 2-input C-element compared to 1-of-4 as shown in Figure 4.2. On the other hand, its XOR gates have only 2-inputs. In 65nm technology the completion detection delay of 32-bit two-phase dual-rail encoded transmission is 217ps. For example, if the data receiving block runs at 5GHz, it has to wait more than one clock cycle only because of the completion detection. The delay of completion detection for N-bit two-phase 1-of-4 encoded transmission is the sum a 4-input XOR gate delay and $(\log_2 N - 1) \cdot 2$ -input C-element delay. In case of two-phase dual-rail encoded N-bit transmission, the delay of completion detection is the sum of a 2-input XOR gate delay and $\log_2 N \cdot 2$ -input C-element delay. So, the larger the channel bit width is, the longer is the overall time spent in completion detection, because each detection circuit becomes a tree of logic elements.

Traditionally, optimal repeater insertion together with pipelining is the method to achieve high throughput in global voltage-mode on-chip interconnects. If such a pipelined interconnect is delay-insensitive, each pipeline stage including the transmitter and receiver themselves, requires area, power, and time consuming completion detection logic. At the receiver side completion detection is needed to indicate the validity of the arrived data and at the transmitter side to indicate the acceptance of the transmitted data, since in a pipelined channel an acknowledgment is sent per group instead of for the whole channel. For example, in [56] an acknowledgment is sent per each 1-of-4 group, this helps to reduce the speed penalty due to large detection logic

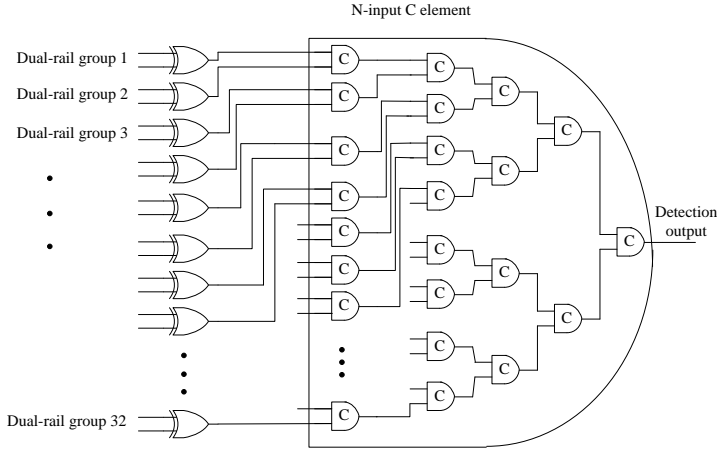


Figure 4.2: Completion detector of 32-bit Two-Phase Dual-Rail Transmission

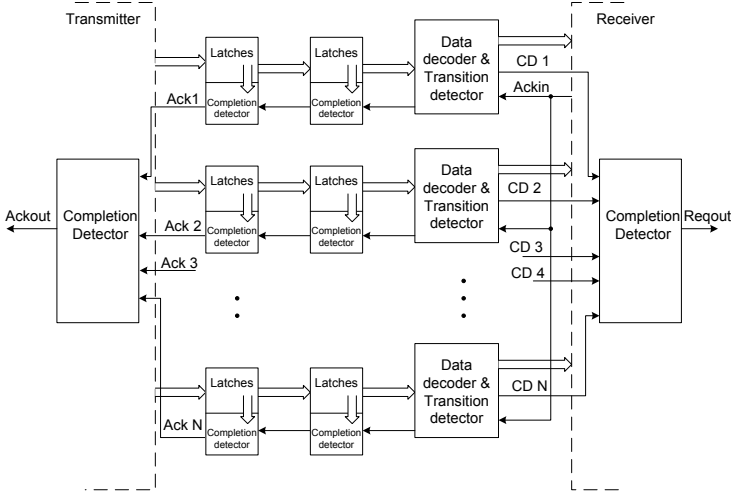


Figure 4.3: Completion detection in a pipelined voltage-mode link

at each pipeline stage. A block level diagram of a delay-insensitive pipelined voltage-mode link is illustrated in Figure 4.3 showing delay causing completion detection blocks at the pipelines latches, receiver and transmitter.

A delay-insensitive current sensing link does not require repeaters nor pipelining to boost its throughput, indicating that completion detection is carried out only once at the receiver, and therefore it achieves higher performance and better power efficiency than a pipelined voltage mode interconnect. This has been proved in Chapter 3 where, however, wire currents are first converted to voltages, and the actual completion detection is carried out in the

voltage mode, resulting in a significant speed penalty. Hence, a high speed completion detection technique which uses the wire currents directly without conversion to voltage mode and carries out completion detection in current mode is proposed and presented in this chapter. Unlike with the conventional completion detection logic, delay of the proposed scheme does not increase with the link bit width, it is bit width independent.

4.2 High-Speed Completion Detection Technique

The proposed completion detection technique uses directly the current on each data wire and carries out completion detection in the current mode. The idea is to sum the currents on all the data wires of a channel and then compare this sum current to a reference current. Implementation of this technique requires only current mirrors, a current source, and a current comparator. The comparator takes as inputs the sum current and the reference, and outputs a full-swing completion detection signal. This signal becomes *high* when the sum current is greater than the reference current, indicating the validity of every received data signal. Unlike with the conventional voltage mode scheme, the speed of the proposed scheme is not affected by the channel bit width, because the current summation is carried out by wiring and its delay is only due to comparing currents.

The completion detection circuit is shown in Figure 4.4. It supports detection of dual-rail (1-of-2) and 1-of-4 encoded data, both of which use $2N$ wires to convey N -bit data, so that the number of active wires per transmission is N in the 1-of-2 case and $N/2$ in the 1-of-4 case. The diode connected transistors $Mw(1)$ to $Mw(2N)$ (one transistor per wire) are used to input the currents on the wires and mirror them to $Ms(1)$ to $Ms(2N)$ transistors, respectively, which are connected together to generate the sum current $I(sum) \equiv Code \times (S-1) \times N \times I(w)$. Here $I(w)$ is the nominal current on a single data wire, N is the number of bits, S is the current down-scaling factor ($S \geq 1$) indicating the current drive ratio between the transistors $Mw(i)$ and $Ms(i)$, and $Code$ is either 1 (1-of-2 code) or 0.5 (1-of-4 code) indicating the number of active wires per transmitted bit. By using a scaling factor (S) larger than 1 the power consumption of the circuit can be efficiently reduced. The reference current, $I(ref)$, is generated using an addition based process invariant current source [57]. Its value is $I(ref) \equiv Code \times (S-1) \times (N-0.25) \times I(w)$.

This value is chosen in order to compare after the wire current of the last bit reaches 75% of its current. The comparator transistor $MpC1$ mirrors $I(sum)$ to $MpC2$, and $MnC1$ mirrors $I(ref)$ to $MnC2$. The comparator output becomes *high* (*low* otherwise), when the current of $MpC2$ is greater than that of $MnC2$. Due to process and supply voltage variations, the sum and reference currents may vary from their nominal values affecting the reliability of completion detection. In order to have correct operation, condition 4.1 has to be fulfilled.

$$\Delta(I(sum)) + \Delta(I(ref)) < (S - 1) \times \left(\frac{I(w)}{2}\right) \quad (4.1)$$

$$\Delta(I(sum)) \cong Code \times (S - 1) N \times I(we) \quad (4.2)$$

For $I(sum)$ the variation can be expressed as in Equation 4.2, where $I(we)$ is the worst case variation of the current on a single wire. For $I(ref)$, according to an extensive Monte Carlo analysis of the circuit for $N = 2$ to 64-bit, $I(w) = 200\mu A$ to $300\mu A$, and $S = 1$ to 5, it is safe to assume that the variation is within the bound expressed in Equation 4.3.

$$\Delta(I(ref)) < (S - 1) \times \frac{I(w)}{6} \quad (4.3)$$

Substituting Equations 4.2 and 4.3 into Equation 4.1 and solving for N yields the following constraints:

$$N < \frac{1}{3} \times SNR \quad (4.4)$$

$$N < \frac{2}{3} \times SNR \quad (4.5)$$

where SNR is the signal-to-noise ratio of a single data wire and expressed in Equation 4.6. Condition 4.4 is for the 1-of-2 (Code = 1) and Condition 4.5 is for 1-of-4 (Code = 0.5) codes.

$$SNR = \frac{I(w)}{I(we)} \quad (4.6)$$

The higher the SNR is the larger number of bits (N) can be reliably transmitted and detected. Furthermore, for a given SNR , a 1-of-4 encoded channel can be twice as wide as a 1-of-2 encoded channel, because the number of active

wires is half of that of the 1-of-2 case. The relation between N and SNR for the 1-of-2 and 1-of-4 encoded channels is shown in Figure 4.5.

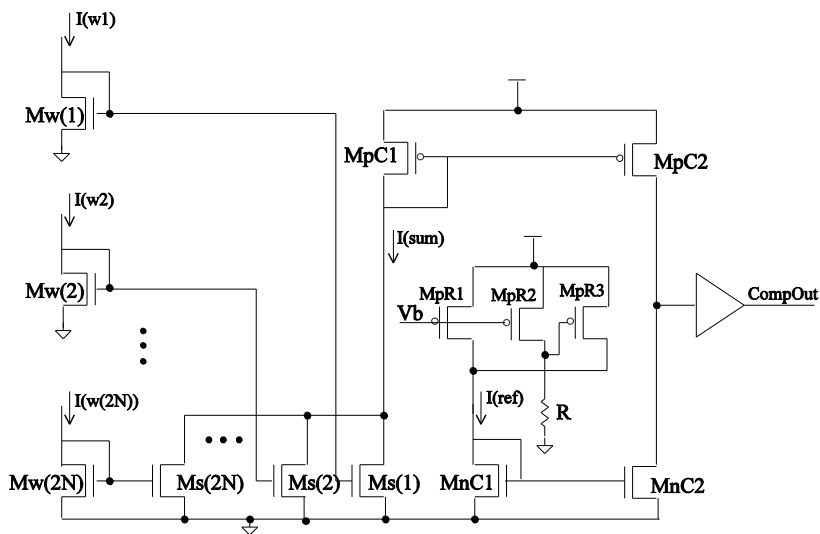


Figure 4.4: High-speed completion detection circuit.




Figure 4.5: Bit-width versus SNR of a wire for a reliable detection.

4.3 Case Studies

In this section, the redesign of *PMCM* (Section 3.2) and *Dualdiff* (Section 3.3) interconnects in order to use the proposed completion detection technique is presented. In the initial design of these two interconnects, the wire currents

have been converted to voltages and then conventional completion detection has been carried out in voltage mode. In the analysis section, the performance improvement due to the novel completion detection will be presented.

4.3.1 1-of-4 Encoded Current Sensing Interconnect

As already discussed in Section 3.2, in the $PMCm$ interconnect the two-phase bundled-data voltage-mode data and request signals are converted into pulsed 1-of-4 multilevel current sensing signaling at the transmitter side. At the receiver side, delay-insensitive current sensing signaling is turned back into bundled-data voltage-mode communication. On the wires information is represented as current rather than voltage transitions, one of the four data wires draws current to indicate the presence of a new two-bit data symbol. The current detected at the receiver has three different values: 0, I_1 , and I_2 . The values I_1 and I_2 are used when the voltage-mode request signal $Reqin$ at the transmitter side is *low* and *high*, respectively, reflecting the adopted two-phase communication protocol. The value 0, in turn, means that there is no symbol on a wire, representing the idle period. Since the design of encoder, driver, receiver and data decoder of this interconnect is the same as in Section 3.2, only the completion detection implementation is discussed here. The $PMCm$ interconnect which uses the proposed completion detection technique is called $PMCmFCD$.

The completion detection circuit for 4-bit transmission using $PMCmFCD$ interconnect is shown in Figure 4.6. This detector requires two current comparators because of the power saving scheme of the $PMCm$ interconnect. That is, during the idle period of the transmission the currents on the wires are switched off. This switching off should not affect the state of the two-phase bundled-data $Reqout$ signal, which is the output of the completion detector. The main comparator, composed of $Mp2$ and $Mn2$, compares the sum of wire currents with a reference current. For N -bit transmission, this reference current I_{ref} is in the range of $S \times N/2 \times I_1 < I_{ref} < S \times N/2 \times I_2$, where S is the wire current scaling factor. If the current through the transistor $Mp2$ is greater than the reference current in the transistor $Mn2$, the output $C1$ goes *high*, otherwise it goes *low*. $C1$ is latched to the output of the completion detector only when there is current on the wires. To determine the availability of current on the wires, an additional comparator is required. This comparator compares the sum of wire currents with a small reference current I_{ref1} , which

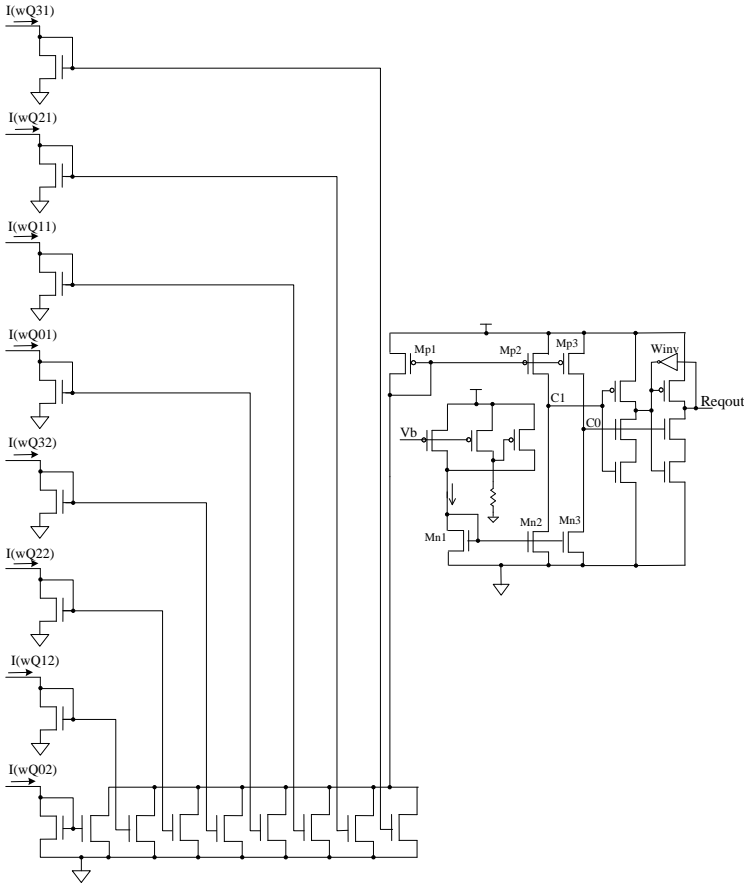


Figure 4.6: Completion detection circuit of 4-bit *PMCMFCD* link.

is in the range of $0 < I_{ref1} < S \times N/2 \times 0.5 \times I_1$. As long as the current in the transistor *Mp3* is greater than I_{ref1} (the reference current that is mirrored to transistor *Mn3*), the output *C1* is latched to *Reqout*. If the current in *Mp3* is less than I_{ref1} , the output *C0* becomes low, which in turn causes the latch to enter the hold mode, and there is no change in the *Reqout* signal. *Winv* in Figure 4.6 is a weak inverter which is used as a keeper.

The major performance improvement due to the use of this completion detection scheme becomes significant when the number of bits transmitted in parallel increases. In *PMCM*, where the completion detection is carried out in the voltage mode after converting the wire current, one 1-of-4 group (2-bit transmission) detector requires eight current comparators, two 4-input NAND gates, a 2-input NAND gate and a 2-input C-element. For N -bit transmission, it requires $(N/2)$ times the components of one 1-of-4 group detector in addition to the $(N/2)$ -input C-element. Thus, the speed penalty due to the completion

detection becomes considerable especially in a high performance system. For example, in 65nm technology, the delay of a 64-bit transmission detector of *PMCM* is 124ps, whereas in *PMCMFCD* it is only 52ps. Detailed analysis will be presented in Section 4.5.

4.3.2 Dual-rail Encoded Differential Current Sensing Interconnect

In this interconnect, differential current sensing signaling and two-phase dual-rail encoding are integrated in an area and power efficient manner. Both current directions and current values are used simultaneously to get both delay-insensitivity and differential signaling features. A change in the current level on the wire indicates arrival of new data (delay-insensitivity), while the direction of the current flow reveals the logical value of the transmitted bit. As already discussed, retrieving the data validity signal is necessary because the receiving module has a bundled-data encoded interface. The arrival of new data on the wire is indicated by either of the two wire current levels I_1 and I_2 , enabling the use of the proposed completion detection technique as a more appropriate choice. This is because it is fast, consumes less power and area compared to the *Dualdiff* interconnect (Section 3.3) where each wire current is first sensed, then converted to voltage and detection is carried out in the voltage mode. The *Dualdiff* interconnect which uses the proposed completion detection technique is here called *DualdiffFCD*.

In *Dualdiff* the current on the wire becomes I_2 when the request input from the sending module is *high* and it becomes I_1 when the request input is *low*. During the idle period of transmission, the wire current is switched off to save power. The completion detection circuit of *DualdiffFCD* is shown in Figure 4.7 for 4-bit transmission. It is the same as the completion detection circuit of *PMCMFCD* since both use three current levels, 0, I_1 and I_2 .

4.3.3 Acknowledgment Transmission

The designs of acknowledgment transmission circuits are the same for the two case study interconnects. The voltage mode bundled-data acknowledgment signal (*Ackin*), sent by the receiver module, is converted into a current mode signal during transmission and back into a voltage mode signal (*Ackout*) at the transmitter side. Current sensing differential signaling is used.

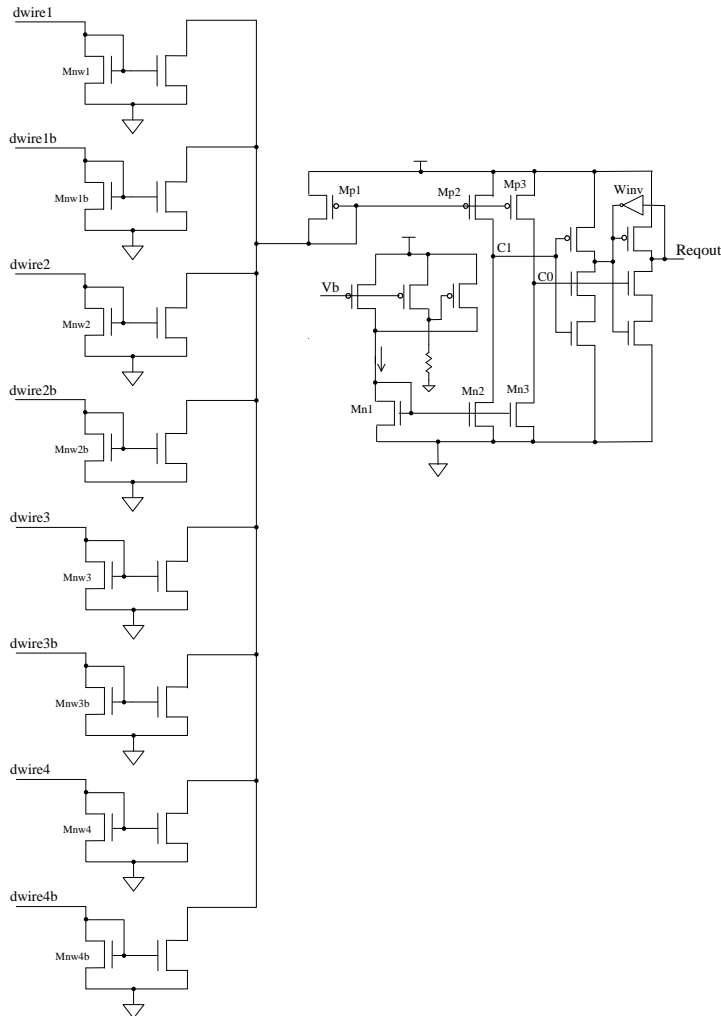


Figure 4.7: Completion detection circuit of 4-bit *DualdiffFCD* link.

A source-coupled differential current-steering driver, shown in Figure 4.8, is used. It is fast because it has an extremely sharp transient response. It has also an advantage of reducing the AC component of the power supply noise because the circuit draws constant current from the supply. The complementary outputs of the driver are attached to the differential pair of wires. The other end of the transmission is parallel terminated into a positive voltage. When *Ackin* makes a transition from *low* to *high*, there is current in *wire0* and no current in *wire1*. When *Ackin* makes a transition from *high* to *low* there is current in *wire1* and no current in *wire0*. Diode connected *Mpt0* and *Mpt1* transistors are used as termination loads. The transconductance of these transistors is regulated through the use of *Mpr0* and *Mpr1*. The receiver is a

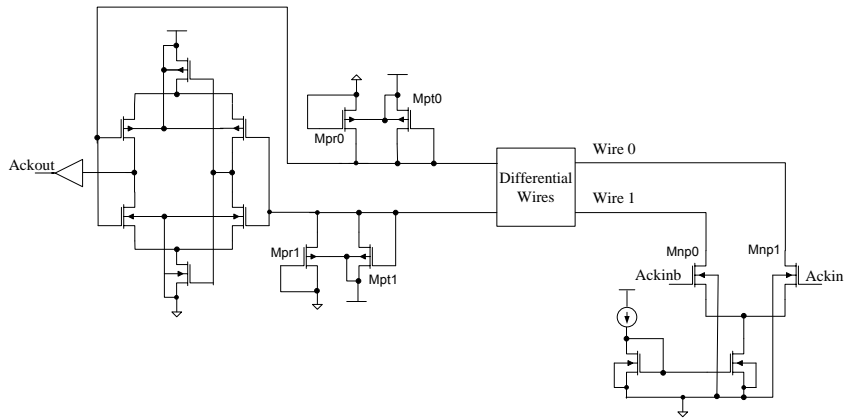


Figure 4.8: Acknowledgment signal transmission.

high speed self-biased differential amplifier, which has a high differential-mode gain.

4.4 Reference Cases

Four interconnects are designed to serve as reference cases. Two of them are used to determine the performance improvement enabled by the presented high speed completion detection technique. They are the ones presented in Sections 3.2 and 3.3, where the wire currents are converted into voltages and completion detection is carried out in the voltage mode. The purpose of the other two is to analyze the contributions of the current sensing signaling along with the high speed completion detection. They use optimally pipelined voltage-mode signaling. The pipeline stages are inserted at distances where optimal throughput can be achieved. One uses two-phase 1-of-4 encoding ($1of4VmP$) and the other uses two-phase dual-rail encoding ($DualVmP$). The implementation details of the $1of4VmP$ interconnect are mostly the same as in Section 3.2.5. The only difference is that here the pipeline stages are inserted at optimal distances, where the highest possible throughput can be achieved. The encoding, decoding and completion detection of the $DualVmP$ circuits are the same as the ones presented in Section 3.1. Their gate-level implementations were shown in Figures 3.2 and 3.3. In the two pipelined links, one acknowledgment wire per 2-bit transmission is used to minimize the delay due to completion detection. With this configuration, completion detection is carried out per each 2-bit group at each pipeline stage and for all bits only at the

receiver and transmitter sides.

