

High-Speed and Low-Energy On-Chip Communication Circuits

by

Jae-sun Seo

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctorate of Philosophy
(Electrical Engineering)
in The University of Michigan
2010

Doctoral Committee:

Associate Professor Dennis Sylvester, Chair
Professor David Blaauw
Associate Professor Igor L. Markov
Assistant Professor David D. Wentzloff

© Jae-sun Seo 2010
All Rights Reserved

To my family with love and gratitude

ACKNOWLEDGEMENTS

This work has been supported and contributed by many people surrounding me.

Professor Dennis Sylvester has been everything but a wonderful advisor throughout my graduate studies. He continuously provided tremendous support, driving strength and motivation on my research. He has always been a very approachable mentor and a great communicator who loves to share his experience and knowledge.

I have worked on every research project throughout my Ph.D. with Professor David Blaauw, and he has always been enthusiastic about research matters. He focused on even the small details that are easily missed, and he always offered essential ideas and feedback.

Professor Igor Markov has been a great collaborator on a joint project that I worked on. He brought in great amount of energy to the project and was a superb investigator. David Wentzloff graciously agreed to be on my defense committee.

During internships in the industry, I gained invaluable experience and learned abundant knowledge, hence I would like to appreciate my mentors: Dr. Ram Krishnamurthy, Dr. Himanshu Kaul from Intel Corporation, and Dr. Ron Ho from Sun Microsystems. Including the aforementioned people, I also thank all the co-authors included in the various projects I have been working on.

I had pleasure to communicate and collaborate with the labmates in our research group. Mingoo Seok has been my most frequent conversationalist, providing helpful feedback to each other on any research idea we came up with. Sanjay Pant and Carlos Tokunaga always delivered the necessary knowledge about chip designs to me. I also enjoyed either co-working on projects or having research discussions with Scott Hanson, Mike Wieckowski, Zhiyoong Foo, Daeyeon Kim, and Yoonmyung Lee.

Last, but not the least, I owe special gratitude to my family, for always standing on my side and providing unconditional support throughout the long graduate life.

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGEMENTS	iii
LIST OF FIGURES	vii
LIST OF TABLES	xi
CHAPTERS	
1 Introduction	1
1.1 Implications of Scaling on Wire Resistance, Capacitance, and Inductance	2
1.2 Performance and Energy of Wires	3
1.3 Clock consideration for global interconnects	5
1.4 Contribution of This Work and Organization	7
2 Edge Encoding	9
2.1 Motivation and Previous Work	9
2.2 Edge Encoding Approach	12
2.2.1 Basic Idea	12
2.2.2 Theoretical Energy Savings	13
2.2.3 Edge Encoding Technique	16
2.2.4 Dual-edge flip-flops and timing concerns in edge encoding	18
2.3 Experimental Results	21
2.3.1 Zero Latency (ZL) Scheme	22
2.3.2 One Cycle Latency (OCL) Scheme	24
2.3.3 Leakage Power Comparison	26
2.3.4 Sensitivity to Variation	27
2.4 Summary	30
3 Alternating Repeater Technique	32
3.1 Related Work	32
3.2 Alternate Repeater Concept	33
3.3 Bus Comparisons	35
3.3.1 Equal Delay and Area	36

3.3.2	Process Skew and Repeater Placement Sensitivity	38
3.3.3	Multi-Cycle Bus	40
3.4	Summary	41
4	Crosstalk-Aware Pulse Width Modulation based Signaling	42
4.1	Related work	42
4.2	PWM-based signaling concept	43
4.3	Proposed Encoder and Decoder Circuits for PWM-based signaling	47
4.3.1	Mono-PWM encoder and decoder	47
4.3.2	Hybrid-PWM encoder and decoder	48
4.4	Crosstalk and Variability Considerations	50
4.4.1	Crosstalk aware signaling	50
4.4.2	Mitigating variability and self-calibration	51
4.5	Measurement Results	54
4.6	Summary	58
5	Self-Timed Regenerators	62
5.1	Consideration on Repeater-less Signaling	62
5.2	Self-timed Regenerator Design	63
5.2.1	Transmission Line Configuration	63
5.2.2	Self-Timed Regenerator (STR) Circuit Operation	64
5.2.3	Sizing of the Circuit	66
5.2.4	Effect of STR on Signal Integrity	68
5.3	Experimental Results	68
5.3.1	Repeater and STR Design Scheme	69
5.3.2	Power, Area, and Peak Current	71
5.3.3	Performance	72
5.3.4	Low V_t Repeaters and Leakage power	73
5.4	Clock Network Application	74
5.5	Summary	76
6	High Bandwidth Low Swing Signaling	77
6.1	Motivation and Previous Work on Low Swing Signaling	77
6.2	Concept and Advantages of the Proposed RZ signaling	79
6.3	Matlab Analysis	81
6.4	Transceiver Circuit Design in SDR Scheme	83
6.4.1	Transmitter Design	84
6.4.2	Hysteresis Receiver	85
6.4.3	Biasing of wire and receiver	86
6.5	Adaptive Pre-Emphasis in DDR Scheme for Further Bandwidth Improvement	87
6.6	Measurement Results	89
6.7	Summary	92
7	Effect of Long Wires on Technology Mapping	96