4.5 Simulation Results and Analysis

4.5.1 Wire Model

The wire properties are set according to ITRS 65nm technology node for global wiring [58]. The RLC values of the wires are extracted using field solvers from microstrip configuration. In the model, both wire width and separation distance are set to 210nm (two times minimum pitch for global wiring) and the wire thickness is set to 242nm. The resistance and inductance matrices are extracted using FastHenry [45], while the capacitance matrices are extracted using Linpar [46].

A number of data transmission wires with different bit widths and with both capacitive and inductive coupling are modeled to analyze the performance enhancement due to the proposed completion detection scheme. Parallel data transmission of 2, 4, 8, 16, 32 and 64-bits are considered. To implement these transmissions in the current sensing interconnects 2, 4, 8, 16, 32, 64 and 128 parallel wires with coupling are modeled and used in the simulation. Their acknowledgment transmission wire is assumed to be shielded at both sides with grounded wires and hence there is no coupling between the data and acknowledgment wires. In case of the two pipelined reference links 5, 10, 20, 40, 80, and 144 parallel wires with both capacitive and inductive coupling are modeled and used in the simulations. The required number of wires is different for the pipelined links since an acknowledgment is sent for each 2-bit group. Performance, power consumption and energy dissipation per bit of the interconnects are examined by varying the wire length from 1 to 5mm during simulations.

4.5.2 Simulations Setup

All simulations are carried out in Cadence Analog Spectre using 65nm CMOS technology from STMicroelectronics and 1V supply voltage. A number of simulations are carried out in order to analyze and compare the performance and power consumption of the presented on-chip interconnects including the references. In order to avoid confusion between the interconnects the following naming conventions are used. The current sensing two-phase 1-of-4 encoded

and two-phase dual-rail encoded differential links with the proposed completion detection are called *PMCmFCD* and *DualdiffFCD*, respectively. The current sensing two-phase 1-of-4 encoded and two-phase differential dual-rail encoded interconnects, which use the conventional completion detection technique are called *PMCm* and *Dualdiff*, respectively. The voltage-mode pipelined two-phase 1-of-4 and dual-rail encoded interconnects are called *1of4VmP* and *DualVmP*, respectively.

The *PMCmFCD* and *DualdiffFCD* interconnects are simulated for 2-bit transmission by varying the wire length from 1 to 5mm. They are also simulated for 2-, 4-, 8-, 16-, 32- and 64-bit transmission with the wire length of 2mm. The 2mm distance is chosen because there is a manufactured prototype in 65nm technology from Intel where the distance between neighboring tiles is 2mm [64]. Furthermore, *PMCm* and *Dualdiff* interconnects are redesigned and simulated in 65nm technology for 2-, 4-, 8-, 16-, 32- and 64-bit transmission and 2mm communication distance. The improvement in performance due to both fast completion detection technique and current sensing signaling was analyzed through design and simulations of the *1of4VmP* and *DualVmP* interconnects. The simulations are carried out for 2-, 4-, 8-, 16-, 32- and 64-bit transmission and 2mm long links. The distance between the pipeline registers is determined in such way that the interconnect achieves optimal throughput, that is, when the sum of encoder, driver and wire delays (both forward and backward) becomes greater than the delay of completion detection at the pipeline stage. It is 0.4mm for both 1-of-4 and dual-rail encoded transmission as determined from the simulation.

4.5.3 Performance Analysis

Latency and throughput are used as main parameters to analyze the performance of all the interconnects. The forward latency of the *PMCmFCD* and *DualdiffFCD* links are shown in Figure 4.9. This latency is the sum of the encoder, driver, wire, receiver, and completion detection delays. The *DualdiffFCD* link has a smaller forward latency and the difference in latency between the two links becomes larger with the increasing wire length. This is because the dual-rail link uses differential current sensing signaling which is faster than the single-ended current sensing signaling used in *PMCmFCD*. The backward latency is the same for both links since they use the same acknowledgment transmission interconnect. The backward latency is smaller than the forward




Figure 4.9: Latency of links with high-speed completion detection.

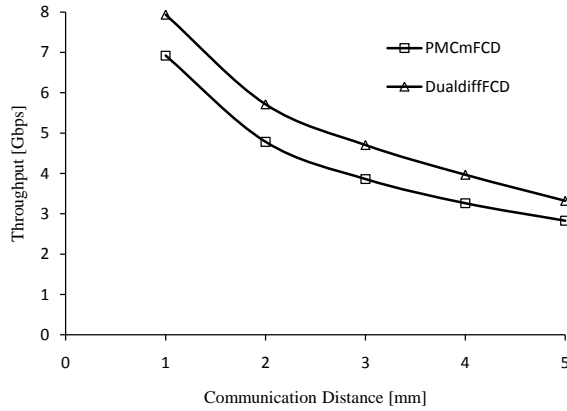


Figure 4.10: Throughput of links with high-speed completion detection.

latency since there is no encoding, decoding or completion detection involved.

The throughput of the *PMcMFCd* and *DualdiffFCD* links is analyzed for 2-bit transmission because it is the smallest possible transmission for a 1-of-4 encoded channel. The result is shown in Figure 4.10. The *PMcMFCd* and *DualdiffFCD* links achieve throughputs of 6.920Gbps and 7.936Gbps, respectively for the 1mm communication distance. When the communication distance increases to 3mm their throughputs become 3.788Gbps and 4.705Gbps, respectively, which indicates 45% and 41% reduction in throughput, compared to the 1mm case.

The major contribution of the high speed completion detection scheme comes into picture when the transmission bit width increases. The throughput of the two-phase 1-of-4 encoded links (*PMcMFCd*, *PMcM* and *1-of-4VmP*)




Figure 4.11: Throughput of two-phase 1-of-4 encoded 2mm interconnects.

is shown in Figure 4.11 for different bit widths and the wire length of 2mm. The difference in throughput between *PMCmFCD* and the other two becomes larger when the bit width increases, showing the advantage of the proposed high speed completion detection technique. For example, in the 64-bit case, throughput of *PMCmFCD* becomes 1.29 times that of *PMCm* and 2.07 times that of *1of4VmP*. The throughput gap between *PMCmFCD* and *PMCm* shows the performance improvement due to the high speed completion detection, because the difference between the two links is only in the implementation of completion detection. The difference between *PMCmFCD* and *1of4VmP* reveals the advantage of current sensing signaling along with the proposed completion detection technique.

In the case of two-phase dual-rail encoded interconnects, the *DualdiffFCD* interconnect achieves a higher throughput than the others as expected and the gap increases with the bit width. Its throughput for 64-bit transmission becomes 1.54 and 1.72 times that of *Dualdiff* and *DualVmP*, respectively. The throughput of these interconnects is shown in Figure 4.12. It can be seen that the current sensing interconnect with the presented fast completion detection technique is a better alternative. It improves the performance of the delay-insensitive links significantly, especially for wider bit-parallel transmissions, compared to the conventional implementation based on pipelined voltage-mode signaling.

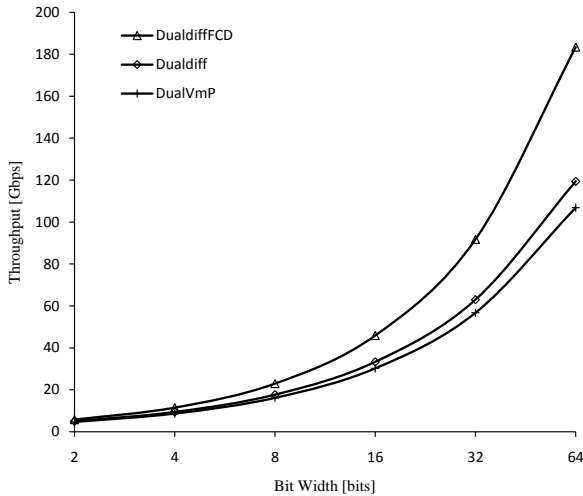


Figure 4.12: Throughput of two-phase dual-rail encoded 2mm interconnects.

4.5.4 Power Analysis

The average total power consumption of the three 1-of-4 encoded 2mm interconnects for different bit widths is shown in Figure 4.13. The *PMCm* interconnect consumes the highest power compared to the other two. For example, in 64-bit transmission *PMCm* consumes 10.6% and 20.1% more power than that of *PMCmFCD* and *1of4VmP*, respectively. Similarly, the dual-rail encoded interconnect with conventional completion detection, *Dualdiff*, consumes the highest power compared to the other two. It consumes 19.5% to 26.2% more power than that of *DualdiffFCD*. For 64-bit transmission, the *Dualdiff* interconnect consumes 32.2% more power than that of the *DualVmP*. *DualdiffFCD* consumes slightly more power than *DualVmP*, for instance, for 32- and 64-bit transmissions it consumes 6.9% and 4.8% more power compared to *DualVmP* power consumption. The power consumption of the three dual-rail encoded interconnects is shown in Figure 4.14.

In order to determine the power efficiency of these links, the energy per bit dissipated by the individual interconnects is examined. The energy per bit metric combines the power consumption and performance figures, which allows to judge the efficiency of these interconnects in a reliable manner. The energy per bit of the 1-of-4 and dual-rail encoded links with 2mm wire length and 2 - 64-bit transmission is shown in Figures 4.15 and 4.16, respectively. For 2- and 4-bit transmission, the conventional links, *1of4VmP* and *DualVmP* have a better power efficiency. When the bit-width is more than 4 bits, *Dualdiff*-




Figure 4.13: Power consumption of 1-of-4 encoded 2mm interconnects.




Figure 4.14: Power consumption of dual-rail 2mm encoded interconnects.

FCD is the most energy efficient and closely followed by the *PMCmFCD* link. Regarding the two voltage-mode reference links, energy efficiency deteriorates with increase in bit width because the completion detection circuit becomes an increasingly large tree of logic elements causing a significant power overhead. Hence, the current sensing delay-insensitive links with the proposed completion detection technique not only boost the performance of communication but also improve the power efficiency. This makes them most appropriate global interconnects in nanometer technologies where delay variations and power consumption are the major concerns.




Figure 4.15: Energy per bit dissipation of 1-of-4 encoded 2mm Links




Figure 4.16: Energy per bit dissipation of dual-Rail encoded 2mm Links

4.5.5 Area Comparison

Minimizing the area overhead of any module in a chip has been one of the concerns especially in the current giga-scale integration era. The link area, which consists of global wires and their signaling circuits, takes a significant portion of the total chip area. In this regard, it is wise to examine the areas of the proposed links and compare them with the conventionally implemented links. The areas of 2mm long 64-bit wide interconnects calculated from the link schematics, are presented in Table 5.3. The two current sensing links with high speed completion detection, *PMCMFCD* and *DualdiffFCD*, require a smaller silicon and metal area than their voltage-mode counterparts. The current sensing links require only 82% of the wiring area of the pipelined links.

The reason for this is that in pipelined links one acknowledgment wire is needed for every 2-bit transmission to minimize the delay incurred due to completion detection at each pipeline stage. Only one acknowledgment wire is required for the whole link in the current sensing links since completion detection needs to be performed only at the receiver. The current sensing links also require less silicon area compared to the pipelined voltage-mode ones because there are no area consuming pipeline stages.

Table 4.1: Area comparison of interconnects.

Link Name	Active Area [μm^2]	Wiring Area [μm^2]
PMCmFCD	508	55020
1of4VmP	704	66780
DualdiffFCD	599	55020
DualVmP	662	66780

4.6 Chapter Summary

A performance boosting technique by using a high speed completion detection circuit for delay-insensitive on-chip interconnects has been presented. Delay-insensitive data transfer is a necessity for global links in nanometer-scale technologies where delay variations are inevitable. The performance and power analysis shows that using the presented high speed completion detection technique improves throughput and power efficiency compared to the conventional implementations. It also requires less silicon and metal area. Therefore, using the presented technique leads to a realization of high throughput and power efficient global communication.

Chapter 5

Energy Efficient Semi-Serial Interconnect

Design and analysis of a high-throughput self-timed serial on-chip communication link is presented. Using fully bit-parallel interconnects that are presented in the previous two chapters for long-range communication links incurs considerable area overhead, routing difficulty, severe crosstalk noise and significant leakage power, making serial links a better alternative. The analysis between parallel and serial links in [103] and [104] shows the tradeoff between link length, latency, dynamic and leakage power as well as active and wiring area. For a given throughput the serial link is always preferable in terms of wiring area and incurs less routing congestion than parallel links. The serial link also takes smaller active area and consumes less leakage and dynamic power than the parallel link for long global communication [104].

In source-synchronous serial communication a clock is injected into the data stream at the transmitting side and the clock signal is recovered at the receiver side. Such clock-data recovery (CDR) circuits often require a power-hungry PLL, which may also take a considerable amount of time to converge on the proper clock frequency and phase at the beginning of each transmission. If the receiver and the transmitter operate in different clock domains, the transaction must also be synchronized at both ends, incurring additional delay and power consumption. One such link is presented in [103], it uses wave-pipelined multiplexed routing technique and its performance is limited by the clock skew and delay variations. In [34], circuits that had originally been used for off-chip communications [79, 80] were adopted to design a serial on-chip

link. It uses output multiplexed transmitter architecture due to its ability to deliver better performance than input multiplexing. However, this comes at the expense of much higher output capacitance that grows linearly with the bit-width. Both transmitter and receiver use multi-phase DLL circuits and clock calibration is required at the receiver side. A prototype chip has been fabricated in 180nm CMOS technology and a 3mm long link has achieved a throughput of 8Gbps. Total power consumption or energy per bit of this link is not reported. Another high-speed serial link was presented in [77], where the serializer/deserializer are based on a chain of MUXes. The link is single-ended and employs wave-pipelining. As a timing reference, constant delay elements are used instead of clock. Furthermore, the operation is based on the assumption that the introduced unit delays for the serializer and deserializer are the same. However, getting the same delay is almost impossible in the sub-100 nanometer CMOS technology due to considerable PVT variations, which in turn makes the reliability of communication using this approach questionable. A test chip was manufactured using 180nm CMOS technology and the measured throughput was 3Gbps. The power consumption or energy per bit dissipation of this link is not reported.

An alternative approach is to use self-timed communication which employs handshake instead of clocks. This enables robustness to delay variations through the use of delay-insensitive encoding and data transfer. A high data rate asynchronous bit-serial link for long-range on-chip communication is presented in [78]. It uses two-phase LEDR data encoding, fast asynchronous shift registers for both serializer and deserializer and wave-pipelined differential current-mode signaling. Due to direct integration of LEDR and differential signaling, this communication requires four wires per one link, increasing the required area and energy per bit of the bit-serial link. It achieves one-gate delay data bit cycle (67Gbps throughput in 65nm CMOS technology). It has been shown that a 4mm long link with 16-bits word length dissipates 150mW.

In today's SoCs power dissipation is a major design constraint that limits battery life and reliability, emphasizing the need for low-power on-chip communication links, which is the major motivation of this work. The argument is that it is possible to achieve 67Gbps throughput using one bit-serial link by restricting the data cycle to gate delay. However, the power dissipation of such a link is unacceptably high [78]. On the other hand, it is possible to achieve the same or even higher throughput with smaller power consumption by hav-

ing few bit-serial links in parallel which are designed from simple customized circuits and combination of techniques. Thus, the serial link presented in this chapter adopts the low-power approach.

This Chapter is organized as follows. In the next Section the need for long-range on-chip communication link in NoC is discussed. The proposed serial link communication protocol and detailed design of its circuits (serializer, deserializer, driver, receiver and data validity decoder) are presented in Section 5.2. Spice-level simulation results and analysis of the link performance, power, energy and area are discussed in Section 5.3. Comparison of fully bit-parallel, bit-serial and semi-serial links in terms of performance, energy and area is presented in Section 5.4. Finally, Section 5.5 presents the summary of this Chapter.

5.1 Long-Range Link in NoC

Most of NoC research has been focused on microarchitecture improvement and routing algorithms. However, selecting an appropriate topology is also one of the most critical decisions because it bounds critical performance metrics such as the network's zero-load latency and its capacity [66] and affect energy efficiency. The most common NoC topologies that have been used so far are 2-D mesh [63]. For example, the 80-node Intel's TeraFLOPS [64], the 64-node chip multiprocessor from Tiler [61], and the 167-processor computational platform [59] are implemented using 2-D mesh network. These networks have short wires in the architecture, but they have long network diameter. This causes energy inefficiency because of extra hops and furthermore consumes area. For instance, the 16-tile MIT RAW on-chip network consumes 36% of total chip power [65] and intel's TeraFLOPS link and routers consume 28% of the tile power [64].

There are NoC topologies which require long-range links such as torus (not folded-torus) [66], flattened butterfly [67], Spidergon [69], and concentrated mesh with replicated subnetworks and express channels [68]. It has been shown in [67] that using flattened butterfly with high-radix router offers lower latency and power consumption than 2-D mesh. That is, the latency and power consumption of flattened butterfly has been reduced by 28% and 38%, respectively compared to 2D-mesh. In [68], detailed area and energy models for different on-chip networks have been developed and their design

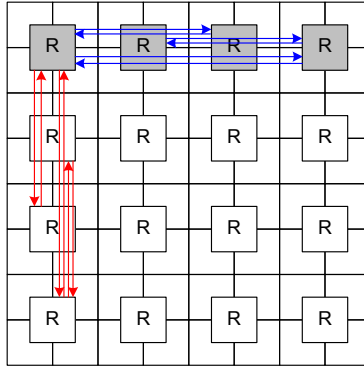


Figure 5.1: 64 nodes flattened butterfly topology [67]

tradeoffs are analyzed. It has been shown that concentrated mesh with replicated subnetworks and express channels provides a 24% improvement in area efficiency and a 48% improvement in energy efficiency over other topologies (mesh, folded-torus, concentrated mesh, fat-tree and tapered fat-tree). The express channel contributes 23% area and 38% energy efficiency. The area overhead is negligible because the express channels are routed over processor tiles in otherwise unused metal tracks and use preexisting router ports. The significant energy efficiency improvement is due to the decrease in completion time and due to the increased routing efficiency. That is, it is more efficient to route packets over express channels than through intermediate routers. Both flattened butterfly and concentrated mesh with express channels require long-range links which span more than one tile (see Figure 5.1 and 5.2). There are also other NoC topologies such as Spidergon [69], [70] and torus which are efficient and require longer links (Figure 5.3). Furthermore, in [71] topologies with fewer hops and longer channels have been proposed as promising solutions for energy and area efficient on-chip interconnection networks. It has also been demonstrated that adding few additional long-range links in a mesh network reduces the average packet latency significantly and improves the achievable throughput substantially [73]. Experiments involving real data traffic from telecom applications shows that the insertion of long-range links provides 36.3% improvement in critical traffic load, and 61.4% reduction in packet latency [73]. All these show the importance and need for high-speed and low-power long-range link, where its length spans two or more tiles.

Using fully bit-parallel communication in long-range NoC links that tra-




Figure 5.2: 64 nodes concentrated mesh with express channel [72]




Figure 5.3: 16 nodes NoC topologies requiring long-range links

verse two or more tiles becomes costly because it requires larger chip area, introduce routing difficulties, severe crosstalk noise and considerable leakage power (due to large driver/receiver to communicate through long lossy wires). Most of these issues can be addressed by using a high-speed serial link. Therefore, a high-throughput and low-power long-range serial on-chip link that can be used in NoC topologies which require long-range links inherently or when customized is presented. This in turn increases the overall network throughput and decreases the power consumption besides minimizing traffic congestion.

5.2 High-Throughput Serial On-Chip Interconnect

A high-throughput and low-power serial on-chip communication link employing integration of pulse dual-rail data encoding, wave-pipelining, pulse signaling and differential current-mode signaling is presented. Two-phase pulse dual-rail encoding is performed at low cost using two *AND* gates, one for data bit '1' and the other for '0'. This encoding enables usage of pulse signaling along with differential signaling directly. Furthermore, both the latency and the power consumption are reduced because data decoding logic is not needed at the receiver. The ability to detect each bit through pulse signaling in the wave-pipelined communication makes the link delay-insensitive and also enables acknowledging the transmission per word instead of per bit, improving throughput and saving energy.

In the presented serial link customized circuits and logics for serialization/deserialization and fastest possible stoppable local clock in the serializer are implemented. High-speed differential pulse current-mode signaling circuits are also designed. In addition, fast and robust data validity decoding circuits are designed. With this, one serial link achieves 9.09Gbps throughput. The serial link consists of serializer and deserializer, dual-rail encoder, driver, receiver and data validity decoder, as shown in Figure 5.4. In the subsequent subsections, the communication protocol, design details of the link circuits and the signaling technique will be discussed.




Figure 5.4: High-throughput serial on-chip communication link.

5.2.1 Communication Protocol

Similarly with the other interconnects, presented in this thesis, it is assumed that the sender and receiver modules have two-phase bundled-data interface. As soon as there is a request from the sender module which informs the data to be sent are ready and stable, the data will be loaded into the shift register.

In addition to the data, the *Stop* bit is also loaded which will be used to stop the shifting in the deserializer without the need for additional control logic such as data bit counter. The locally generated clock starts running after parallel data loading is completed. It is used for data shifting and dual-rail data encoding. It is a stoppable clock that runs only when there is data in the shift register to be transmitted and stopped at all other time, saving the communication power significantly. Data is shifted at the negative edge of the clock and encoded when the clock is in *high* state. The counter counts at the negative edge of the clock and signals the completion of data shifting when it reaches the maximum count value, which in turn stops the clock.

Dual-rail and differential pulse current-mode signaling is used for data transmission through the wire. Acknowledgment is sent per word instead of per bit thanks to the devised delay-insensitive wave-pipelining in the wire. In the receiving side, the transmitted data is retrieved directly from the receiver without the need for data decoding logic. The extracted data validity indicator is used as a clock for shifting the data in the deserializer. Shifting is performed at both edges of the data validity indicator signal. The arrival of *Stop* bit at the last flip-flop of the deserializer indicates that shifting is completed and the data are ready for parallel bit out. At this point, request to the receiving router will be sent. The deserializer shift register will be cleared when an acknowledgment from the data receiving module is received.

The overall communication protocol of the serializer is shown as a timing diagram in Figure 5.5. *Req2L* and *Ack2L* are the two-phase bundled-data request and acknowledgment signals of the sender. *Req_pulse* is used to enable parallel data loading in the shift register. *SRout* is the serializer's shift register data output and *Clk* is the locally generated clock. *Count* is the counter output signal and becomes *high* when all data are shifted out from the serializer's shift register. *Ack_pulse* is generated from *Count*, and it is used to stop the clock. It also saves communication time between data bursts by allowing data loading to be performed whilst waiting for acknowledgment to arrive from the receiving side. *Reset* is a locally generated signal which is used to reset the counter's registers besides allowing the clock to start running again by putting down the *Ack_pulse* to *low*.

The deserializer consists of a shift register and interfacing circuit between the serial link and receiving router. Its communication protocol is shown in Figure 5.6. The data receiver output *Wdout* is shifted in the deserializer shift




Figure 5.5: Serializer communication protocol.

register at both edges of data validity indicator signal, *DVIout*. The shifting process will be stopped when the *Stop* bit reaches to the shift register's last flipflop. *Req2R* and *Ack2R* are the bundled-data interface between the link and receiving block (Figure 5.4). *RstH* signal resets the deserializer's shift register.

5.2.2 Serializer and Pulse Dual-Rail Encoding

The bit parallel data from the sender is serialized using a novel shift register which uses the locally generated clock to shift the stored data. As shown in Figure 5.7, the serializer consists of shift register, counter, clock generating circuit and other interfacing elements. The design of the shift register is based on True Single-Phase-Clocked (TSPC) flip-flops [86] and customized to have parallel data loading ability. TSPC is chosen because of its ability to embed logic, parallel data loading in this case, with very little delay overhead. In addition, it has much smaller setup time and propagation delay compared to other dynamic flip-flops, making it the most suitable to realize high-speed shift registers. The customized TSPC circuit with parallel loading is shown in




Figure 5.6: Deserializer communication protocol.

Figure 5.8. In the loading phase, transistors Mns and Mnr are used to load bit '1' and bit '0', respectively and transistor Mps decouples D from node $L1$ (preventing error when D is '0' and data to be loaded is '1'). The tri-state weak inverter is used as a keeper for the loaded data. There are two 3-input upper asymmetric C-elements ($C1$ and $C2$) in the serializer circuit, shown in Figure 5.7, that are used to generate the local clock and keeper enable signals. The output of the two NOR gates act as the active-low reset signal for $C1$ and $C2$.

One-hot counter is designed from shift register so that its delay becomes equivalent to the data shift register in the serializer. As in the serializer's shift register, the counter shifts its one-hot code at the negative edge of the clock. Its shift register is designed from TSPC flip-flops which are customized to support active-low reset as shown in Figure 5.9. For N -bit word counter, N TSPC flip-flops are connected in series and the last flip-flop's output is inverted and feedback to the first flip-flop's input.

As already discussed in Chapter 1, the delay-insensitive data transfer, such as the dual-rail encoded interconnect, is a necessity in global interconnects of a nanometer SoC [87]. The delay-insensitivity makes the data transfer robust, because the sender and the receiver modules can communicate reliably regardless of delays in the transceivers and wires. Delay-insensitive data encoding technique requires $2N$ wires to transmit N -bit data. Pulse dual-rail encoding, where the presence of a new valid bit is represented by a pulse instead of voltage transitions or levels, is formulated and used in the presented serial link. This encoding enables straightforward use of pulse signaling. Furthermore,

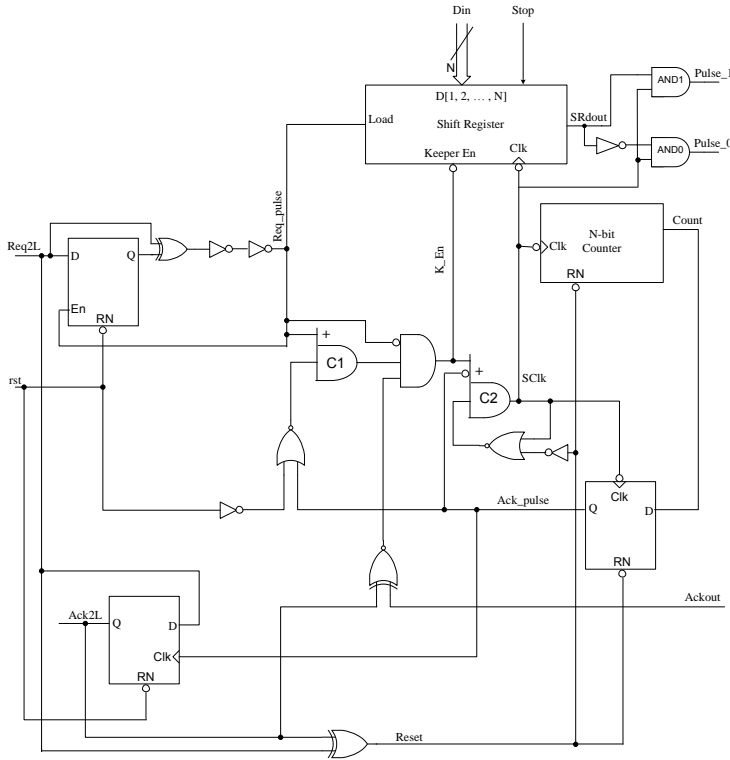


Figure 5.7: Serializer and pulse dual-rail encoder.

it has simpler and faster encoding/decoding logic when it is used along with differential signaling than the transition based protocols. When the clock is *high*, the dual-rail encoder, shown in Figure 5.7, encodes each bit into Pulse and No Pulse (P , NP) pair depending on its value. For example, when the output of the shift register is bit '1', and the clock is *high*, there is a pulse at the output of $AND1$ and no pulse at the output of $AND0$, as shown in Figure 5.10. Since there is no pulse in both wires between transmission of two consecutive bits, the receiver is able to detect each bit. That is, each bit can propagate at its own speed and can be detected reliably at the receiver regardless of the propagation delay variations.

5.2.3 High-Speed Differential Pulse Current-Mode Signaling

In pulse signaling only a small portion of the wire is charged during pulse propagation, significantly reducing the amount of capacitance need to be charged and hence, saving considerable amount of power over level-based signaling. It




Figure 5.8: Shift register’s TSPC flip-flop with parallel data loading.

has been shown that the use of pulse signaling can save up to 50% of energy compared to level-based signaling with repeater insertion [88]. Furthermore, it has been demonstrated through analytical models that more than 70% power saving could be achieved by combining pulse signaling with wave-pipelining technique without penalties of data throughput [89]. Since the main goal of this work is to achieve both high-speed global communication and low-power consumption, pulse signaling along with wave-pipelining is employed. In addition, differential current-mode signaling is used because of its high performance, better energy efficiency and noise immunity features [39]-[42]. Integration of dual-rail encoding and differential signaling has been realized using only two wires per link instead of four (two for dual-rail and two for differential signaling). This further reduces both power and required area of the link.

In addition to power saving, pulse current-mode signaling mitigates the effect of dispersion due to its return-to-zero signaling scheme in which sharp current pulses are used to transmit data and receiver termination is employed. To make use of these promising advantages, the wires need to be modeled with consideration of the lossy on-chip environment. Wider and thicker wires with larger spacing than the minimum is preferred to ease attenuation and preserve pulse integrity. This can be realized with smaller area overhead in a serial link than in parallel links.




Figure 5.9: Counter's resettable TSPC flip-flop.




Figure 5.10: Pulse dual-rail encoder input and output signals.

Driver

In this link, a source-coupled differential current-steering driver, shown in Figure 5.11, is used. It is fast because it has an extremely sharp transient response. The driver has also an advantage of reducing the AC component of the power supply noise because the circuit draws constant current from the supply when it is operating. This driver is naturally suited to drive a balanced differential pair of wires. The complementary outputs of the driver are attached to the two wires. The other end of the transmission is parallel terminated into a positive voltage. Depending on the output of the dual-rail encoder, in other words input to the driver, current will be steered in one of the wires from the current source. When bit '1' is transmitted, a voltage pulse drives the gate of *Mnp1* this in turn steers current pulse in *wire1* and no current in *wire0*. And when bit '0' is transmitted, the current pulse is steered

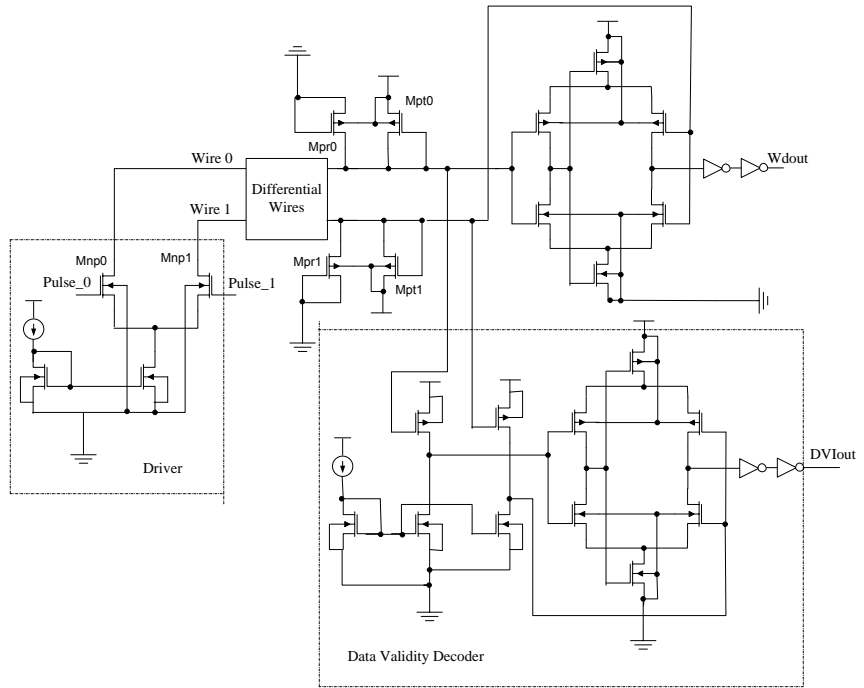


Figure 5.11: Driver, receiver and data validity decoder of serial link.

through $Mnp0$, which in turn steers the current pulse in $wire0$ and no current in $wire1$.

Receiver and Data Validity Decoder

The termination load and receiver design is shown in Figure 5.11. Diode connected $Mpt0$ and $Mpt1$ transistors are used as termination load. In addition to termination, they are also used to mirror the wire current which will be needed to decode out data validity indicator. The transconductance of these transistors has been regulated through the use of $Mpr0$ and $Mpr1$. The receiver needs to have high common-mode noise rejection capability in order to take full advantage of differential signaling. Due to this, a high-speed self-biased differential amplifier is used. The differential amplifier used in this design has less sensitivity to process, temperature and supply voltage variations. It operates at high speed because its output switching currents are significantly greater than its quiescent current. Furthermore, the adopted amplifier has higher differential-mode gain than conventional amplifiers and a large common-mode input range because its bias condition adjusts itself to accommodate the input swing [86, 90].

In delay-insensitive transmission, decoding of data and data validity indicator at the receiving end is necessary. The transmitted data is received and decoded out directly in the receiver without the need for separate data decoding logic. This is due to the novel integration of pulse and differential signaling. The remaining issue is data validity indicator, which will also be used as a clock to shift the data bit in the deserializer. From the encoding, it is known that there will be current only in one of the wires when there is valid bit transmission and no current in both wires between two consecutive bit transmissions. Each wire's current is compared with a reference current using a current comparator and the output of the two current comparators is fed to a differential amplifier. The output of the differential amplifier is the data validity indicator (*DVIout*). This way of completion detection makes the communication robust to both delay variations and noise because it takes into account both wires and the used differential amplifier, which has a high common-mode noise rejection ratio. Both edges of *DVIout* signal indicate the availability of valid and new data at the receiver output. The circuit of the data validity decoder is shown in Figure 5.11.

5.2.4 Deserializer

The deserializer consists of a shift register and interfacing circuit (between the receiving module and the deserializer) as shown in Figure 5.12. In shift register data is shifted out at both edges of the *DVIout* signal. The shift register is designed from double-edge-triggered flip-flops. This flip-flop is designed by tying together the outputs of a negative and a positive edge-triggered TSPC flip-flops, obtaining multiplexer function for free. It stores dynamically during opposite clock phases and drives its output actively on both clock edges. The circuit of a double-edge-triggered flip-flop is shown in Figure 5.13. *Mnrs1* and *Mnrs2* transistors are used for resetting the flip-flop. When *Stop* bit reaches the last FF, the shifting will be stopped and data can be read out in parallel. The bundled-data two-phase request signal for parallel data receiving module is generated using a D-FF as shown in the interfacing circuit (Figure 5.12). As soon as an acknowledgment is received, the deserializer's shift register will be cleared (resetted).

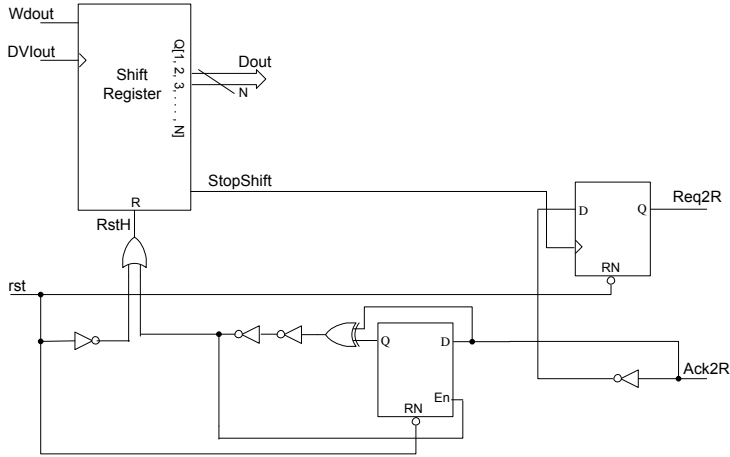


Figure 5.12: Deserializer circuit.

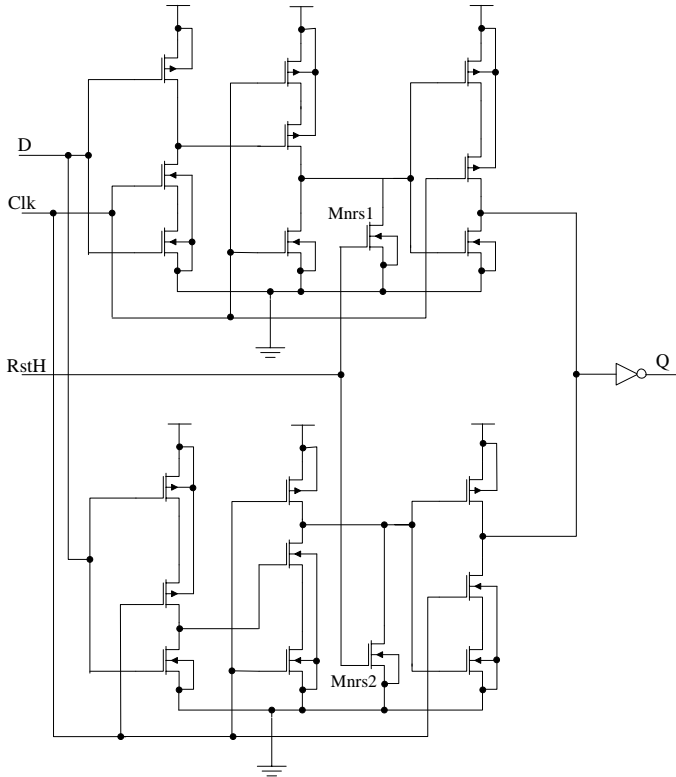


Figure 5.13: Double edge-triggered TSPC flip-flop.

5.2.5 Acknowledgment Transmission

As already discussed, acknowledgment is sent from the receiver per word instead of per bit. The same signaling technique as in data transmission is used

except that there is no wave-pipelining, as it is not necessary (there is only one bit to transmit at a time). Since the receiver has a two-phase handshaking interface, a pulse is generated and transmitted at each transition edge of the acknowledgment signal. The pulse generator circuit, shown in Figure 5.14, generates a pulse for low-to-high and high-to-low transitions. Its driver and receiver circuits are similar with the data transmission circuit, shown in Figure 5.11, but the wires are narrower. This is due to the fact that the performance of acknowledgment transmission is not detrimental for the throughput of the link because the transmission takes place once per word.




Figure 5.14: Pulse generator for acknowledgment signal transmission.

5.3 Simulation Results and Analysis

The performance, power consumption, and energy per bit of the presented serial link is discussed in this section. In addition to the bit-serial link, semi-serial (two, four and eight bit-serial) links are simulated and analyzed. In the bit-serial case, simulation is carried out for different wire lengths (1 to 8 mm) and in case of semi-serial, the simulation is performed for 4mm long communication. All simulations are performed using 65nm CMOS technology from STMicroelectronics with supply voltage of 1 V. Depending on the link circuits operating condition requirements, low-power low-vt or low-power high-vt transistors are used.

5.3.1 Wire Model and Simulation Waveforms

A distributed RLC-model is adopted to accurately model signaling over long on-chip wires. Furthermore, both capacitive and inductive coupling is added between wires to take into account crosstalk noise. The wire properties were set according to ITRS 65nm technology node for global wiring [58]. In the serial data transmission wire modeling, wide and upper metal layers are assumed. Its wire width and separation distance were set to $1\mu\text{m}$ and $1.5\mu\text{m}$,

respectively. In the acknowledgment wire modeling, both wire width and separation distance were set to 210nm. The RLC values of the wires were extracted using field solvers, FastHenry [45] and Linpar [46].

The simulation waveforms of major signals of the serial link are shown in Figure 5.15. The two-phase bundled-data handshake signals of the two communicating parties are *Req2L*, *Ack2L* and *Req2R* and *Ack2R* (see Figure 5.4). The parallel load enable signal of the serializer's shift register is *Req_pulse*. The locally generated stoppable clock is *Clk* and the serialized data output of the serializer is *SRout*. The pulse dual-rail encoder outputs (the differential driver inputs) are *Pulse_1* and *Pulse_0*. *Wdout* and *DVIout* are the receiver and data validity decoder outputs, respectively. Signal which informs the receiving router the availability of all data for parallel output is *StopShift*.

5.3.2 Performance

The throughput of the presented bit-serial link is 9.091Gbps (110ps bit cycle) for all simulated link lengths and it is limited by the capacity of the clock generation circuit. It is known that the fastest clock that can be generated using a ring oscillator is bounded to $6 - 8\tau_4$ [80], resulting in 90 to 125ps clock period for 65nm technology. Bit shifting is performed at the negative edge of the clock and pulse dual-rail encoding during *high* state of the clock. The delay between consecutive data words is minimized since data loading to the shift register is performed while waiting for the acknowledgment to arrive from the receiver. As soon as the acknowledgment signal arrives, clock starts running and then shifting out the bit for transmission. The throughput of the semi-serial link increases linearly with the number of bit-serial parallel links. For instance, from eight parallel bit-serial links, a throughput of 72.728Gbps was achieved.

5.3.3 Power and Energy Consumption

The bit-serial link is simulated for 32-bit word and the length of the link is varied from 1 to 8mm. Its overall (including all link circuits) average power consumption and energy per bit is listed in Table 5.1. The power consumption of 32-bit serializer and deserializer is 2.198mW and 1.416mW, respectively. The power consumption of the link and its energy dissipation per bit do not increase steeply with the wire length because of the pulse signaling, where




Figure 5.15: Simulation waveforms of serial link.

only a small portion of the wire needs to be charged. In addition, due to wider spacing of the wires the link has smaller coupling capacitance which allows the required driving current to be smaller and this in turn reduces the power consumed by the link.

It is known that if N number of the presented bit-serial links are used in parallel, the overall throughput of the channel becomes N times the throughput of one bit-serial link. But what about the power efficiency or energy dissipation per bit. In order to answer this question, a semi-serial link is designed and simulated. This link consists of eight bit-serial links each of which has 8-bit word length. The simulation results of the semi-serial link are then

Table 5.1: Power and energy dissipations of bit-serial communication.

Link Length [mm]	Average Total Power [mW]	Energy per Bit [pJ/bit]
1	4.231	0.465
2	4.442	0.488
3	4.623	0.508
4	4.785	0.526
5	4.947	0.544
6	5.088	0.559
7	5.203	0.572
8	5.317	0.584

compared with 64-bit one bit-serial link. The wire length is varied from 1 to 8mm for both semi-serial and bit-serial links simulations. As shown in Figure 5.16, the semi-serial link has smaller energy dissipation than the bit-serial link. To be more precise, the energy dissipation of the semi-serial link is less than one-third of the bit-serial link. There are two reasons for this: first, proportional increase in channel bandwidth is much higher than the increase in power consumption. Second, equal amount of power is dissipated in both links due to serializer/deserializer and their control circuits, because the size of the transmitted word is the same and there is no need to replicate the control and clock generating circuits. Only proper buffering is required since the locally generated clock is also used for pulse dual-rail data encoding. In addition, even if the semi-serial link has eight wires, the dynamic power consumption on the wire is reduced significantly due to the use of both pulse signaling and wave-pipelining. Therefore, semi-serial link is a better alternative for long-range high-throughput and energy efficient on-chip communication.

Assuming that the sending and receiving modules have 64-bit wide data and they are placed 4mm apart each other, three different semi-serial links are simulated. The energy per bit of these links is presented in Table 5.2. While the throughput doubles and quadruples with the number of parallel links, the power consumption (energy) does not. The reason is that the control and clock generating circuits of the serializer are shared, and there is no need to replicate them for each link. This decreases the energy dissipation per bit of the semi-serial link as the number of parallel bit-serial links increases (Table 5.2). In case of eight parallel links, a throughput of 72.728Gbps is achieved




Figure 5.16: Energy per bit of 64-bits word bit-serial and semi-serial links.

with 16.596mW power consumption. If this result is compared with [78], which achieves a throughput of 67Gbps with a power consumption of 150mW for 16-bit word and 4mm long bit-serial link, the power consumption of the proposed semi-serial link is almost one-tenth of [78] besides achieving a slightly higher throughput.

Table 5.2: Performance and energy dissipation of semi-serial links.

Links	Throughput [Gbps]	Power Consumption [mW]	Energy per Bit [pJ/bit]
2	18.182	7.958	0.437
4	36.364	11.730	0.322
8	72.728	16.419	0.226

5.4 Fully Bit-Parallel vs Serial Links

With increasing number of complex, non-uniform sized nodes in a NoC, high-throughput low-power and area efficient long-range links become necessity. For example, the node size is 2mm by 1.5mm the same as in TERAFLIPS [64] and the NoC consists of 20 nodes. To connect the farthest nodes in regular mesh structure or for end-around channel of a torus an 8mm long link is required. In order to analyze the trade-off between bit-serial and fully bit-parallel long-range channel of a NoC, fully bit-parallel link with optimal repeater insertion is designed as it is a conventional way of global on-chip interconnect implementation. Also, a semi-serial link consisting of eight bit-serial links in parallel is designed and simulated.

The presented serial link is delay-insensitive, and therefore the optimally

repeated bit-parallel link is also designed as a delay-insensitive link using two-phase LEDR encoding. The LEDR encoding is chosen due to its simpler and faster completion detection and data decoding logics than two-phase dual-rail encoding. In the fully bit-parallel link, the wires were modeled narrower than the wires in the serial link to minimize the wiring area overhead. That is, the pitch of the wire is two times the minimum pitch of a global wire according to ITRS [58].