7.1	Motivation	96
7.2	Analysis of Single Paths	98
7.3	Evaluating utility of large cells in technology mapping	99
7.3.1	Methodology	100
7.3.2	Experimental Results	101
7.4	Summary	108
8	Conclusions	109
	BIBLIOGRAPHY	112

LIST OF FIGURES

Figure

1.1	ITRS Trends of transistor delay, local interconnect delay, and global interconnect delay with technology scaling.	2
1.2	Scaling trends of unit-length wire resistance, ground capacitance, and lateral coupling capacitance for a semi-global metal layer with minimum width and spacing at advanced technology nodes.	3
1.3	Die micrograph of two recent multi-core microprocessors.	6
2.1	Conventional and proposed wire switching scenario in adjacent wires.	13
2.2	Ideal wire energy savings due to MCF reduction based on 65nm interconnect dimensions [14].	14
2.3	Block and timing diagrams of ZL scheme.	17
2.4	Flip-flop placement in conventional and ZL edge encoding scheme.	18
2.5	Block and timing diagrams of OCL scheme.	19
2.6	Flip-flop placement in conventional and OCL edge encoding scheme.	20
2.7	Schematic of time-borrowing pulsed dual-edge flip-flop.	20
2.8	4-bit cyclic bus model.	21
2.9	Comparison of the ZL edge encoding scheme to conventional busses in worst-case/average energy and clock frequency for flop distances (L1) of 2-5mm.	23
2.10	Energy-clock frequency comparison for a 3mm OCL edge-encoded bus.	24
2.11	Energy breakdown of a 5mm wire for conventional and OCL edge encoding scheme.	25
2.12	Leakage power comparison between a conventional and OCL edge-encoded bus. Ten repeaters are used in both cases and wirelength is 3mm.	26
2.13	Delay selection in staggered firing bus [36].	28
2.14	Skew selection in skewed repeater bus [16].	29
2.15	Sensitivity of improvements against process, supply, and temperature (PVT) variation for OCL edge-encoded bus (with MS-FF and TB-FF), staggered firing bus with 10ps and 20ps guard-banding (GB) [36], and skewed repeater bus [16] (flop distance: 3mm).	31
3.1	Worst-case switching pattern for conventional bus design.	34
3.2	Worst-case switching pattern for proposed alternating repeater bus.	35
3.3	Impact on delay with alternating repeater technique.	35

3.4	(a) Energy-delay and (b) Peak Current-Delay comparisons for 5mm bus in 65nm technology.	36
3.5	Sizing optimization for same repeater area.	37
3.6	Energy, delay and peak current reductions for the same repeater area with alternating repeaters for 5mm bus in 65nm technology.	37
3.7	Energy / delay gains across process skews for 5mm alternating repeater bus in 65nm technology.	38
3.8	Energy, delay and peak current reductions for the same repeater area with alternating repeaters for 5mm bus with nonequidistant repeater placement.	39
3.9	A core-to-core/core-to-cache 64b bus for a high-performance multi-core microprocessor.	40
3.10	64b driver/repeater block layout	40
4.1	(a) Conventional bus with repeaters (b) Proposed PWM-based bus with repeaters (both schemes have same footprint).	44
4.2	Concept of PWM-based signaling for proposed mono-PWM and hybrid-PWM schemes.	45
4.3	Timing diagram of monotonic PWM scheme.	46
4.4	Timing diagram of hybrid PWM system.	47
4.5	Encoder and decoder circuits for mono-PWM signaling.	48
4.6	Encoder and decoder circuits for hybrid-PWM signaling.	49
4.7	Due to crosstalk from adjacent bits, certain data patterns can result in changing pulse width over long wires.	51
4.8	Self-calibration methodology for each control signal.	52
4.9	Energy vs delay comparison of conventional, mono-PWM, and hybrid-PWM bus system.	53
4.10	Die photograph (1.5mm X 0.7mm).	54
4.11	Average energy comparison of conventional, mono-PWM, and hybrid-PWM scheme for microprocessor address, data and LFSR traces.	55
4.12	Measured waveforms of mono-PWM using on-chip oscilloscope.	56
4.13	Comparison of crosstalk-aware signaling on decoder input pulse with spread for all possible data patterns.	57
4.14	Global and local variation experiments are shown. (a) Sensitivity of mono-PWM system to global voltage and temperature variation. (b) Contour plot showing functionality and performance with local supply variation at 40°C. Additional guardbanding improves robustness at the expense of performance gains.	59
4.15	The performance and energy spread of 20 chips before and after self-calibration is shown.	60
4.16	Comparison of mono-PWM scheme delay distribution with self-calibration. Self-calibration reduces σ/μ of 21 chips by 2.7X.	60
4.17	Comparison of hybrid-PWM scheme delay distribution with self-calibration. Self-calibration reduces σ/μ of 21 chips by 2.5X.	61
5.1	Interconnect with transmission line behavior.	64