To illustrate real life applications, 64-bit wide data transmission is used in all links. As can be seen from Figure 5.17, the throughput of a fully bit-parallel link is greater than both bit-serial and semi-serial links for 1mm and 2mm long communication distances. However, its throughput decreases rapidly with the increase in communication distance, and it becomes less than the semi-serial link throughput starting from 3mm. Consider a NoC, which consists of 2mm by 2mm nodes. In such a network, the semi-serial link gives higher throughput than the fully bit-parallel one when communicating nodes are two or more hops away. For 8mm long communication, the semi-serial link achieves 1.97 times the throughput of a bit-parallel link. The throughput of the fully bit-parallel link is considerably affected by the time required to carry out completion detection, which needs to be done to maintain its delay-insensitive behavior.

The energy dissipated per bit transmission of these three links is shown in Figure 5.18. The semi-serial link dissipates the least energy at all lengths of the wire and the bit-parallel link dissipates more energy than the two links when wire length is 4mm or longer. The energy dissipation of the bit-parallel link rises sharply with wire length, because the required number of repeaters and dynamic power consumption due to the wire capacitance increases. It consumes 15.3 times more energy than the semi-serial link for 8mm long communication. Hence, the semi-serial link is a better alternative than both bit-serial and bit-parallel links for long-range NoC links since it gives higher throughput with smaller energy dissipation.

The required active and the wiring areas of the three 64-bit word links with 8mm long transmission are shown in Table 5.3. The active area taken by the fully bit-parallel link is much larger than the other two. That is, it is 75 times larger than the semi-serial link area. This is because the LEDR encoder and the repeaters take 98% of the total active area of the bit-parallel link. One-bit encoder consists of three double-edge triggered flip-flops, five inverters and one XOR gate. The repeater size is 96 times the minimum sized

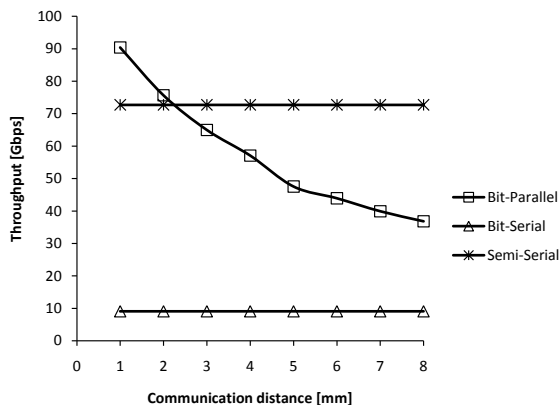


Figure 5.17: Throughput of 64-bit word serial and parallel links.




Figure 5.18: Energy per bit of 64-bit word serial and parallel links.

inverter. The wiring area of the fully bit-parallel link is also larger than the other two. The semi-serial link takes only 73% of the wiring area compared to the bit-parallel link. Thus, the bit-parallel link costs much more chip area than both the semi-serial and bit-serial links.

Table 5.3: Area comparison between serial and parallel links.

Link Type	Active Area [μm^2]	Wiring Area [μm^2]
Bit-Serial	84	33040
Semi-Serial	91	313040
Bit-Parallel	6854	428400

5.5 Chapter Summary

In this chapter design and analysis of a high-throughput and low-power serial on-chip communication link is presented. The combination of pulse dual-rail encoding, wave-pipelining, pulse signaling and differential current-mode signaling besides customization of serializer/deserializer circuits leads to a realization of high-throughput serial link with low power consumption. This link is a promising candidate for long-range NoC channels, which are needed inherently due to topologies or through customization of regular 2D networks. In addition, its delay-insensitive data transfer makes it more appropriate for nanoscale NoC interconnects where delay variations are inevitable. A 32-bit word 8mm long bit-serial link achieves a throughput of 9.09Gbps with 0.58pJ/bit energy dissipation while a 64-bit word 4mm long semi-serial link achieves a throughput of 72.72Gbps with 0.22pJ/bit energy dissipation. A semi-serial link which consists of eight of the presented bit-serial links outperforms the fully bit-parallel link both in throughput and energy efficiency as well as active and wiring area when the communication distance is 4mm and longer.

Chapter 6

Comparison of the Designed Interconnects

In the previous three chapters (3, 4, and 5), design and analysis of delay-insensitive and high-performance on-chip interconnects have been presented. These interconnects are suitable for any kind of point-to-point on-chip communication, such as in a SoC to connect nearby or far away system blocks and in a NoC between two routers. The purpose of this chapter is to make a generalized summary of the presented interconnects as well as comparisons between them. In order to do so, all interconnects are redesigned and simulated in 65nm CMOS technology from STMicroelectronics with 1V supply voltage.

6.1 Summary of the Interconnects

The presented four interconnects use different encoding/decoding, completion detection and signaling techniques. Each approach has its own advantages and limitations but all have the same goals: delay variation robustness, high performance and energy efficiency. The generalized summary of these interconnects is presented in Table 6.1. *LEDRCm*, *PMCmFCD*, and *DualdiffFCD* interconnects are designed for fully bit-parallel transmission. Whereas the *Bit-Serial* interconnect is for serial transmission of bits and the *Semi-Serial* is made up of few bit-serial interconnects in parallel. In *Bit-Serial* and *Semi-Serial* interconnects the bits are wave-pipelined on the wire and acknowledgment is transmitted per word.

Table 6.1: Summary of interconnects

Interconnect	Encoding	Signaling	Advantages	Limitations
<i>LEDR_{Cm}</i>	Level encoded dual-rail (LEDR)	Binary current sensing	No resetting transitions. Level based detection rather than transition. Better performance and energy efficiency due to current sensing signaling.	It is not possible to use the proposed fast completion detection scheme and to shut wire currents off in order to save power
<i>PM_{Cm}FCD</i>	Two-phase 1-of-4	Pulsed Multilevel current sensing	No resetting transitions. Performance enhancement due to current sensing signaling and fast completion detection. Power saving due to shutting off wire currents.	Extra design effort is required to determine the optimal current margins
<i>DualdiffFCD</i>	Two-phase dual-rail	Multilevel and differential current sensing	No resetting transitions. Area and power efficient integration of dual-rail encoding and differential signaling. High common-mode noise rejection. Performance and energy efficiency improvement due to differential current sensing signaling and fast completion detection.	Extra design effort is required to determine the optimal current margins
<i>Bit-Serial Semi-Serial</i>	Pulse dual-rail	Wave-pipelining, pulse signaling, and differential current mode signaling	Fast encoding. No data decoding logic. Per word acknowledgment without losing bit-level delay insensitivity. High common-mode noise rejection. High throughput and energy efficient due to integration of signaling techniques as well as circuit optimizations.	Not energy efficient for short-range links

6.2 Comparison of the Interconnects

In this section, comparison between *LEDRCm*, *PMCmFCD*, *DualdiffFCD*, *Bit-Serial*, and *Semi-Serial* interconnects in terms of performance, energy and area will be carried out for a number of transmissions with different bit widths. Except in the case of the serial link, acknowledgment signal transmission is performed using the same signaling scheme and circuits (the same as the one presented in Section 4.3.3). In the serial communication, acknowledgment is sent per word using differential pulse current-mode signaling. In order to have a proper comparison between the bit-parallel interconnects, their drivers are designed so that they have the same output current values. Since the *LEDRCm* interconnect uses binary signaling while the other two use multilevel current sensing signaling, the driver output current of *LEDRCm* is set to one of the current levels (I_2) of the other two interconnects.

6.2.1 Performance

In the *LEDRCm* interconnect the proposed high-speed completion detection technique is not implemented, because data validity decoding can be done only by detecting transitions, as LEDR encoding is based on data phase and state. The *PMCmFCD* and *DualdiffFCD* interconnects use the high-speed current-mode completion detection circuit. The throughputs of the three interconnects are listed in Table 6.2 for 2-bit transmission and 1-5 mm long communication distances. Two-bit transmission is considered in order to have a proper comparison with the 1-of-4 encoded link, *PMCmFCD*, since the smallest possible transmission in the 1-of-4 encoded link is two bits. When the communication distance increases the throughput of all the interconnects decreases. The *DualdiffFCD* interconnect achieves the highest throughput among the three interconnects because of its differential signaling scheme. The throughput of *PMCmFCD* interconnect is higher than that of the *LEDRCm* link. The throughput of bit-serial and semi-serial links is not affected by the wire length as they use wave-pipelining, and an acknowledgment is sent per word instead of per bit. The *Bit-Serial* link achieves a throughput of 9.09Gbps and the throughput of *Semi-Serial* link increases linearly with the number of parallel bit-serial links.

The current sensing interconnects have also been designed and simulated for a number of different bit widths from 2 to 64 bits. The communication

Table 6.2: Throughput of interconnects.

Communication Distance [mm]	Throughput [Gbps]		
	<i>LEDRCm</i>	<i>PMCmFCD</i>	<i>DualdiffFCD</i>
1	5.988	6.92	7.936
2	4.662	4.784	5.714
3	3.809	3.861	4.705
4	3.205	3.262	3.968
5	2.597	2.828	3.322

distance is assumed as 2mm for all the interconnects. Their throughput is shown in Figure 6.1. For 2- and 4-bit transmissions there is no big difference between the throughputs. Starting from 8-bit transmission the current sensing interconnect’s throughput becomes higher than the pipelined voltage-mode links. The throughput gap between the current sensing and pipelined voltage-mode links increases with the bit width. For example, the throughput of *DualdiffFCD* is 1.4 and 1.7 times the throughput of *DualVmP* for 8- and 64-bit transmissions, respectively. The *DualdiffFCD* link achieves the highest throughput because of the differential signaling, followed by the *PMCmFCD* link. In 64-bit transmission the throughput of *DualdiffFCD* is 1.19 and 1.59 times *PMCmFCD* and *LEDRCm* links throughput, respectively. For 32- and 64-bit transmissions the *LEDRCm* link’s throughput is considerably lower than that of the two current sensing links due to the difference in completion detection.

The purpose of the serial link is to communicate blocks of data over a long distance. Therefore the comparison with the bit-parallel current sensing links is carried out for 64-bit 5mm long transmission. Among the serial communication links one bit-serial and three semi-serial links are considered. The *Semi-Serial8*, *Semi-Serial12*, and *Semi-Serial16* links in Table 6.3 are links consisting of 8, 12, and 16 bit-serial links in parallel, respectively. The throughputs of these interconnects are presented in Table 6.3. Compared to the *Bit-Serial* and *Semi-Serial8* links, the two bit-parallel links perform better.

6.2.2 Power Efficiency

The energy per bit dissipation of *LEDRCm*, *PMCmFCD*, and *DualdiffFCD* interconnects are determined for 1 to 5mm long transmissions. As it can be




Figure 6.1: Throughput versus bit width of links

Table 6.3: Throughput of 64-bit word 5mm long links.

Link Name	Throughput [Gbps]
<i>PMCmFCD</i>	90.523
<i>DualdiffFCD</i>	106.312
<i>Bit-Serial</i>	9.09
<i>Semi-Serial8</i>	72.72
<i>Semi-Serial12</i>	109.08
<i>Semi-Serial16</i>	145.44

seen in Figure 6.2, the *LEDRCm* interconnect dissipates the highest amount of energy per bit for all communication distances. One of the reasons for this is that it uses voltage-mode completion detection which contains a large number of power-consuming logic gates. The *PMCmFCD* interconnect dissipates lowest energy per bit at 1 and 2mm long communications and after that its energy consumption becomes higher than that of the *DualdiffFCD* interconnect. This is due to the single-ended signaling used in the *PMCmFCD* interconnect.

As the purpose of these links is to provide transfer of data between two points at a global distance with minimum possible energy, their energy per bit consumption is analyzed for 2 to 64-bit 2mm long transmission. Besides the three current sensing links, the conventional two-phase dual-rail and 1-of-4 encoded and optimally pipelined voltage-mode interconnects' energy dissipation per bit is examined and presented. The *LEDRCm* interconnect dissipates the highest energy per bit and the *DualdiffFCD* link dissipates the least energy




Figure 6.2: Energy per bit versus communication distance of interconnects




Figure 6.3: Energy per bit versus transmission bit width of links

starting from 4-bit transmission (Figure 6.3). For instance, 64-bit transmission using *LEDRCm* interconnect dissipates 0.507pJ/bit which is 2.96 times higher than the energy per bit dissipation of *DualdiffFCD* link. For 2-bit transmission the pipelined voltage-mode interconnects are preferable, consuming the lowest energy. Starting from 8-bit transmission both *PMCmFCD* and *DualdiffFCD* links consume the lowest energy, thanks to the current-mode completion detection scheme.

The *Bit-Serial* and *Semi-Serial8* links dissipate 0.886pJ/bit and 0.190pJ/bit energy, respectively for 64-bit 2mm long transmission. The semi-serial link dissipates almost the same energy as the *PMCmFCD* link and 11.1% more energy compared to the *DualdiffFCD* link.

6.2.3 Area

In today’s interconnect-centric and ultra integration era, area is among the fore front design parameters. Comparison of silicon and wiring area is performed for the *LEDRCm*, *PMCmFCD*, *DualdiffFCD*, *Bit-Serial* and *Semi-Serial8* interconnects for 64-bit 2mm long communication. The active area taken by the *LEDRCm* is 237% and 186% more than the area of *PMCmFCD* and *DualdiffFCD* interconnects. The reason is that one bit LEDR encoding requires three double-edge triggered flip flops and one 2-input XOR gate. In addition, 63 2-input C-elements are needed for the completion detection. The active area required by *PMCmFCD* and *DualdiffFCD* links is also smaller than the pipelined voltage-mode links but larger than the serial links. The *LEDRCm*, *PMCmFCD* and *DualdiffFCD* links take only 82% and 58% wiring area of the pipelined voltage-mode and *Semi-Serial8* links, respectively (Table 6.4).

Table 6.4: Area comparison of 64-bit 2mm long transmission links.

Link Name	Silicon Area [μm^2]	Wiring Area [μm^2]
<i>LEDRCm</i>	1714	55020
<i>PMCmFCD</i>	508	55020
<i>DualdiffFCD</i>	599	55020
<i>1of4VmP</i>	704	66780
<i>DualVmP</i>	662	66780
<i>Bit-Serial</i>	84	8260
<i>Semi-Serial8</i>	91	78260

6.3 Chapter Summary

Generalized summary of the three bit-parallel and serial interconnects that were presented in the previous three chapters has been discussed. Comparisons of performance, energy per bit and area have also been carried out. Among the bit-parallel interconnects dual-rail encoded differential current sensing interconnect achieves the highest throughput with lowest energy per bit dissipation. The LEDR encoded current sensing interconnect has the poorest performance and highest energy per bit dissipation. The dual-rail differential current sensing interconnect takes slightly larger active area. The semi-serial link outperforms the bit-parallel ones starting from 5mm communication distance in terms of throughput, showing its potential for long-range communication.

Chapter 7

Circuit Techniques for PVT Variation Tolerance

As part of an integrated circuit, on-chip interconnects experience two types of variations: physical and environmental. A physical variation is due to the manufacturing process imperfections. Whereas, environmental variations occur during the operation of a circuit and includes dynamic variations in the supply voltage and temperature. Precise control of the manufacturing process is worsening with technology scaling due to smaller dimensions, smaller number of doping atoms and aggressive lithographic techniques. This becomes a major concern since it causes uncertainty in electrical characteristics of devices and interconnecting wires which consequently affect the reliability of the system. Variability in the operating environment also affects the reliability of on-chip interconnects. As the variations increase, techniques which reduce their impacts while providing the highest performance for a given power constraint are necessary at the system, architecture, and circuit levels [111]. In this chapter, circuit level techniques which ensure signal integrity of a current sensing on-chip interconnect in the presence of PVT variations are developed and implemented. Since all the interconnects that are presented in this thesis are delay-insensitive, the developed signal integrity technique considers only the signal amplitude variation.

This chapter is organized as follows. In the next section, signal integrity problems of a current sensing interconnect due to process and environmental variations are discussed. Brief discussion about post-manufacture variation adaptation technique is presented in Section 7.2. Process variation tolerance

technique along with its algorithm, methodology, and circuit realization is presented in Section 7.3. The runtime environmental variation monitoring and management technique is discussed along with its implementation in Section 7.4. Simulation results of the presented PVT variation tolerance techniques as well as analysis of power, delay and area overheads are presented in Section 7.5. The summary of the Chapter is presented in the last Section.

7.1 Signal Integrity of Current Sensing Interconnect

For convenience a current sensing interconnect can be divided into three parts: driver, wire, and receiver. Figure 7.1 shows a current sensing interconnect structure along with its electrical parameters which can be affected by PVT variations. In a current sensing interconnect the receiver compares wire current with a reference in order to retrieve the data transmitted from the far-end. If the variation of input and/or reference current is out of the allocated margin then it can lead to erroneous output. It can be possible to allocate large current margins by considering worst-case variations, however; this has power consumption costs associated, especially with the increase in the number of variability sources. Thus, it is wise to deal with variations in these two currents and devise techniques at the circuit level, which can tolerate their process and environmental induced variations thereby enhancing the reliability of the interconnect with low power overhead.

7.1.1 Effects of Process Variation

Both wire and reference currents may deviate from their nominal values due to uncertainties in front-end and back-end manufacturing processes. The front-end process comprises of manufacturing steps that are involved in creating




Figure 7.1: Variable parameters of current sensing interconnect


devices, while the back-end is responsible for creating the interconnecting wires between the devices. The primary causes of variations in device electrical parameters are threshold voltage (V_{th}) variation, line-edge roughness (channel length and width variations), oxide thickness variation and dopant fluctuations [114], [115], [116]. These variations make the output current of the driver I_{win} to be different from its nominal value (Figure 7.1).

It is known that sub-100nm CMOS transistors are velocity saturated, i.e., there is a linear dependence between I_D and V_{GS} in the strong-inversion region [126]. Also, threshold voltage is strongly impacted by channel length and operational voltage V_{DS} . The output current of the driver I_{win} under velocity saturation can be expressed by the Equation 7.1 [126], [127]. E_C is the critical electric field at which the carrier velocity becomes saturated. From this equation it can be seen that variations in threshold voltage, channel length and width and/or oxide thickness make I_{win} to fluctuate from its intended value.


$$I_{DSat} = \frac{W}{L} \frac{\mu_{eff} C_{ox}}{2} V_{DSat} (V_{GS} - V_{TH}) \quad (7.1)$$

$$\text{where } V_{DSat} = \frac{(V_{GS} - V_{TH}) E_C L}{(V_{GS} - V_{TH}) + E_C L} \quad \text{and} \quad C_{ox} = \frac{\epsilon_{ox}}{t_{ox}}$$


In order to examine the effect of front-end process variations on signal integrity of *LEDRCm*, *PMCmFCD*, and *DualdiffFCD* interconnects, Monte Carlo process runs of 1000 are carried out using 65nm technology statistical model from STMicroelectronics. The communication distance is assumed to be 2mm long. The variation of output current of (*LEDRCm*) interconnect driver is shown in Figure 7.2(a). In this simulation, only the front-end variability in the data encoder and the driver is considered. Its worst-case variation from its mean is $20\mu A$. The effect of electrical parameter variations of the termination transistor on I_{win} is also simulated and shown in Figure 7.2(b). Variability in the termination load causes additional I_{win} variation, taking the total I_{win} variation to $40\mu A$. The variation in the receiver's input current $I_{rec,in}$ due to fluctuations in the data encoder, driver and termination transistors is shown in Figure 7.2(c). Its worst-case variation is about $50\mu A$, and hence requiring a current margin greater than that, in between $I_{rec,in}$ and receiver's reference currents.



(a) I_{win} variation due to encoder and driver device parameters variability.



(b) I_{win} variation due to encoder, driver and termination device parameters variability.



(c) $I_{rec,in}$ variation due to encoder, driver and termination device parameters variability.


Figure 7.2: LEDRCm interconnect I_{win} and $I_{rec,in}$ variations

The variation of *PMCMFCD* I_{win} considering the process variations of the data encoder and the driver devices is shown in Figure 7.3(a). In this case its I_{win} worst-case variation is $40\mu\text{A}$. Variations in its I_{win} and $I_{rec,in}$ when the termination transistor effect is included are shown in Figures 7.3(b) and 7.3(c), respectively. Worst-case variation of I_{win} is $40\mu\text{A}$, while $I_{rec,in}$ is about $50\mu\text{A}$.


The *DualdiffFCD* interconnect's I_{win} and $I_{rec,in}$ variations due to front-end variabilities were also examined. Its I_{win} variation due to transmitter side device parameters uncertainties is shown in Figure 7.4(a) and its worst-case variation was about $50\mu\text{A}$. Variations of I_{win} and $I_{rec,in}$ due to manufacturing variabilities of encoder, driver and termination transistors are shown in Figures 7.4(b) and 7.4(c), respectively. I_{win} variation increases by $25\mu\text{A}$ due to the termination transistor. The current variations in this interconnect were larger than those in *PMCMFCD* and *LEDRCm* interconnects.

The fluctuations in the back-end processes cause variations in geometry and material properties of the wire structure. Studies show that among the back-end process steps, erosion and dishing during chemical-mechanical polishing (CMP) process has strong impact on wire parasitics. This is due to the systematic pattern or spatial effects (metal density, width and space) [119]. In general, dishing strongly affects wide lines, while erosion is worse for narrower oxide and dielectric spacing between lines. In medium size features, the two effects combine, so that both dishing and erosion contribute to overall copper thickness reduction. The strong correlation between metal width and thickness variation due to CMP has also been proved from test chip measurements in [118]. The effect of line loss from dishing and erosion can be considerable and directly impacts the resulting electrical parameters of interconnecting wires. In [117], increase in resistance was observed on wide lines due to dishing and on high pattern densities due to dielectric erosion. Parasitic resistance, capacitance and inductance of the wire vary because of variations of metal and inter-layer dielectric (ILD) thickness and width as well as due to variations of material properties such as resistivity. It has been shown, that a 10% increase in width leads to about 10% increase in total capacitance, 12% increase in coupling capacitance, and 10% reduction in resistance [125]. It has been demonstrated, that parasitic RLC variations affect circuit performance [120]-[124].


In a current sensing interconnect, wire parasitics variation may impact



(a) I_{win} variation due to encoder and driver device parameters variability.




(b) I_{win} variation due to encoder, driver and termination device parameters variability.




(c) $I_{rec,in}$ variation due to encoder, driver and termination device parameters variability.


Figure 7.3: $PMCmFCD$ interconnect I_{win} and $I_{rec,in}$ variations



(a) I_{win} variation due to encoder and driver device parameters variability.



(b) I_{win} variation due to encoder, driver and termination device parameters variability.



(c) $I_{rec,in}$ variation due to encoder, driver and termination device parameters variability.

Figure 7.4: *DualdiffFCD* interconnect I_{win} and $I_{rec,in}$ variations




Figure 7.5: Interconnect model for analysis

signal integrity by causing variation in driver's output current and receiver's input current. Driver's output current I_{win} , is affected by the variation in its load (effective impedance which includes wire parasitics and termination load). The receiver's input current $I_{rec,in}$, is usually different from the driver's output current due to the non-ideal behavior of the interconnecting wire. These parameters are different from the ones estimated at design time due to variation of wire parasitics. The simple model of a current sensing interconnect demonstrating the variation in $I_{rec,in}$ is shown in Figure 7.5. Let us assume that the characteristic impedance of the lossy transmission line is Z_o and the current through the line is I_{win} . The near and far-end voltages are V_{win} and V_{wout} , respectively where V_{wout} and $I_{rec,in}$ can be expressed as follows:

$$V_{wout} = V_{win} - I_{win}Z_o \quad (7.2)$$

$$I_{rec,in} = \frac{V_{wout}}{R_t} \quad (7.3)$$

From these two equations, it can be seen that $I_{rec,in}$ can be different from its nominal value determined at design time, because of variations of wire parasitics (Z_o) and termination device parameters. It has already been demonstrated from the Monte Carlo simulations that variations of termination transistor parameters cause additional variations in $I_{rec,in}$. The receiver's reference current also deviates from its nominal value due to variation of its devices electrical parameters. The conventional approach to ensure the reliability of the interconnect is taking into account effects of all these variations at design time and allocating large enough current margin. However, this has considerable power consumption cost, this is especially significant for multilevel current sensing interconnects. Hence, one has to further explore alternative power efficient technique.

7.1.2 Runtime Supply Voltage and Temperature Variations

The impact of delivering increasing currents to the huge number of active devices on a chip, and the effect of parasitics on both on-chip and package power delivery wires, leads to deviation of V_{DD} and GND signals from their nominal values. Increasing operating frequencies and power densities in sub-100nm high performance ICs leads to an increase in voltage drops in the power grid. For instance, a voltage drop of 18% of the nominal voltage has been reported in *POWER6TM* dual-core processor fabricated in 65nm SOI processes [128]. In the multicore scenario, clock-gating, power-gating and other power saving techniques have undesired consequences like increase in the variations of current drawn by different cores leading to additional supply voltage fluctuations. The semiconductor industry has already moved to dynamic voltage drop (DVD) analysis [129] in order to account for the contribution of power density, variations in switching activity profile and impact of inductance and decaps. DVD also captures the impact of spatial and temporal switching events. This move shows the importance of taking into account the unavoidable temporal and spatial voltage drop fluctuations which may lead to signal integrity problems for on-chip communications if they are not addressed well. Moreover, temperature variability rises due to the distributed nature of an integrated circuit, and due to the fact that some components dissipate more power than others. In the silicon substrate, heat generated at one point spreads and causes an increase in temperature at nearby points. Temperature variations also occur with time, as the subsystems switch between idle and active periods. While designing the current sensing interconnects, temperature variations must be considered as they affect the device and wire characteristics.

Runtime supply voltage and temperature variations cause fluctuation in device's drain current. To characterize the drain current fluctuations induced by environmental variations, BSIM4 MOSFET current equations are used [130].

$$I_{ds} \propto \frac{I_{ds0}}{1 + R_{ds}I_{ds0}/V_{dseff}} \quad (7.4)$$

$$I_{ds0} \propto \frac{V_{gstseff}\mu_{eff}V_{dseff}(1 - A_{bulk}V_{dseff}/2(V_{gstseff} + 2V_T))}{1 + V_{dseff}/E_{SAT}L_{eff}} \quad (7.5)$$

Where I_{ds} , I_{ds0} , R_{ds} , V_{dseff} , V_{gsteff} , A_{bulk} , μ_{eff} , V_T , E_{SAT} , and L_{eff} are the drain current with short-channel effects, drain current of a long channel device, parasitic drain-to-source resistance, effective drain-to-source voltage, effective gate overdrive ($V_{GS} - V_t$), parameter to model the bulk charge effect, effective carrier mobility, thermal voltage, electric field at which the carrier drift velocity saturates and effective channel length, respectively. MOSFET channel current is a function of both gate and drain voltages. Either of these voltages, or both, are affected by supply voltage variations depending on the circuit configuration. These variations in turn affect the drain current. Absolute values of threshold voltage, carrier mobility, and saturation velocity degrade with the increase in temperature [130], [131]. The degradation of threshold voltage with temperature tends to increase the drain current due to the increase in gate overdrive ($V_{GS} - V_t$), whereas degradation in carrier mobility tends to reduce the drain current as can be seen from Equation 7.5. Hence, overall variation of I_{win} is determined from cumulative variation of V_{GS} and V_{DS} caused by supply voltage fluctuation and the variation of the dominant device parameter when the temperature varies. Furthermore, the resistivity of a wire increases with temperature [132], increasing the parasitic resistance of the interconnecting wires which in turn decreases $I_{rec,in}$. It is also affected by the supply voltage and temperature variations at the transmitter end (Figure 7.1).

As an example, the effect of voltage and temperature variations on *PM-CmFCD* interconnect's $I_{rec,in}$ and reference current have been examined in Cadence Analog Spectre using 65nm CMOS technology from ST Microelectronics. The voltage and temperature was swept by 25mV and 10°C, respectively. In $I_{rec,in}$ analysis, supply voltage and temperature at the transmitter end were varied, while in reference current analysis, voltage and temperature at the receiver end were varied. Both analyses show that voltage fluctuation has much more pronounced effect on current variation than temperature (Figure 7.6 and Figure 7.7). For instance, around the nominal operating point a $\pm 100\text{mV}$ change in supply voltage at the transmitter end causes about $\pm 40\mu\text{A}$ variation in $I_{rec,in}$ whereas a temperature increase of 100°C causes only about $\pm 13\mu\text{A}$ variation.




Figure 7.6: $I_{rec,in}$ versus supply voltage and temperature




Figure 7.7: I_{ref} versus supply voltage and temperature

7.2 Post-Manufacture Variation Adaptation

Process variations may cause signal integrity problems in a current sensing interconnect. This negatively impacts the manufacturing yield. Techniques are needed to alleviate the effects of such variations. The traditional assumption of worst-case variation and guard-banding technique which uses large current margins has high power consumption costs. The other approaches can be classified into two: circuit optimization techniques such as V_t modulation, and post-manufacture circuit tuning techniques. In [133] post-manufacture variation adaptation technique to keep the delay and leakage power of the circuit within an acceptable range has been proposed. This technique relies on a hardware framework that supports self-test and performs self-adaptation using optimization algorithms of design parameters.

The process variation tolerance technique proposed in this chapter also uses a post-manufacture self-adaptation mechanism. The receiver's input and reference current variations are the result of manufacturing fluctuations which are static. Wear-out and ageing also cause variation but they are time-dependent on the scale of months and years. Hence, in the proposed technique, inter-

connect's signal integrity test and calibration are performed at every power start-up of the system to tackle process, wear-out and aging related variations. If an error is detected then the receiver, driver or both are reconfigured according to the developed algorithm and methodology. This makes the link adaptive to the effect of variations and thus enabling continuous and reliable operation of the interconnect. It also results in lower power consumption when compared to the worst-case approach.

7.3 Calibration for Process Variation Tolerance

The post-manufacture calibration technique has two advantages. The first and the most important one is ensuring tolerance to process variation and reliable communication by making the link adaptive to the effects of variation. The second one is reducing power consumption. Rather than assuming worst-cases and allocating large current margin which causes unnecessary power consumption, the margin is adjusted at every power start-up during the calibration phase by detecting the existing amount of variation. Based on the detection, receiver and driver reconfiguration will be performed. This is an efficient technique since it saves power by optimizing the margin and at the same time guarantees reliability. An error detection scheme as well as reconfiguration algorithms and methodology are developed. Furthermore, reconfiguration control and communication circuits are designed and simulated for a multilevel current sensing interconnect.

7.3.1 Algorithm and Methodology

The interconnect's signal integrity test is initiated by the receiver. When error is detected, receiver reconfiguration will be carried out first; if it is not enough to handle the variation then driver reconfiguration will be followed. Upon successful completion of calibration processes, the interconnect will be ready for data transmission phase. However, if both reconfigurations are failed the link will be declared as 'faulty/do not use'. The flows of the interconnect calibration process are shown in Figure 7.8.

The calibration process is formulated by considering a three-level current sensing interconnect (0 , I_1 , and I_2) the same as in *PMCMFCD* and *DualdiffFCD* interconnects. The reason for choosing three-level is that *PMCMFCD* and *DualdiffFCD* interconnects have superior performance and better power

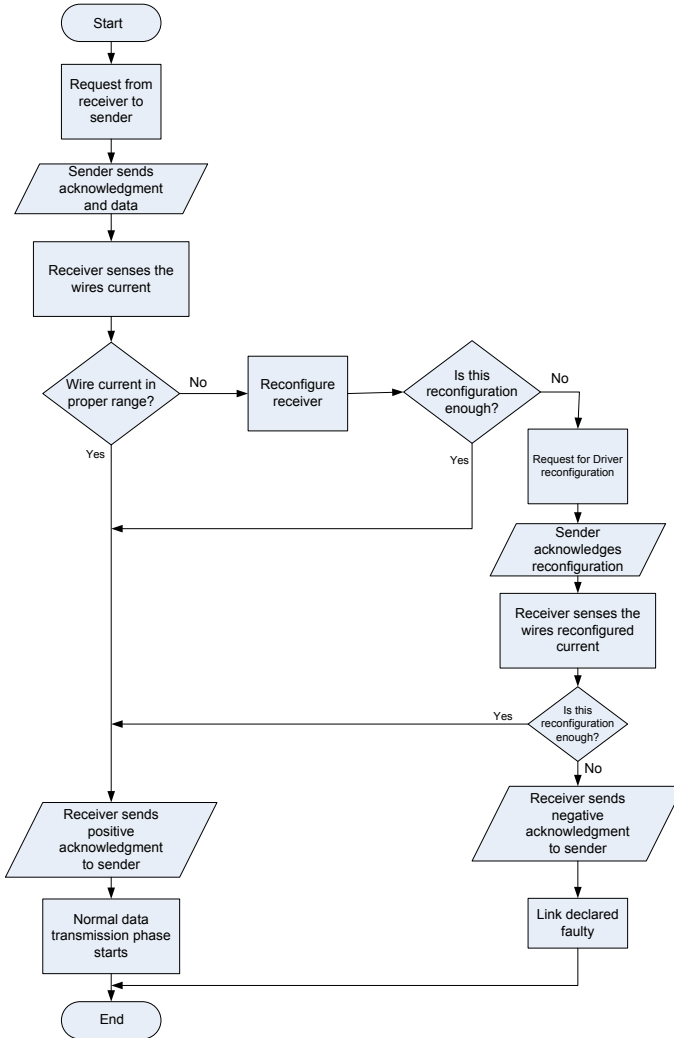


Figure 7.8: Interconnect calibration flow chart

efficiency than the *LEDRCm* which uses binary current sensing signaling (see Chapter 6). In fact, the formulated calibration process is scalable to any current sensing interconnects including binary current sensing signaling (0 and I). In a three-level current sensing interconnect Equations 7.6 to 7.8 should be satisfied in order to ensure its signal integrity. Besides data wires, four additional wires are needed to carry out the calibration (Figure 7.9). *Calib_Ack*, *Calib_Req*, *C1* and *C2* wires are needed for handshaking and communicating the results of reconfigurations between sending and receiving modules during the calibration phase.




Figure 7.9: Interconnect with calibration wires

$$I_{ref1} < I_1 < I_{ref2} \quad (7.6)$$

$$I_2 > I_{ref2} \quad (7.7)$$

$$I_{ref2} > I_{ref1} \quad (7.8)$$

Reconfigurable driver and receiver current sources of a three-level current sensing interconnect are shown in Figure 7.10. When the receiver is ready to accept data it closes the switch of I_{min2} and I_{min1} . These switches stay closed all the time because to achieve the minimum required performance I_1 and I_2 wire currents must be greater than I_{min1} and I_{min2} currents, respectively. From Equation 7.6 and 7.8, it can be deduced that I_2 is always greater than I_1 which in turn means that I_{min2} is greater than I_{min1} . Initially the switches of I_{var1} and I_{var2} will also be closed. These two switches are needed to tolerate I_1 and I_2 wire current variation respectively. Depending on the variations, these two switches might be opened as a result of receiver reconfiguration. Additional current source I_{var21} is required in order to avoid an error wherein a considerable increase (variation) in I_1 could mislead the receiver to interpret it as I_2 . In this case I_{ref2} will also be increased. If I_1 increases it is more likely that I_2 also increases, because the drivers of I_1 and I_2 are placed closer to each other making them highly spatially correlated. If I_2 is not increased like in I_1 then the driver reconfiguration can be used to increase it, if necessary. Initially the switch of I_{var21} is open and will be closed when needed.

The receiver's output signals are $comp1$ and $comp2$. In a reliable interconnect, when I_1 is being transmitted, $comp1$ should be *high* and $comp2$ should be *low*. When I_2 is transmitted both $comp1$ and $comp2$ should be *high*. These two cases will be checked for all wires during the calibration process. In the

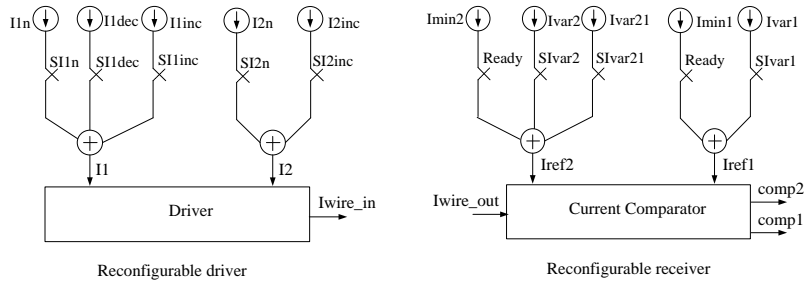


Figure 7.10: Reconfigurable driver and receiver

driver, the switches $SI1n$ and $SI2n$ are always closed as they are the minimum currents that are required to achieve the desired communication performance. Initially, switch $SI1dec$ is also closed and will be opened if the receiver requests for decreasing I_1 . Switches $SI1inc$ and $SI2inc$ are open and will be closed when necessary.

The calibration process will be initiated by the receiver when it sends a request through $Calib_Req$ wire to the sender. The sender sends an acknowledgment through $Calib_Ack$ wire and current I_1 through all data wires immediately after it gets the first request transition (low-to-high). Outputs $comp1$ and $comp2$ of all wires will be checked after the acknowledgment signal arrives at the receiver. The algorithm for I_1 calibration is presented as pseudocode in Algorithm 1. There are three possible scenarios depending on the amount of variation. In *Case 1*, I_1 is in proper range. In *Case 2*, I_1 becomes less than expected and receiver cannot detect it. Finally in *Case 3*, I_1 becomes much larger than expected and receiver might detect it as I_2 .

Upon getting a high-to-low request signal transition, the sender drives I_2 through data wires and acknowledgment through $Calib_Ack$ wire. The algorithm for I_2 calibration is presented as pseudocode in Algorithm 2. There are also three possible variation cases for I_2 . In *Case 1*, I_2 is in proper range. In *Case 2*, I_2 is less than expected and this causes error because the receiver can detect it as I_1 . In the last case, I_2 may even be less than I_1 and the receiver cannot detect any transmission in this case.

The success of calibration and the links reliability is confirmed when there is a second low-to-high request transition through $Calib_Req$ wire. A calibration failure is indicated by sending a second pulse either on $C1$ or $C2$ wire. The calibration process can be classified into best, average, worst and failure

Algorithm 1 Calibration of I_1

Receiver: Req = 0 \rightarrow 1; // request for start of calibration;
Sender: Ack = 0 \rightarrow 1 and transmit I_1 through data wires;
Receiver: gets acknowledgment;
Case 1:
Receiver: if ($comp1 = 1$ and $comp2 = 0$) then
 reconfiguration is not required;
 Req = 1 \rightarrow 0; // I_1 is reliable and request for I_2
Case 2:
Receiver: if ($comp1 = 0$ and $comp2 = 0$) then
 opens *Ivar1* switch;
 check $comp1$ and $comp2$;
 if ($comp1 = 1$ and $comp2 = 0$) then
 receiver reconfiguration is successful;
 Req = 1 \rightarrow 0; // I_1 is reliable and request for I_2
 else // Receiver reconfiguration is not sufficient
 sends a pulse on *C1* wire; // request for I_1 increase
Sender: Ack = 1 \rightarrow 0; // sends acknowledgment
Receiver: gets acknowledgment;
 check $comp1$ and $comp2$;
 if ($comp1 = 1$ and $comp2 = 0$) then
 driver reconfiguration successful;
 Req = 1 \rightarrow 0; // I_1 is reliable and request for I_2
 else
 sends a second pulse on *C1*; // I_1 calibration not successful
Case 3:
Receiver: if ($comp1 = 1$ and $comp2 = 1$) then
 close switch of *Ivar21*;
 check $comp1$ and $comp2$;
 if ($comp1 = 1$ and $comp2 = 0$) then
 receiver reconfiguration successful;
 Req = 1 \rightarrow 0; // I_1 is reliable and request for I_2
 else
 sends pulse on *C1* and *C2* // request for I_1 decrease
Sender: Ack = 1 \rightarrow 0; // sends acknowledgment
Receiver: gets acknowledgment;
 check $comp1$ and $comp2$;
 if ($comp1 = 1$ and $comp2 = 0$) then
 driver reconfiguration is successful;
 Req = 1 \rightarrow 0; // I_1 is reliable and request for I_2
 else
 sends a second pulse on *C1* // I_1 calibration not successful

Algorithm 2 Calibration of I_2

Sender: Ack = 0 \rightarrow 1 or 1 \rightarrow 0; // depending on previous state transmits I_2 ;

Receiver: gets acknowledgment;

Case 1:

Receiver: if ($comp1 = 1$ and $comp2 = 1$) then
reconfiguration is not needed;
Req = 0 \rightarrow 1; // I_2 is in proper range and calibration completed successfully and link is reliable

Case 2:

Receiver: if ($comp1 = 1$ and $comp2 = 0$) then
opens $Ivar2$ switch; // receiver reconfiguration
check $comp1$ and $comp2$;
if ($comp1 = 1$ and $comp2 = 1$) then
Req = 0 \rightarrow 1; // calibration completed successfully
else
sends a pulse on $C2$ wire; // request for I_2 increase

Sender: Ack = transition; // sends acknowledgment

Receiver: gets acknowledgment;
check $comp1$ and $comp2$;
if ($comp1 = 1$ and $comp2 = 1$) then
Req = 0 \rightarrow 1; // calibration completed successfully
else
sends a second pulse on $C2$; // calibration not successful

Case 3:

Receiver: if ($comp1 = 0$ and $comp2 = 0$) then
sends a pulse on $C2$; // request for I_2 increase

Sender: Ack = transition; // sends acknowledgment

Receiver: gets acknowledgment;
check $comp1$ and $comp2$;
if ($comp1 = 1$ and $comp2 = 1$) then
Req = 0 \rightarrow 1; // calibration completed successfully
elseif ($comp1 = 1$ and $comp2 = 0$) then
opens $Ivar2$ switch;
check $comp1$ and $comp2$;
if ($comp1 = 1$ and $comp2 = 1$) then
Req = 0 \rightarrow 1; // calibration completed successfully
else
sends a second pulse on $C2$; // calibration not successful

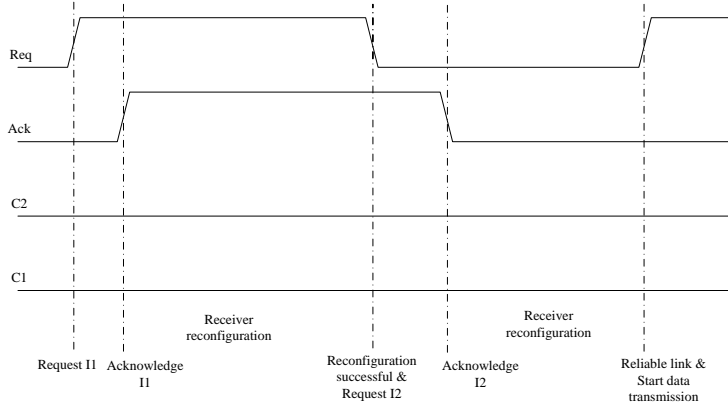


Figure 7.11: Best case : driver reconfiguration is not required

cases depending on the number of steps required and the final result (successful or failure). In the best case either there is no need for reconfiguration at all, or only receiver reconfiguration is enough (Figure 7.11). In the average case only one driver reconfiguration either on I_1 or I_2 besides receiver reconfiguration is required. There are three possible ways in which the average case can be manipulated: by increasing I_1 , decreasing I_1 and increasing I_2 (Figure 7.12).

In the worst case, two driver reconfigurations are needed in addition to the receiver reconfiguration (Figure 7.13). There are two possible ways in which the worst case can be handled: by increasing both I_1 and I_2 and decreasing I_1 and increasing I_2 . If both receiver and driver reconfigurations failed to make the link adaptive to the variation, then the link is in failure state (Figure 7.14). At this stage the error will be reported to a higher level error controlling system. This leads to a more power efficient error detection and correction scheme, because higher level error controlling mechanisms will come into play only when necessary. There are three failure scenarios: does not compensate for the variation by decreasing/increasing I_1 or increasing I_2 besides receiver reconfigurations. Upon completion of the calibration process successfully, data transmission phase will start.

7.3.2 Reconfiguration Control and Communication Circuits

In general the circuits can be classified into two parts: driver and receiver side circuits. The receiver side circuit detects the receiver's outputs and performs

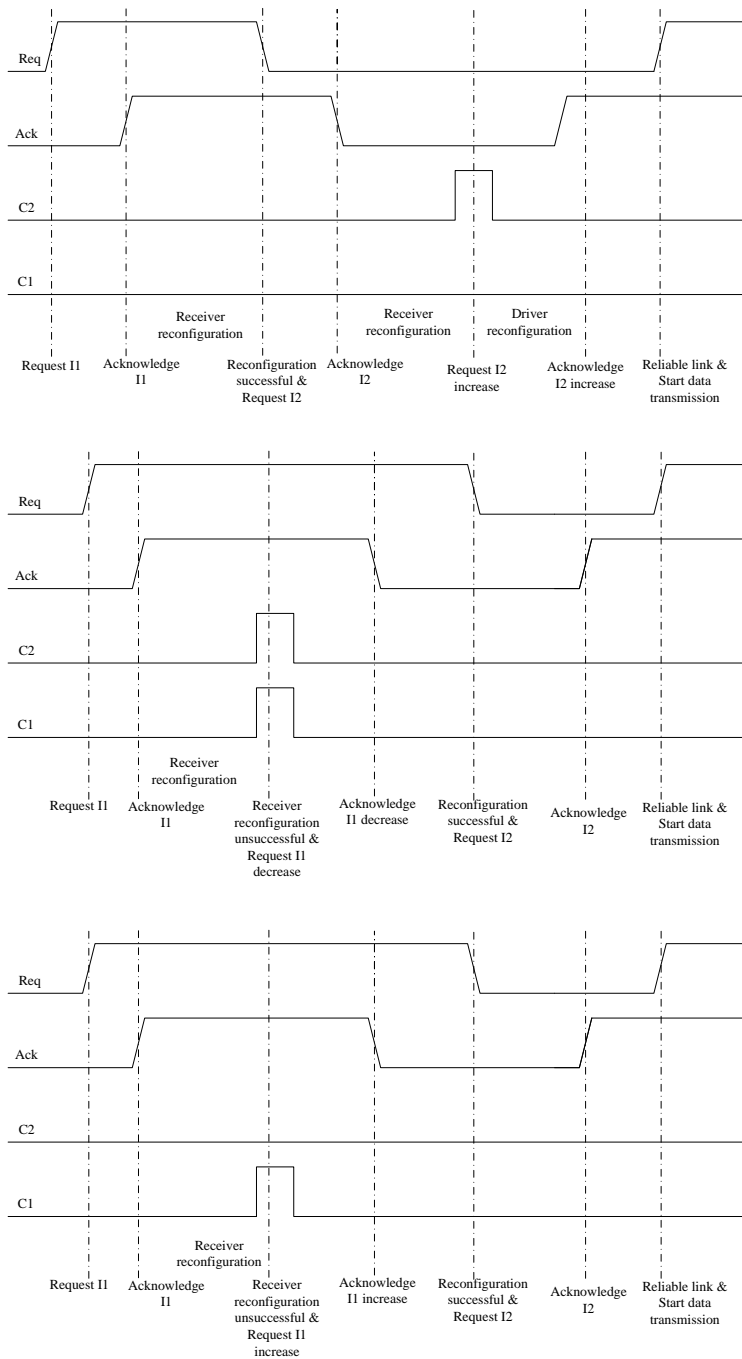


Figure 7.12: Average case calibrations

the needed reconfiguration at the receiver by increasing or decreasing reference currents. It also sends requests for I_1 and I_2 transmissions, and for driver

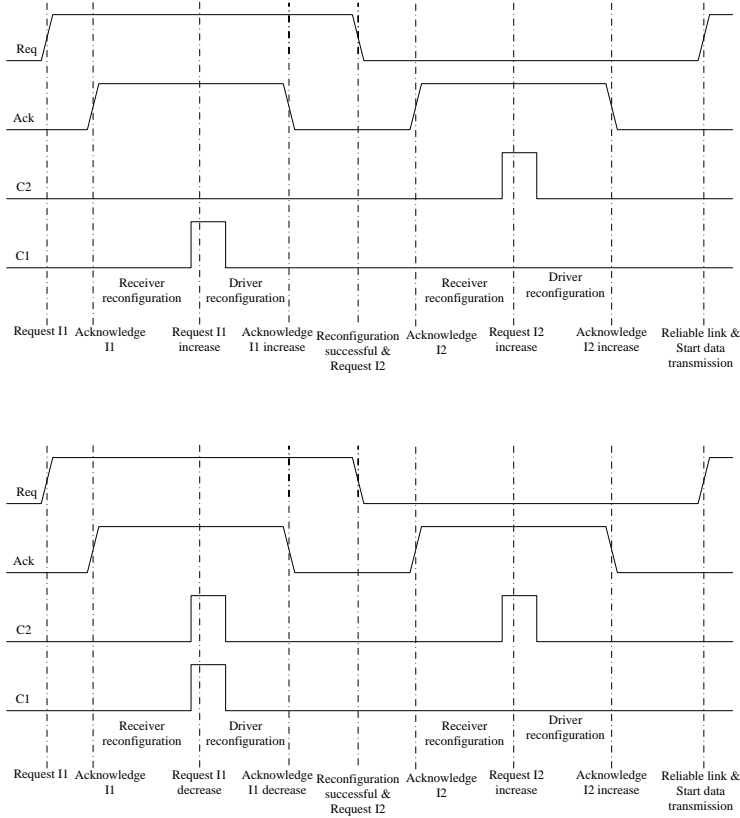


Figure 7.13: Worst case calibrations

reconfiguration when required. In addition, it communicates the calibration results with the transmitter. The driver sends either the nominal or reconfigured current through data wires and an acknowledgment signal through the *Calib_Ack* wire depending on the state of handshaking signals. The input and output signals of driver and receiver reconfiguration control circuits are shown in Figure 7.15. The block level diagram is intended to provide a clear distinction between reconfiguration control and communication signals.

The receiver and driver reconfiguration control circuits are shown in Figures 7.16 and 7.17, respectively. Immediately after power start-up, the receiver raises its *Ready* signal to *high* whenever it is ready to accept data, this makes *Reqin* signal to have a transition from low-to-high. Upon getting *Reqout* transition the sender transmits I_1 and *Ackin* (low-to-high transition) through data and *Calib_Ack* wires, respectively. The receiver checks its output signals, *comp1* and *comp2*, when it gets a transition in *Ackout*. Receiver reconfiguration will be performed by controlling the reference current's switches using

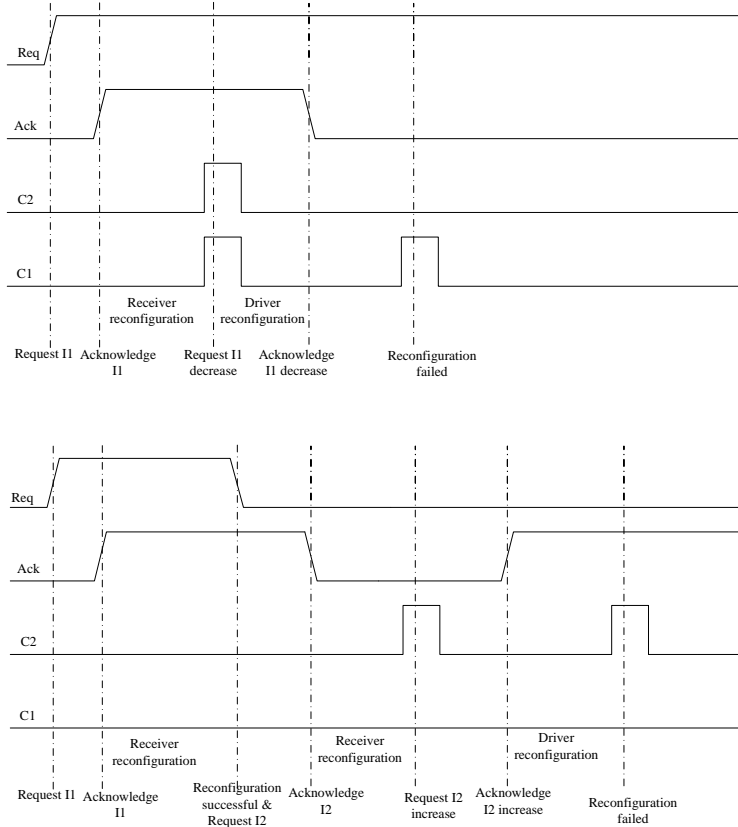


Figure 7.14: Calibration failure cases

$SIvar1$, $SIvar2$ and $SIvar21$ signals (Figure 7.16). If the receiver reconfiguration is able to fix the signal integrity problem, then the receiver communicates

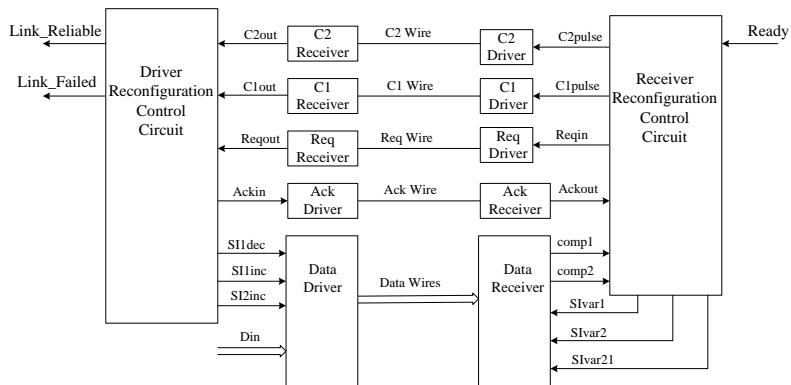


Figure 7.15: Calibration control and communication signals

it to the transmitter using *Reqin* signal transition. If not, it sends a request for driver reconfiguration using signals *C1pulse*, *C2pulse* or both depending on the required reconfiguration through their respective wires. Based on the receiver reconfiguration result, transmitter sends either I_2 or reconfigured I_1 . The driver performs reconfiguration of I_1 using *SI1dec* or *SI1inc* signals (Figure 7.17). When the receiver requests for I_2 reconfiguration, it will do so using *SI2inc*. In case both sides of reconfigurations fail, a second pulse will be generated on *C1pulse* or *C2pulse* depending on the current level under the reconfiguration process. Upon getting a second pulse the transmitter raises *Calibration_Failed* signal to *high*, indicating the failure of the calibration to adapt with the variation. The success of calibration, in other words process variation tolerance of the interconnect is approved by the receiver when it generates a second low-to-high transition in *Reqin* signal (Figure 7.16). The *Link_Reliable* signal will be *high* when the transmitter gets a second low-to-high transition on *Reqout* signal (Figure 7.17).

The d(latch) and d(comp) in Figure 7.16 are delay elements and their delay correspond to the delay of latch and current comparator, respectively. The *t2p* block in Figure 7.16 and *p2t* block in Figure 7.17 are transition-to-pulse and pulse-to-transition converters, respectively and their implementation is shown in Figure 7.18.

7.4 Runtime Management of Voltage and Temperature Variations

In a global on-chip communication link, transmitter and receiver are placed far apart from each other. At runtime, the supply voltage and temperature of a transmitter can be different from the receiver depending on the spatial switching activities and hotspot localities. These in turn deviate the transmitter's output current, and consequently the receiver's input current from its nominal value. In a current sensing interconnect, where a receiver compares its input current with a reference current, this variation may affect reliability if it is out of the allocated margin. The usual trend is assuming worst-case variation and allocating large enough current margins between the receiver's input and reference currents which leads to additional power consumption.

An alternative power efficient technique, is monitoring the variation at runtime and adjusting the interconnect circuits when a signal integrity problem is

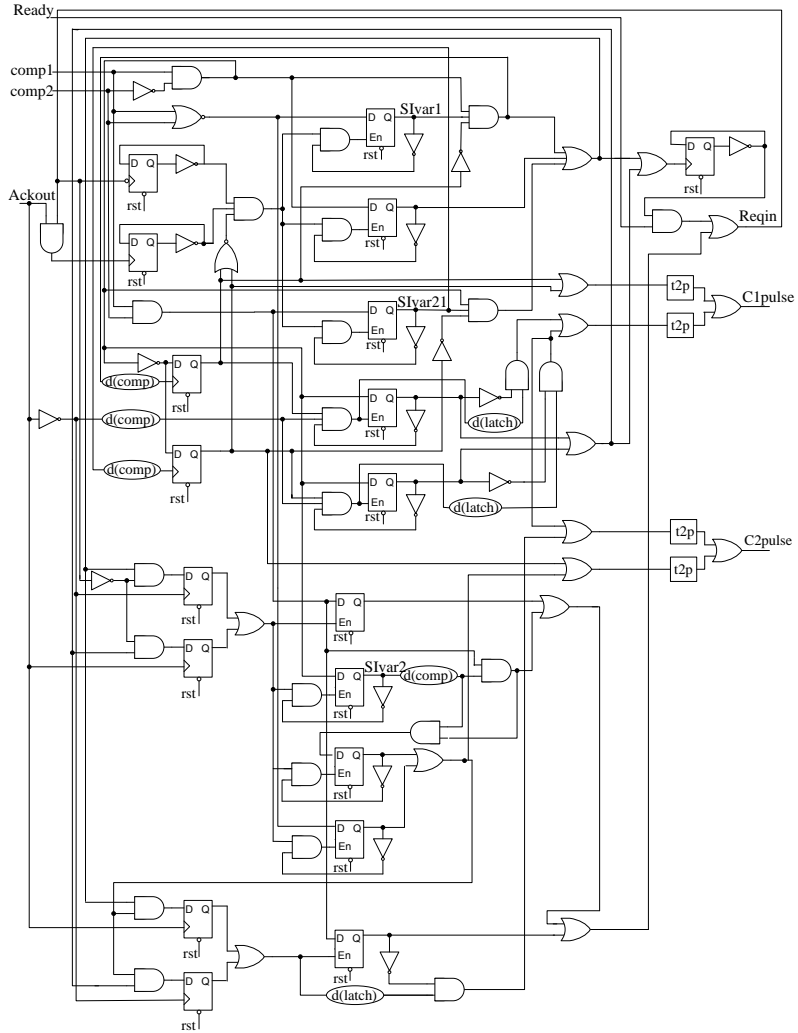


Figure 7.16: Receiver reconfiguration control circuit

detected. This enables power efficient, runtime error detecting and correcting scheme. To do so, circuit level variation sensing mechanism along with its implementation is devised and presented in this section. When an error is detected due to runtime variation, reconfiguration of the interconnect circuits and retransmission of the data will be carried out.

7.4.1 Sensing Effects of Voltage and Temperature Variation

Sensing is an important task of any adaptive system that compensates for variation. A sensor monitors the runtime operating conditions of a system.

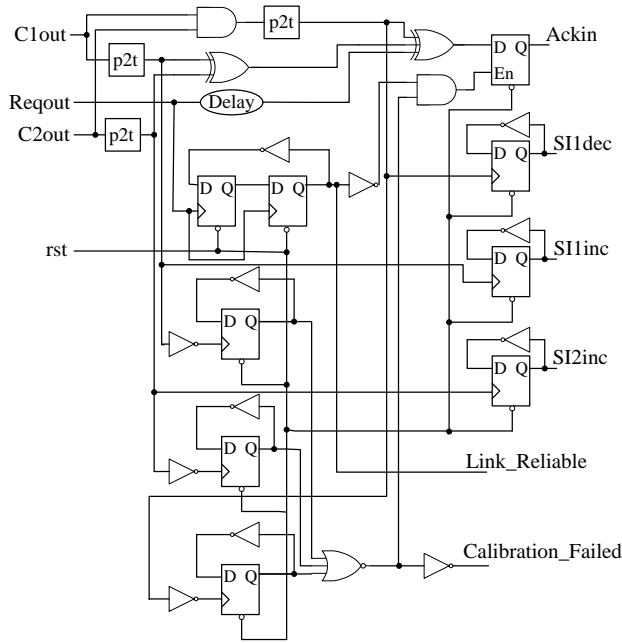


Figure 7.17: Driver side calibration control




Figure 7.18: Transition-to-pulse and pulse-to-transition converters

Here runtime variation of $I_{rec,in}$ along with receiver’s reference current will be monitored. If the variation causes error, the error will be reported to both the transmitter and receiver, besides reconfiguring the receiver to adapt with the variation. The receiver output will be erroneous if $I_{rec,in}$ becomes lower than expected or receiver’s reference current increases more than needed. There are three causes for the error: large supply voltage drop at the transmitter side, significant increase in temperature at the receiver or both. So, effects of supply voltage fluctuation at the transmitter and temperature variation at the receiver are considered in the design of the sensing circuit.

Runtime monitoring of Voltage and Temperature (VT) variations is carried out using two additional wires which run adjacent to data transmission wires (Figure 7.19). With every new data transmission, the same amount of current

as in data wires will be transmitted through the sensing wires by changing the direction of current. The sensor circuit at the receiver compares the sensing wire's current with the receiver's reference current. If the sensing wire's current is greater than the reference, the sensor output stays low, indicating there is no variation which causes error. If the sensing wire's current is less than the reference, then sensor output goes high, thus detecting error.

7.4.2 Sensor Circuit Implementation

The sensor circuit is based on current subtraction. It subtracts the receiver's reference current from the sensing wire's current. The sensor circuit is shown in Figure 7.20. The current direction changes in the sensing wires with every new data transmission in the channel. For example, if I_{vts1} flows towards the receiver, then I_{vts2} flows towards the transmitter and vice versa in the next transmission. Due to this, there is current either in $Mn1$ or $Mn2$ at any time. Current I_{vts1} and I_{vts2} are mirrored to $Mn5$ and $Mn6$, respectively. The reference current is mirrored to $Mp4$. The current comparator, which is based on current subtraction, compares I_{ref} with either I_{vts1} or I_{vts2} depending on the one which flows towards the receiver. The comparator output is buffered to make the $Cout$ signal full swing.

The purpose of the current direction sensor is to know the arrival of new data and check for its reliability when it is valid and stable. The current direction sensor circuit is shown in Figure 7.20 and was used as part of a receiver in *Dualdiff* interconnect. Consider the top current direction sensor in Figure 7.20. Transistor $Mp1$ provides negative feedback to transistor $Mn3$. It switches the gate of $Mn3$ on and off as required and helps in modulating the input impedance. Transistor $Mp2$ provides a constant current bias and

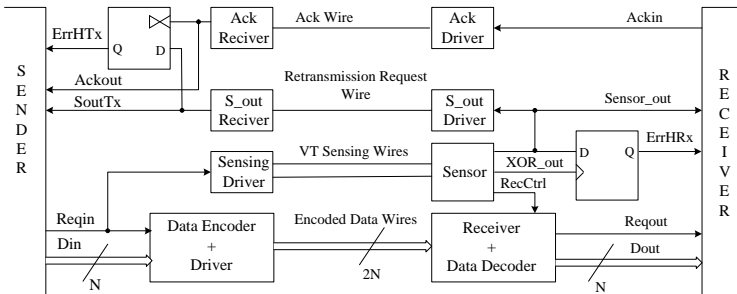


Figure 7.19: Interconnect with VT runtime variations detector

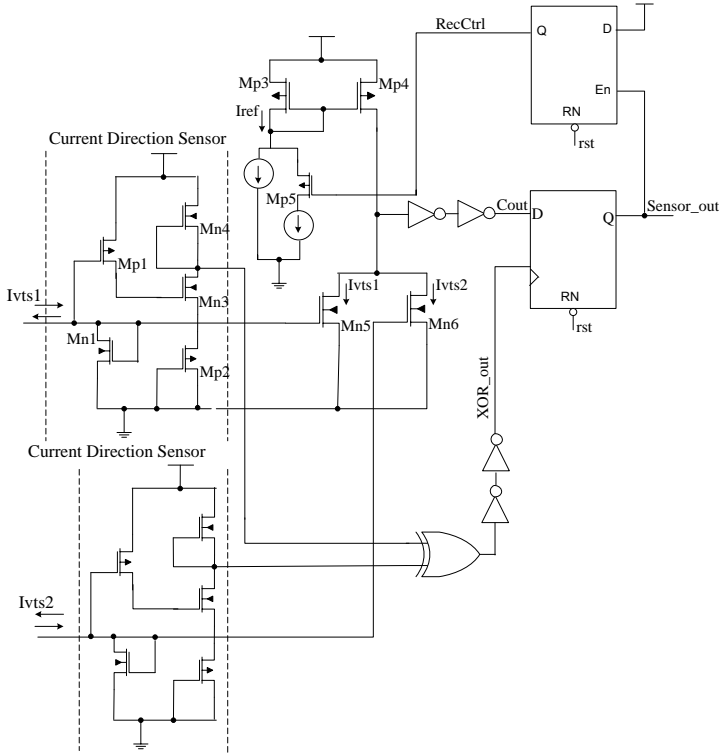


Figure 7.20: VT variation sensor and reconfiguration circuit

thus regulates the transconductance of $Mn3$. The source terminal of transistor $Mn3$ is connected to the I_{vts1} wire. When current flows towards the driver, $Mn3$ switches to on state and pulls the output of the current sensor to *low*. When current is sourced by the driver, the source voltage of $Mn3$ rises thus switching it off. In this case, current flows through the load transistor $Mn4$ to the output, making the output voltage of the current direction sensor *high*. The output of the two direction sensors are used as inputs to the XOR gate. The additional delay due to XOR gate and buffers, ensures the stability of $Cout$ before XOR_out becomes *high*. If either I_{vts1} or I_{vts2} is less than I_{ref} , then $Cout$ becomes *high*. This output will be latched to $Sensor_out$ when XOR_out makes a transition to *high*, thus detecting the error due to runtime voltage and temperature variations.

7.4.3 Reconfiguration and Retransmission

The receiving block always checks the output of the sensor circuit $Sensor_out$ before it uses or forwards the received data. When $Sensor_out$ becomes *high*

it suspends the use of data until *Sensor_out* returns to *low*. Reconfiguration of the receiver, more specifically decreasing the reference current *Iref*, will be carried out when the sensor detects error. The latch will be enabled when *Sensor_out* is *high* which in turn makes *RecCtrl* *high* and then switches *Mp5* to a non-conducting state. In order to check the success of reconfiguration in withstanding the variation effect, retransmission request will be sent to the transmitter through the retransmission request wire (Figure 7.19). *Sensor_out* signal is used as a retransmission request signal. When the transmitter gets a retransmission request signal, it sends the data again and changes the current directions in the sensing wires. *XOR_out* signal makes a transition to *high* due to current direction changes in sensing wires, allowing *Cout* to be latched to *Sensor_out*. If the reconfiguration is successful, *Sensor_out* becomes *low* and the receiver and transmitter resume their normal data transmission phase. Otherwise, the error will be reported to the transmitter, receiver and the higher level error controlling system. When the reconfiguration fails, *ErrHRx* signal becomes *high*, and then the receiver sends negative acknowledgment through *Ackwire*, which in turn makes *ErrHTx* *high*, thereby informing the failure of reconfiguration to the transmitter. The transmitter translates *Ackout* signal transition as negative acknowledgment when *SoutTx* is *high* (Figure 7.19). If there is a transition in *Ackout* signal and *SoutTx* is *low* then the transmitter sends the next data. Whereas it retransmits the same data if there is no transition in *Ackout* and *SoutTx* is *high*. When the reconfiguration has failed to tolerate the effect of variations, the link will be labeled as temporarily failed until the higher level error controlling system fixes the variation, for instance by decreasing the switching activities to make the voltage drop in an acceptable margin.

7.5 Simulation Results and Analysis

Simulations of *PMcmFCD* and *DualdiffFCD* interconnects consisting of the calibration and VT runtime variation management circuits were designed and performed in Cadence Analog Spectre using 65nm CMOS technology from STMicroelectronics and 1V supply voltage. The interconnect length was set to 2mm, the same as in inter router link length of Intel 80-Tile TeraFLOPS processor [64]. The wire properties were set according to ITRS 65nm technology node for global wiring. The RLC values of the wires were extracted

using field solvers for microstrip configuration. The resistance and inductance values were extracted using FastHenry [45], while the capacitance values were extracted using Linpar [46]. In calibration wires (*Calib_Req*, *Calib_Ack*, *C1* and *C2*) voltage-mode signaling with repeater insertion was used.

The time taken and average power consumed during the calibration process is listed in Table 7.1. The calibration delay for best-case, where driver reconfiguration was not required, was 2.66ns. This is the minimum delay incurred due to calibration at every power start-up of the system. The best-case requires five communications between the sender and the receiver. They are separate request and acknowledgment for both I_1 and I_2 transmission besides communicating the robustness of the link. The average-case delay, which requires one driver reconfiguration in addition to receiver reconfiguration, was 4.19ns. The delay of the worst-case calibration was 5.72ns, which requires both driver and receiver side reconfiguration for I_1 and I_2 . The average power consumed during the calibration process is low. However, its peak power is high, (Table 7.1) but it occurs only for a very short period.

Table 7.1: Calibration delay and power consumption.

Classification	Delay [ns]	Average Power [μW]	Peak Power [μW]
Best-case	2.66	164	1188
Average-case	4.19	179	1392
Worst-case	5.72	243	1345

In order to examine the power saving benefits of the presented calibration technique, three 64-bits wide *PMCMFCD* interconnects consisting of the calibration circuits were designed. The first interconnect was designed by guard-banding for worst-case variation. Let us quantify it as $\pm 3\sigma$ variation which accounts for about 99.7% of the overall variation range. The second interconnect is designed by allocating current margins for $\pm 2\sigma$ variations, which covers about 95.4% of the variation. The third one is designed with current margins for $\pm 1.5\sigma$ variations. The second and third interconnects have saved 7.88% and 14.21% power, respectively over the interconnect with worst-case margin allocation. This proves that allocating smaller margins and relying on the proposed calibration technique will lead to a better power efficiency than the conventional worst-case design.

The required additional area due to the calibration circuits was determined

from the Cadence schematic. It requires $25\mu\text{m}^2$ active and $2940\mu\text{m}^2$ wiring areas. The calibration area overhead for *PMCmFCD* interconnect has been calculated and it decreases for larger bit width transmissions (Table 7.2). For example, the active area overhead for 4-bits and 64-bits *PMCmFCD* interconnects are 38.96% and 4.67%, respectively.

Table 7.2: Calibration area overhead.

Bit Width [bits]	Active Area Overhead [%]	Wiring Area Overhead [%]
2	50.91	38.88
4	38.96	26.92
8	26.23	16.66
16	15.76	9.72
32	8.80	5.07
64	4.67	2.63

Simulation waveforms for an average-case calibration process is shown in Figure 7.21. As it can be seen from the simulation waveforms, *Reqin* signal goes *high* when it gets *Ready* signal from the receiver and *Ackin* goes *high* after it gets a *high* transition on the *Reqout* signal. The receiver’s reconfiguration control checks *comp1* and *comp2* and detects error, because both *comp1* and *comp2* were *low*. Then it reconfigures the receiver by turning on the switch of *Ivar1*, making *SIvar1* signal *high* by decreasing *Iref1*. But this is not enough to control the effect of variation. Then the receiver requests for an increase in I_1 by sending *C1pulse*. Upon getting a pulse on *C1out*, the transmitter reconfigures the driver, increases I_1 by closing switch *SI1inc*. It also transmits the reconfigured I_1 along with acknowledgment. The receiver checks *comp1* and *comp2* signals when *Ackout* makes transitions to *low*. Then both the signals are in proper range, which confirms the reliability of I_1 transmission. The receiver requests for I_2 transmission by sending a high-to-low transition in *Calib_Req* wire. When *Reqout* makes a transition from high-to-low, the transmitter sends signals I_2 and acknowledgment (transition to *high*). The receiver checks signals *comp1* and *comp2* when it gets a transition in *Ackout*, both are *high*, indicating I_2 is in proper range. The receiver then declares the interconnects reliability by sending a low-to-high transition in *Calib_Req* wire. The transmitter declares the link’s reliability upon getting a second low-to-high transition on *Reqout* signal by raising the *Link_reliable* signal to *high*.




Figure 7.21: Simulation waveforms of average-case calibration

After the calibration phase is completed, normal data transmission along with runtime VT variation sensing is performed. The delay of the VT variation sensor circuit was measured and it is 184ps. The power consumed due to monitoring the variation was $489\mu\text{W}$. This is the only power overhead if there is no retransmission due to error. As can be seen from Table 7.3, this power overhead decreases for larger bit width transmissions. For example, in *DualdiffFCD* interconnect 4 and 64-bit transmissions the power overheads are 17.2% and 1.5%, respectively. The power overhead is reasonable and affordable for 16-bits and wider transmissions.

Table 7.3: Power overhead of VT variation management.

Bit Width [bits]	PMCMFCD [%]	DualdiffFCD [%]
2	26.8	26.3
4	21.6	17.2
8	10.9	10.2
16	6.2	5.6
32	3.2	2.9
64	1.6	1.5

It is possible to have power saving rather than power overhead using the proposed VT management circuits. For example, instead of the usual worst-case guardbanding for 10% V_{DD} variation (which is equivalent to 3σ according to the ITRS roadmap [58]), allocating margin for 6.67% V_{DD} variation covering 95% of the variations range can be sufficient. The rest will be relied on the proposed runtime monitoring and reconfiguration. This approach results in 2% power savings in a 64-bit *PMCMFCD* interconnect instead of consuming additional power.

The required active and wiring areas for this technique (including retransmission and reporting error to higher level error controlling system) are $5.83\mu\text{m}^2$ and $2100\mu\text{m}^2$, respectively. The portion of area taken by the VT management circuits have been determined for *PMCMFCD* interconnect and listed in Table 7.4. In 64-bit *PMCMFCD* interconnect, it requires only 1.13% of the active area and 3.81% of the wiring area. The area overhead becomes smaller for larger bit width transmissions, thus showing its appropriateness for real life applications.

Table 7.4: VT variation management area overhead.

Bit Width [bits]	Active Area Overhead [%]	Wiring Area Overhead [%]
2	19.47	31.24
4	12.95	20.83
8	7.65	12.49
16	4.18	7.14
32	2.20	3.67
64	1.13	1.89

Simulation waveforms are shown in Figure 7.22 for a nominal supply voltage of 1V, and assuming that there is a supply voltage drop at the transmitter




Figure 7.22: Simulation waveforms of VT variation tolerance

which causes receiver’s input current variation that leads to an error. The sensor then detects an error and it is flagged by making *Sensor_out* high. This in turn leads to reconfiguration of the receiver by making *RecCtrl* signal high and at the same time a retransmission request is sent to the transmitter, by making *SoutTx* signal high. Sensing after reconfiguration and arrival of the retransmitted data asserts the reliability of the link by making *Sensor_out* signal low.

7.6 Chapter Summary

In this Chapter, circuit techniques for PVT variations tolerance for current sensing on-chip interconnects are presented. The technique for process variation tolerance is based on detecting signal integrity of the interconnect and performing calibration at every power start-up of the system. If an error is detected, the receiver's reference and/or input currents will be adjusted through receiver and driver reconfigurations. This makes the interconnect adaptive to the effects of process, wear-out and aging related variations, thereby enabling its continuous and reliable operation.

In a current sensing interconnect, using traditional worst-case guardbanding to tolerate environmental variations is costly and may not even be sufficient as the amount of runtime variation in high performance ICs are increasing. This makes runtime VT variation management technique an alternative and a better approach. The presented technique for runtime VT variations tolerance is based on monitoring the effect of their runtime variations and then reconfiguring the receiver when an error is detected. After the reconfiguration, request for data retransmission will be sent. The power and area overhead of this technique is low especially for larger bit width transmissions. It has been even proved that power can be saved compared to the worst-case approach. Therefore, the presented techniques are among the promising circuit level variation tolerance techniques for wider bits and high-performance on-chip interconnects.

Chapter 8

Conclusions

Global on-chip interconnects get slower with technology scaling and dissipate more power. At the same time, parameter variations are increasing, causing signal propagation delay to be uncertain, which in turn affects the performance and reliability of communication significantly. These global on-chip communication bottlenecks have been the primary focus of this thesis, and to tackle these challenges several design techniques have been developed and integrated. Based on these techniques, high performance, power efficient and variation tolerant on-chip interconnects have been presented.

Through design of self-timed delay-insensitive data transmission schemes, the reliability challenge due to delay variations has been addressed. Few delay-insensitive codes have been selected to be used in the interconnects after examining many codes in terms of their delay and power overheads. To further reduce the overhead, high speed and power efficient signaling techniques have been introduced. Area and power efficient integration of delay-insensitive and differential transmission has also been implemented in one of the presented interconnects to boost its noise immunity in addition to delay variation insensitivity. The feasibility and superiority of the designed interconnects have been verified through simulations and comparison with conventionally implemented delay-insensitive interconnects.

In bit-parallel delay-insensitive data transmission, the larger the channel bit width is, the longer is the overall time spent in completion detection, because each detection circuit becomes a tree of logic elements. In order to overcome this challenge, a novel high speed completion detection technique along with its circuit realization has been designed. Its speed is independent

of the channel bit width. This detection technique has been implemented in two of the presented interconnects. Simulation results showed that significant improvement in performance and energy dissipation has been achieved, especially for large bit width transmissions, compared to conventionally designed interconnects.

As the number of cores integrated on a chip is increasing, high performance, power and area efficient long-range communication links become a necessity. Using fully bit-parallel communication in long-range links that traverse two or more cores become costly because it takes a large chip area and causes routing difficulties, severe crosstalk noise as well as considerable leakage power. In order to address most of these issues, a high-throughput and low power serial link has been designed in this thesis. This serial link design is based on a combination of pulse dual-rail encoding, wave-pipelining, pulse signaling and differential current-mode signaling besides formulating novel control circuits for the serializer and deserializer. A semi-serial link which consists of eight bit-serial links outperforms the fully bit-parallel link both in throughput and energy efficiency as well as in active and wiring area when the communication distance is 4mm and longer.

Furthermore, circuit level techniques to ensure signal integrity of the designed interconnects in the presence of PVT variations have been devised and implemented. To guarantee signal integrity despite process, wearout, and ageing caused variations, interconnect calibration technique at every power start-up of a system has been formulated. An error detecting mechanism as well as a calibration algorithm and methodology have been developed and verified through simulations. This technique has a delay, power and area overhead. The delay overhead occurs only once at every power start-up. The power and area overhead is low for large bit width transmissions. The runtime power supply voltage and temperature variation tolerance technique is based on monitoring the effect of runtime variations and then reconfiguring the receiver when an error is detected. The power and area overhead of this technique is low especially for large bit width transmissions. It has even been proved that power can be saved compared to the worst-case margin allocation approach. Therefore, these schemes make the interconnects adaptive to the effect of both static and dynamic variations, enabling continuous and reliable operation of the interconnects.

8.1 Future Work

Intuitively it is known that using the presented interconnects on a NoC, which consists of M nodes communicating with each other through the network, will increase the overall network throughput and decrease the average latency per packet. The network can also accept more traffic with congestion free communication and dissipates less power per packet than a NoC which uses conventionally designed interconnects. In order to quantify these improvements, it is necessary to develop a clockless NoC simulator since all existing open source NoC simulators are cycle-based. Hence, one future direction of the work is developing an event-based NoC simulator and examine network performance, energy and other related parameters with synthetic and application specific traffic patterns. Furthermore, analysis of different topologies' performance and energy efficiencies using the designed bit parallel interconnects between neighboring routers and the semi-serial interconnect as long-range links will be carried out. The other direction of the future work is integrating the developed PVT variation tolerance techniques of the interconnects with data link layer tolerance schemes. Determining the advantages and limitations as well as optimization and verification will also be performed.

Bibliography

- [1] S. Borkar, *Design challenges of technology scaling*. Micro, IEEE, 19(4):23–29, 1999.
- [2] G.E. Moore, *Cramming More Components on Integrated Circuits*. Electronics (38), 8:114–117, 1965.
- [3] Karl Goser, Peter Glösekötter and Jan Dienststuhl, *Nanoelectronics and Nanosystems: From Transistors to Molecular and Quantum Devices*. Springer-Verlag Berlin, 2004.
- [4] R. Kumar and G. Hinton, *A Family of 45nm IA Processors*. 2009 IEEE International Solid-State Circuits Conference, Digest of Technical Papers, Vol. 1, pp. 58-59, Feb. 2009.
- [5] S. Rusu, S. Tam, H. Muljono, J. Stinson, D. Ayers, J. Chang, R. Varada, M. Ratta, S. Kottapalli, and S. Vora, *A 45nm 8-Core Enterprise Xeon® Processor*. 2007 IEEE Journal of Solid-State Circuits, Vol. 45, No. 1, pp. 7-14, Jan. 2010.
- [6] T. Sakurai, *Perspectives on power-aware electronics*. 2003 IEEE International Solid-State Circuits Conference, Digest of Technical Papers, Vol. 1, pp. 26-29.
- [7] N. Magen, A. Kolodny, U. Weiser, and N. Shamir, *Interconnect power dissipation in a microprocessor*. in IEEE/ACM International Workshop on System Level Interconnect Prediction, pp. 7-13, Feb. 2004.
- [8] S. Vangal, A. Singh, J. Howard, S. Dighe, N. Borkar, and A. Alvandpour, *A 5.1GHz 0.34mm² router for network-on-chip applications*. in Symposium VLSI Circuits Dig. Tech. Papers, pp.42–43, Jun. 2007.

- [9] B. P. Wong, F. Zach, V. Moroz, A. Mittal, G. W. Starr, and A. Kahng, *Nano-CMOS Design for Manufacturability*. A John Wiley Sons Inc Publication, 2009.
- [10] M. Orshansky, S. Nassif, D. Boning, *Design for Manufacturability and Statistical Design: A Constructive Approach*. Springer, 2007.
- [11] Dennis Sylvestera, Kanak Agarwalb and Saumil Shaha, *Variability in nanometer CMOS: Impact, analysis, and minimization*. Integration, the VLSI Journal, Vol. 41, No. 3, pp. 319-339, May 2008.
- [12] S. R. Nassif, *Model to Hardware Matching; For nano-meter Scale Technologies*. 2006 International Conference on Simulation of Semiconductor Processes and Devices, pp. 5-8, Sept. 2006.
- [13] K. Bernstein, et al, *High-performance CMOS variability in the 65-nm regime and beyond*. IBM Journal of Research and Development archive, Vol. 50, No. 4/5, pp. 433-449, July 2006.
- [14] V. Agarwal, M. S. Hrishikesh, S. W. Keckler and D. Burger, *Clock rate versus IPC: the end of the road for conventional microarchitectures*. Proceedings of the 27th International Symposium on Computer Architecture, pp. 248-259, 2000.
- [15] P. Guerrier and A. Greiner, *A generic architecture for on-chip packet switched interconnections*. in Design, Automation and Test in Europe Conference and Exhibition, DATE'00, Paris, France, Mar. 2000, pp.250-256.
- [16] A. Hemani et al, *Network on a chip: An architecture for billion transistor era*. in 18th NORCHIP Conference, Turku, Finland, Nov. 2000, pp. 166-173.
- [17] M. Sgroi et al, *Addressing the System-on-a-chip interconnect woes through communication-based design*. in 38th Design Automation Conference, DAC 2001, Las Vegas, NV, June 2001, pp. 667-672.
- [18] *TILE64TM Processor*. [Online]. Available <http://www.tilera.com/products/TILE64.php>

- [19] A. Upadhyay, S. R. Hasan, and M. Nekili, *A novel asynchronous wrapper using 1-of-4 data encoding and single-track handshaking*. The 2nd Annual IEEE Northeast Workshop on Circuits and Systems, pp. 205- 208, June, 2004.
- [20] M. Ferretti, and P. A. Beerel, *Single-Track Asynchronous Pipeline Templates Using 1-of-N Encoding*. Proceedings of the 2002 Design, Automation and Test in Europe Conference and Exhibition, pp. 1008-1015, August, 2002.
- [21] T. Hanyu, T. Takahashi and M. Kameyama, *Bidirectional data transfer based asynchronous VLSI system using multiple-valued current mode logic*. 33rd International Symposium on Multiple-Valued Logic, pp. 99-104, May 2003.
- [22] E. Nigussie, J. Plosila and J. Isoaho, *On Asynchronous Full-Duplex Dual-Rail Link with Multiple-Valued Current-Mode Signaling*. 23rd NORCHIP Conference, pp. 222-225, Nov. 2005.
- [23] E. Nigussie, J. Plosila and J. Isoaho, *Full-duplex link implementation using dual-rail encoding and multiple-valued current-mode logic*. 2006 IEEE International Symposium on Circuits and Systems ISCAS 2006, 4 pp, May. 2006.
- [24] K. Nabors, and J. White, *FastCap: a multipole accelerated 3-D capacitance extraction program*. 2006 IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, Vol. 10, No. 11, pp. 1447 - 1459, Nov. 1991.
- [25] H. Ron, J. Gainsley and R. Drost, *Long wires and asynchronous control*. Proc. 10th IEEE International Symposium on ASYNC, pp. 240-249, Apr. 2004.
- [26] R. Dobkin, R. Ginosar, and C. P. Sotiriou, *High Rate Data Synchronization in GALS SoCs*. IEEE Transactions on Very Large Scale Integration (VLSI) Systems, Vol. 14, No. 10, pp. 1063-1074, October 2006.
- [27] A. Sheibanyrad, and A. Greiner, *Two efficient synchronous \Leftrightarrow asynchronous converters well-suited for networks-on-chip in GALS architectures*. Integration the VLSI Journal, Vol. 41, No. 1, pp. 17-26, 2008.

- [28] W. Ning G. Fen, and W. Fei, *Design of a GALS Wrapper for Network on Chip*. 2009 IEEE WRI World Congress on Computer Science and Information Engineering, Vol. 3, pp. 592-595, March 2009.
- [29] P. Liljeberg, J. Plosila, and J. Isoaho, *Self-timed communication platform for implementing high-performance systems-on-chip*. Integration the VLSI Journal, Vol. 38, No. 1, pp. 43-67, 2004.
- [30] R. Bashirullah, W. Liu, and R.K. Cavin, *Current-mode signaling in deep submicrometer global interconnects*. IEEE Transactions on VLSI Systems, Vol. 11, No. 3, pp. 406-417.
- [31] R. Bashirullah, *Reduced delay sensitivity to process induced variability in current sensing interconnects*. Electronics Letters, Apr. 2006, Vol. 42, No. 9.
- [32] A. Katoch, E. Seevinick and H. Veendrick, *Fast signal propagation for point to point on-chip long interconnects using current sensing*. European Solid-State Circuits Conference, 2002.
- [33] A. Katoch, H. Veendrick and E. Seevinick, *High speed current-mode signaling circuits for on-chip interconnects*. IEEE International Symposium on Circuits and Systems, Vol. 4, pp. 4138-4141, May 2005.
- [34] A.P. Jose, G. Patounakis and K.L. Shepard, *Near speed-of-light on-chip interconnects using pulsed current-mode signaling*. IEEE Symposium on VLSI Circuits Digest of Technical Papers, pp. 108-111, June 2005.
- [35] M.K. Gowan, L.L. Biro and D.B. Jackson, *Power considerations in the design of the alpha 21 264 microprocessor*. Proc. Design Automation Conference, 1998, pp. 726-731.
- [36] E. Nigussie, J. Plosila and J. Isoaho, *Delay-Insensitive On-Chip Communication Link using Low-Swing Simultaneous Bidirectional Signaling*. IEEE Computer Society Annual Symposium on VLSI, pp. 217-222, Mar. 2006.
- [37] M. -H. Oh and D. -S. Har, *A Novel Mechanism for Delay-Insensitive Data Transfer Based on Current-Mode Multiple Valued Logic*. PATMOS 2004, pp. 691-700.

- [38] V. Venkatraman, and W. Burleson, *Robust Multi-Level Current-Mode On-Chip Interconnect Signaling in the Presence of Process Variations*. Proceedings of the 6th International Symposium on Quality of Electronic Design, pp. 522-527, 2005.
- [39] T. Kuboki, A. Tsuchiya, and H. Onodera, *A 10Gbps/channel On-Chip Signaling Circuit with an Impedance-Unmatched CML Driver in 90nm CMOS Technology*. IEEE Asia and South Pacific Design Automation Conference, pp. 120-121, 2007.
- [40] N. Tzartzanis, and W. W. Walker, *Differential current-mode sensing for efficient on-chip global signaling*. IEEE Journal of Solid-State Circuits, Vol. 40, No. 11, pp. 2141-2147, 2005.
- [41] L. Zhang, J. Wilson, R. Bashirullah, and P. Franzon, *Differential current-mode signaling for robust and power efficient on-chip global interconnects*. IEEE 14th Meeting on Electrical Performance of Electronic Packaging, pp. 315-318, 2005.
- [42] A. Maheshwari, and W. Burleson, *Differential current-sensing for on-chip interconnects*. IEEE Transactions on VLSI Systems, Vol. 12, No. 12, pp. 1321-1329, 2004.
- [43] V. K. Venkatraman, *Design and Integration of Current-Mode On-Chip Interconnect Signaling in Nanometer Technologies*. PhD Thesis, University of Massachusetts Amherst, 2007.
- [44] W.J. Dally, B. Towles, *Route Packets, not Wires: On-Chip Interconnection Networks*. Proc. 38th Design Automation Conference, June, 2001.
- [45] M. Kamon, M. J. Tsuk and J. K. White, *FASTHENRY: A multipole-accelerated 3-D inductance extraction program*. IEEE Transactions on Microwave Theory and Techniques, Vol. 42, No. 9, pp. 1750-1758.
- [46] A. Djordjevic, M. Bazdar, T. Sarkar and R. Harrington, *Linpar for Windows: Matix parameters for multiconductor transmission lines*. Software and Users Manual, Version 2.0, Norwood, MA: Artech House Publishers, 1999.
- [47] T. Verhoeff, *Delay-insensitive codes - An overview*. Distributed Computing, 3(1):1-8, 1988.

- [48] R. Venkatesan, J.A. Davis and J.D. Meindl, *Compact distributed RLC interconnect Models - Part IV: unified models for time delay, crosstalk and repeater insertion*. IEEE Transactions on Electron Devices, Vol. 50, No. 4, April 2003.
- [49] A. Narasimhan, S. Divekar, P. Elakkumanan, and R. Sridhar, *A low-power current-mode clock distribution scheme for multi-GHz NoC-based SoCs*. IEEE 18th International conference on VLSI Design, pp. 130-133, 2005.
- [50] L. Benini and G. Micheli, *Networks on Chips: Technology and Tools*. Morgan Kaufmann Publishers, 2006.
- [51] M. Krstic, E. Grass, F. K. Gurkaynak, and P. Vivet, *Globally asynchronous, locally synchronous circuits: Overview and outlook*. IEEE Design and Test of Computers, Vol. 24, No. 5, pp. 430-441, 2007.
- [52] E. Beigne, et al, *An Asynchronous Power Aware and Adaptive NoC Based Circuit*. IEEE Journal of Solid-State Circuits, Vol. 44, No. 4, pp. 1167-1176, 2009.
- [53] A. Sheibanyrad, I. Miro Panades, A. Greiner, *Systematic Comparison between the Asynchronous and the Multi-Synchronous Implementations of a Network on Chip Architecture*. IEEE Design, Automation and Test in Europe Conference and Exhibition, 2007, pp. 1-6, 2007.
- [54] A. J. Martin and M. Nyström, *Asynchronous Techniques for System-on-Chip Design*. Proceedings of the IEEE, Vol. 94, No. 6, pp. 1089-1120, 2006.
- [55] W. J. Bainbridge, and S. B. Furber, *Delay insensitive system-on-chip interconnect using 1-of-4 data encoding*. International Symposium on Asynchronous Circuits and Systems, pp. 118-126, Mar. 2001.
- [56] W. J. Bainbridge, and S. B. Furber, *Chain: a delay-insensitive chip area interconnect*. IEEE Micro, Vol. 22, No. 5, pp. 16-23, 2002.
- [57] A. M. Pappu, X. Zhang, A. V. Harrison and A. B. Apsel, *Process-Invariant Current Source Design: Methodology and Examples*. IEEE Journal of Solid-State Circuits, Vol. 42, No. 10, 2007.

- [58] International Technology Roadmap for Semiconductors 2007, <http://public.itrs.net>.
- [59] D. N. Truong, et al, *A 167-Processor Computational Platform in 65 nm CMOS*. IEEE Journal of Solid-State Circuits, Vol. 44, No. 4, pp. 1130-1144, 2009.
- [60] Y. Hoskote, S. Vangal, A. Singh, N. Borkar and S. Borkar, *A 5-GHz Mesh Interconnect for A TeraFLOPS Processor*. IEEE MICRO, Vol. 27, No. 5, pp. 51-61, 2007.
- [61] S. Bell, et al, *Tile64TM Processor: A 64-core SoC with Mesh Interconnect*. IEEE Solid-State Circuits Conference 2008.
- [62] D. Lattard, et al, *A Reconfigurable Baseband Platform Based on An Asynchronous Network-on-Chip*. IEEE Journal of Solid-State Circuits, Vol. 43, No. 1, pp. 223-235, 2008.
- [63] E. Salminen, A. Kulmala, and T. Hamalainen, *On Network-on-Chip Comparison*. Euromicro Conference on Digital System Design, pp. 503-510, 2007.
- [64] S. R. Vangal, et al, *An 80-Tile Sub-100-W TeraFLOPS Processor in 65-nm CMOS*. IEEE Journal of Solid-State Circuits, Vol. 43, No. 1, pp. 29-41, 2008.
- [65] M. B. Taylor et al, *The Raw Microprocessor: A computational fabric for software circuits and general purpose programs*. IEEE Micro Vol. 22, No. 2 pp. 25-35, 2002.
- [66] W. J. Dally and B. Towles, *Principles and Practices of Interconnection Networks*. Morgan Kaufmann, San Francisco, CA, 2004.
- [67] J. Kim, J. Balfour, and W. J. Dally, *Flattened Butterfly Topology for On-Chip Networks*. 40th Annual IEEE/ACM International Symposium on Microarchitecture, pp. 172-182, 2007.
- [68] J. Balfour and W. J. Dally, *Design tradeoffs for tiled CMP on-chip networks*. 20th ACM International Conference on Supercomputing, pp. 187-198, 2006.

- [69] L. Bononi, and N. Concer, *Simulation and analysis of network on chip architectures: ring, spidergon and 2D mesh*. IEEE Design, Automation and Test in Europe 2006 (DATE'06), Vol. 2, 6 pp., 2006.
- [70] M. Coppola, *Keynote lecture: Spidergon STNoC: the communication infrastructure for multiprocessor architectures*. In International Forum on Application-Specific Multi-Processor SoC, 2008.
- [71] W. J. Dally, *Enabling Technology for On-Chip Interconnection Networks*. Invited talk, 1st IEEE/ACM International Symposium on Network-on-Chip, May 2007 2007.nocsymposium.org/keynote1/dally_nocs07.ppt
- [72] W. J. Dally, *On-Chip Interconnection Networks Low-Power Interconnect*. Special Session, International Symposium on Low Power Electronics and Design 2007, August 2007 [http : //www.islped.org/X2007/DallyISLPE07.pdf](http://www.islped.org/X2007/DallyISLPE07.pdf)
- [73] U. Y. Ogras, and R. Marculescu, *It's a small world after all": NoC performance optimization via long-range link insertion*. IEEE Transactions on Very Large Scale Integration (VLSI) Systems, Vol. 14, No. 7, pp. 693-706, 2006.
- [74] L. Benini and G. Micheli, *Networks on Chips: a new SoC Paradigm*. IEEE Computer, Vol. 35, No. 1, pp. 70-78, Jan 2002.
- [75] M. Krstic, E. Grass, F. K. Gurkaynak, and P. Vivet, *Globally asynchronous, locally synchronous circuits: Overview and outlook*. IEEE Design and Test of Computers, vol. 24, No. 5, pp. 430-441, 2007.
- [76] E. Beigne, et al, *An Asynchronous Power Aware and Adaptive NoC Based Circuit*. IEEE Journal of Solid-State Circuits, Vol. 44, No. 4, pp. 1167-1176, 2009.
- [77] S.-J. Lee, K. Kim, H. Kim, N. Cho, and H.-J. Yoo, *Adaptive Network-on-Chip with Wave-Front Train Serialization Scheme*. IEEE 2005 Symposium on VLSI Circuits, Digest of Technical Papers, pp. 104-107, 2005.
- [78] R. Dobkin, Y. Perelman, T. Liran, R. Ginosar, and A. Kolodny, *High Rate Wave-pipelined Asynchronous On-chip Bit-serial Data Link*. 13th IEEE International Symposium on Asynchronous Circuits and Systems, pp. 3-14, 2007.

- [79] C. K. K. Yang, *Design of High-Speed Serial Link Transceiver*. PhD Thesis, Stanford University, 1998.
- [80] M. J. E. Lee, *An efficient I/O and Clock Recovery for TERABIT Integrated Circuits Design*. PhD Thesis, Stanford University, 2001.
- [81] T. Bjerregaard, *The MANGO Clockless Network-on-Chip: Concepts and Implementation*. PhD Thesis, Technical University of Denmark, 2005.
- [82] E. Beigné, et al, *An Asynchronous NOC Architecture Providing Low Latency Service and its Multi-Level Design Flow*. 11th IEEE International Symposium on Asynchronous Circuits and Systems, pp. 54-63, 2005.
- [83] D. Lattard, et al, *A Reconfigurable Baseband Platform Based on An Asynchronous Network-on-Chip*. IEEE Journal of Solid-State Circuits, Vol. 43, No. 1, pp. 223-235, 2008.
- [84] R. Dobkin, R. Ginosar and A. Kolodny, *QNoC Asynchronous Router*. Integration the VLSI Journal, Vol. 42, No. 2, pp. 103-115, 2009.
- [85] E. Beigne, F. Clermidy, P. Vivet, A. Clouard, and M. Renaudin, *An asynchronous NOC architecture providing low latency service and its multi-level design framework*. 11th IEEE International Symposium on Asynchronous Circuits and Systems, pp. 54-63, 2005.
- [86] W. J. Dally and J. W. Poulton, *Digital Systems Engineering*. Cambridge University Press, 1998.
- [87] S. Borkar, *Designing reliable systems from unreliable components the challenges of transistor variability and degradation*. IEEE Micro, Vol. 25, No. 6, pp. 10-16, 2005.
- [88] P. Wang, G. Pei, and E. C. -C. Kan, *Pulsed Wave Interconnect*. IEEE Transactions on VLSI, Vol. 12, No. 5, pp. 453-463, 2004.
- [89] M. Chen and Y. Cao, *Analysis of Pulse Signaling for Low-Power On-Chip Global Bus Design*. 7th IEEE International Symposium on Quality Electronic Design, 6pp, 2006.
- [90] M. Bazes, *Two Novel Fully Complementary Self-Biased CMOS Differential Amplifiers*. IEEE Journal of Solid-State Circuits, Vol. 26, No. 2, pp. 165-168, 1991.

- [91] Y. I. Ismail, E. G. Friedman, and J. L. Neves, *Figures of Merit to Characterize the Importance of On-Chip Inductance*. IEEE Transactions on VLSI, Vol. 7, No. 4, pp. 442-449, 1999.
- [92] *International Technology Roadmap for Semiconductors 2008*. <http://public.itrs.net>.
- [93] B. Wong, A. Mittal, Y. Cao, and G. W. Starr, *Nano-CMOS Circuit and Physical Design*. IEEE Wiley-IEEE Press, 2004.
- [94] J. M. Rabaey, A. Chandrakasan, and B. Nikolic, *Digital Integrated Circuits: A Design Perspective*. Prentice Hall Ltd, 2 edition, 2003.
- [95] Li-Rong Zheng, *Design, Analysis and Integration of Mixed-Signal Systems for Signal and Power Integrity*. PhD thesis, Royal Institute of Technology(KTH), Sweden, 2001.
- [96] A. Deutsch, G.V Kopcsay, P.W. Coteus, C.W. Surovic, P.E. Dahlen, and D.L. Heckmann, *Frequency-dependent losses on high-performance interconnections*. IEEE Transactions on Electromagnetic Compatibility, Vol. 43, No. 4, pp. 446-465, 2001.
- [97] A. V. Mezhiba and E. G. Friedman, *Power Distribution Networks in High Speed Integrated Circuits*. Kluwer Academic Publisher, 2003.
- [98] E. Rosa, *The self and mutual inductance of linear conductors*. Bulletin of the National Bureau of Standards, Vol. 4, pp. 301-304, 1908.
- [99] A. Ruehli, *Inductance calculations in a complex integrated circuit environment*. IBM Journal of Research and Development, Vol. 16, No. 5, pp. 470-481, 1972.
- [100] C. K. Cheng et al, *Interconnect Analysis and Synthesis*. John Wiley New York, 2000.
- [101] B. Young, *Digital Signal Integrity: Modeling and Simulation with Interconnects and Packages*. Prentice Hall PTR Upper Saddle River, NJ, USA, 2000.
- [102] A. Deutsch et al., *When are transmission-line effects important for on-chip interconnections*. IEEE Transactions on Microwave Theory Tech., Vol. 45, No. 10, pp. 1836– 1846, Oct 1997.

- [103] A. J. Joshi, G. G. Lopez, and A. Davis, *Design and optimization of on-chip interconnects using wave-pipelined multiplexed routing*. IEEE Transactions on VLSI Systems, Vol. 15, No. 9, pp. 990-1002, Sept. 2007.
- [104] R. Dobkin, A. Morgenshtein, A. Kolodony, and R. Ginosar, *Parallel versus serial on-chip communication*. 10th International Workshop on System Level Interconnect Prediction, pp. 43-50, 2008.
- [105] *The International Technology Roadmap for Semiconductors(ITRS), 2005..*
- [106] D. Pamunuwa, H. Tenhunen, *Repeater insertion to minimize delay in coupled interconnects*. VLSI Design, Jan. 2001, pp. 513-517.
- [107] W. J. Bainbridge, W. B. Toms, D. A. Edwards, S. B. Furber, *Delay-insensitive, point-to-point interconnect using m-of-n codes*. Ninth International Symposium on Asynchronous Circuits and Systems, pp. 132 - 140, May 2003.
- [108] B. R. Quinton, M. R. Greenstreet, S. J. E. Wilton, *Asynchronous IC interconnect network design and implementation using a standard ASIC flow*. IEEE International Conference on Computer Design: VLSI in Computers and Processors, pp. 267 - 274, Oct. 2005.
- [109] B. R. Quinton, M. R. Greenstreet, S. J. E. Wilton, *Practical Asynchronous Interconnect Network Design*. IEEE Transactions on Very Large Scale Integration (VLSI) Systems, Vol. 16, No. 5, pp. 579 - 588, May. 2008.
- [110] *The International Technology Roadmap for Semiconductors(ITRS), 2007..* Design, pp. 31, <http://public.itrs.net>.
- [111] O.S. Unsal, et al, *Impact of Parameter Variations on Circuits and Microarchitecture*. IEEE Micro, Vol. 26, No. 6, pp. 30 - 39, Nov.-Dec. 2006.
- [112] A. A. Mutlu, M. Rahman, *Statistical methods for the estimation of process variation effects on circuit operation*. IEEE Transactions on Electronics Packaging Manufacturing, Vol. 28, No. 4, pp. 364 - 375, Oct. 2005.

- [113] K. Bowman, J. Meindl, *Impact of within-die parameter fluctuations on the future maximum clock frequency distribution.* in Proceedings of Custom Integrated Circuits Conference, pp. 229-232, 2001.
- [114] H.-S. Wong, D. J. Frank, P. Solomon, C. Wann, and J. Wesler, *Nanoscale CMOS.* in Proceedings of the IEEE, Vol. 87, No. 4, pp. 537-570, 1999.
- [115] J. A. Croon, G. Storms, S. Winkelmeier, and I. Pollentier, *Line-edge roughness: characterization, modeling, and impact on device behavior.* in Proceedings of International Electron Devices Meeting, pp. 307-310, 2002.
- [116] P. Oldiges, L. Qimghuang, K. Petrillo, M. Leong, and M. Hargrove, *Modeling line edge roughness effects in sub 100 nanometer gatelength devices.* in Proceedings of the International Conference on Simulation of Semiconductor Processes and Devices, pp. 131-134, 2000.
- [117] S. Lakshminarayanan, P.J. Wright, J. Pallinti, *Electrical characterization of the copper CMP process and derivation of metal layout rules.* IEEE Transaction on Semiconductor Manufacturing, Vol. 16, No. 4, pp. 668-676, Nov. 2003.
- [118] X. Qi, A. Gyure, Y. Luo, S. C. Lo, M. Shahram, K. Singhal, *Measurement and characterization of pattern dependent process variations of interconnect resistance, capacitance and inductance in nanometer technologies.* in Proceedings of the 16th ACM Great Lakes symposium on VLSI, pp. 14-18, 2006.
- [119] D. Boning, *Pattern dependent characterization of copper interconnect.* Tutorial, International Conference on Microelectronic Test Structures, Mar. 2003.
- [120] L. Scheffer, S. Nassif, A. Strojwas, B. Koenemann, and N. NS, *Design for manufacturing at 65 nm and below.* IEEE 42nd Design Automation Conference, Tutorial 5, June 2005.
- [121] N. NS, *BEOL variability and impact on RC extraction.* IEEE 42nd Design Automation Conference, June 2005.
- [122] K. Bowman, S. Duvall, and J. Meindl, *Impact of die-to-die and within-die parameter fluctuations on the maximum clock frequency distribution*

- for gigascale integration.* IEEE Journal of Solid-State Circuits, Vol. 37, No. 2, pp. 668-676, Feb. 2002.
- [123] N. N. Hoang, A. Kumar, P. Christie, *The Impact of Back-End-of-Line Process Variations on Critical Path Timing.* 2006 International Interconnect Technology Conference, pp. 193-195.
- [124] V. Mehrotra and D. Boning, *Technology scaling impact of variation on clock skew and interconnect delay.* International Interconnect Technology Conference, June 2001.
- [125] E. Demircan, *Effects of Interconnect Process Variations on Signal Integrity.* 2006 IEEE International System on Chip Conference, pp. 281-284, Sept. 2006.
- [126] J. Rabaey, *Low Power Design Essentials.* Springer Book Series on Integrated Circuits and Systems, 2009.
- [127] Y. Taur and T. H. Ning, *Fundamentals of Modern VLSI Devices.* Cambridge University Press, 1998.
- [128] N. James, P. Restle, J. Friedrich, B. Huott, and B. McCredie, *Comparison of Split-Versus Connected-Core Supplies in the POWER6TM Microprocessor.* 2007 IEEE International Solid-State Circuits Conference, Power Management Papers, pp. 298-300.
- [129] H. Harizi, R. Haussler, M. Olbrich, and E. Barke, *Efficient Modeling Techniques for Dynamic Voltage Drop Analysis.* 44th ACM/IEEE Design Automation Conference, pp. 706-711, June 2007.
- [130] T. H. Morshed, et al, *BSIM4.6.4 MOSFET Model-User's Manual.* Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, 2009.
- [131] Y. P. Tsividis, *Operation and Modeling of the MOS Transistor.* McGraw-Hill, New York, 1999.
- [132] N. Lu, M. Angyal, G. Matusiewicz, V. McGahay, and T. Standaert, *Characterization, Modeling and Extraction of Cu Wire Resistance for 65 nm Technology.* IEEE 2007 Custom Integrated Circuits Conference (CICC), pp. 57-60, Sept. 2007.

- [133] M. Ashouei, *Algorithms and Methodology for Post-Manufacture Adaptation to Process Variations and Induced Noise in Deeply Scaled CMOS Technologies*. PhD Thesis, Georgia Institute of Technology, Dec. 2007.
- [134] Arteris. [Online]. Available: <http://WWW.arteris.com/>
- [135] M. Coppola, C. Pistrutto, R. Locatelli and A. Scandurra, *STNoCTM: An Evolution Towards MPSoC Era*. NoC Workshop, DATE, Mar. 2006.
- [136] M. Coppola, R. Locatelli, G. Maruccia, L. Pieralisi, A. Scandurra, *Spidergon: a novel on-chip communication network*. International Symposium on System-on-Chip, Nov. 2004.
- [137] S. Xu, V. Venkatraman, and W. Burlison, *Energy-Aware Differential Current Sensing for Global On-Chip Interconnects*. 49th IEEE International Midwest Symposium on Circuits and Systems, MWSCAS'06, Vol.1, pp. 718-722, Aug. 2006.
- [138] J. Xu, and W. Wolf, *A wave-pipelined on-chip interconnect structure for networks-on-chips*. 11th Symposium on High Performance Interconnects, pp. 10-14, Aug. 2003
- [139] D. Schinkel, E. Mensink, E. A. M. Klumperink, E. van Tuijl, and B. Nauta, *A 3-Gb/s/ch transceiver for 10-mm uninterrupted RC-limited global on-chip interconnects*. IEEE Journal of Solid-State Circuits, Vol. 41, pp. 297-306, Jan. 2006.
- [140] E. Mensink, D. Schinkel, E. Klumperink, E. van Tuijl, and B. Nauta, *A 0.28pJ/b 2Gb/s/ch Transceiver in 90nm CMOS for 10mm On-Chip Interconnects*. IEEE International Solid State Circuits Conference (ISSCC), Digest of Technical Papers, pp. 414-415, Feb. 2007.
- [141] E. Mensink, *High-Speed Global On-Chip Interconnects and Transceivers*. PhD Thesis, University of Twente, June 2007.
- [142] Haikun Zhu; Rui Shi; Chung-Kuan Cheng; Hongyu Chen, *Approaching Speed-of-light Distortionless Communication for On-chip Interconnect*. Asia and South Pacific Design Automation Conference, 2007, ASP-DAC '07, pp. 684-689, Jan. 2007.

- [143] Joonsung Bae; Joo-Young Kim; Hoi-Jun Yoo, *A 0.6pJ/b 3Gb/s/ch transceiver in 0.18 μ m CMOS for 10mm on-chip interconnects*. IEEE International Symposium on Circuits and Systems, 2008 (ISCAS 2008), pp. 2861-2864, May. 2008.
- [144] Medardoni Simone, Marcello Lajolo, and Davide Bertozzi, *Variation tolerant NoC design by means of self-calibrating links*. Conference and Exhibition on Design, automation and test in Europe (DATE 2008), pp. 1402-1407, Mar. 2008.
- [145] D. Mangano, R. Locatelli, A. Scandurra, C. Pistritto, M. Coppola, L. Fanucci, F. Vitullo, D. Zandri, *Skew Insensitive Physical Links for Network on Chip*. 1st International Conference on Nano-Networks and Workshops, NanoNet '06, pp. 1-5, Sept. 2006.
- [146] C. D'Alessandro, D. Shang, A. Bystrov, A. Yakovlev, and O. Maevsky, *Multiple-rail phase-encoding for NoC*. 12th IEEE International Symposium on Asynchronous Circuits and Systems (ASYNC), pp. 107-116, 2006.
- [147] P.B. McGee, M.Y. Agyekum, M.A. Mohamed, S.M. Nowick, *A Level-Encoded Transition Signaling Protocol for High-Throughput Asynchronous Global Communication*. 14th IEEE International Symposium on Asynchronous Circuits and Systems (ASYNC), pp. 116-127, Apr. 2008.
- [148] Y. Shi, S.B. Furber, J. Garside and L.A. Plana, *Fault tolerant delay insensitive interchip communication*. 15th IEEE International Symposium on Asynchronous Circuits and Systems (ASYNC), pp. 77-84, Apr. 2009.
- [149] 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, Apr. 2003.
- [150] L.-Z. Zhang, Y. Hu, C.-P. Chen, *Wave-pipelined on-chip global interconnect*. Asia and South Pacific Design Automation Conference, 2005 (ASP-DAC 2005), Vol. 1, pp. 127-132, Jan. 2005.
- [151] P. Cocchini, *Concurrent flip-flop and repeater insertion for high performance integrated circuits*. IEEE/ACM International Conference on Computer Aided Design (ICCAD), pp. 268 -273, 2002.

- [152] R. Lu, G. Zhong, C.-K. Koh, and K.-Y. Chao, *Flip-flop and repeater insertion for early interconnect planning*. Proceedings of Design, Automation and Test in Europe Conference and Exhibition, 2002, pp. 690–695, Mar. 2002.
- [153] C. Lin and H. Zhou, *Retiming for wire pipelining in system-on-chip*. IEEE International Conference on Computer Aided Design 2003 (CAD 2003), pp. 215–220, 2003.
- [154] L. Scheffer, *Methodologies and tools for pipelined on-chip interconnect*. Proceedings of the 2002 IEEE International Conference on Computer Design: VLSI in Computers and Processors (ICCD'02), pp. 152–157, Sept. 2002.
- [155] A. Lines, *Asynchronous interconnect for synchronous SoC design*. IEEE Micro, Vol. 24, No. 1, pp. 32–41, Jan./Feb. 2004.
- [156] J. D. Owens, W. J. Dally, R. Ho, D. N. (Jay) Jayasimha, S. W. Keckler, and L.-S. Peh *Research Challenges for On-Chip Interconnection Networks*. IEEE Micro, Vol. 27, No. 5, pp. 96–108, Sep./Oct. 2007.