5.2	Self-timed regenerator circuit. Optimal sizing(unit: μm) for power reduction when 5 STRs are placed for a $0.45\mu m$ wire is shown.	66
5.3	Timing diagrams of STR at rising and falling transition. Speedup due to 2 low V_t transistors is shown.	67
5.4	Structure of global interconnect.	68
5.5	Repeater/STR implementation scheme.	69
5.6	STR and repeater simulation waveforms of $0.3\mu m$ wide interconnect.	70
5.7	STR simulation waveform of $4\mu m$ wide interconnect.	71
5.8	(a) Delay comparison with different numbers of STRs and repeaters (width : $1\mu m$). Sizing is optimized for each different number of STR and repeater. (b) Energy vs. Delay of STR and repeater (width : $0.45\mu m$).	72
5.9	Leakage power with different V_t assignments (width : $0.3\mu m$).	73
5.10	Spine clock distribution network configuration ($W1/W2=2$).	75
6.1	Receiver input signal waveform of [85, 86] and proposed scheme. Fast RZ signaling leads to bandwidth improvement.	79
6.2	Eye height and latency comparison of previous NRZ scheme and proposed RZ scheme.	80
6.3	Distributed interconnect model and 2nd order approximation of the transfer function.	82
6.4	Far-end signal for various pulse shapes.	83
6.5	For a given signal at the far-end, desired signal at the near-end is found.	84
6.6	Proposed transmitter with receiver circuits with waveforms when '001100' patterns is sent over on-chip links. Note that only 01 (rising) or 10 (falling) patterns generate pulses on the wire. Transceiver remains idle with consecutive 0s or 1s.	85
6.7	Further bandwidth improvement using double-data rate(DDR) scheme.	88
6.8	Communication system comparison employing single-data rate (SDR) and double-data rate(DDR).	89
6.9	Simulated waveforms of intermediate nodes in the DDR communication system.	90
6.10	Overall block diagrams of four communication schemes: (a) conventional full-swing repeater scheme (b) single series capacitor scheme [85, 86] (c) proposed SDR scheme (d) proposed DDR scheme.	91
6.11	Measured energy and performance of conventional, previous [85, 86], and proposed scheme.	92
6.12	Energy versus data activity of the proposed work.	93
6.13	Measured waveforms of trasmitter output and receiver input signals.	94
6.14	Bandwidth density and energy per bit comparison between proposed work and literature.	95
6.15	Chip micrograph.	95
7.1	Indiscriminate technology mapping may produce longer wires, adversely affecting delay and routing congestion.	97

7.2	Three schemes for comparison of single paths (a) Logic block (16 3-input NANDs) driving an optimally repeated 5mm wire (b) 16 3-input NANDs are placed along the wire (c) 16 3-input NANDs are decomposed into 24 2-input NANDs and placed along the wire.	98
7.3	Energy versus delay comparison for the three different schemes in Figure 7.2. . . .	100
7.4	Delay breakdown (logic delay, repeater delay, and wire delay) of the three schemes in Figure 7.2 at iso-energy of 1.1pJ.	101
7.5	Flow chart for the methodology of ‘Original’ and ‘No Large Cells’.	102
7.6	Critical path delay comparison of IWLS benchmarks using ‘Original’ and ‘No Large Cells’ approach in 130nm, 90nm, 65nm, and 45nm technology.	103
7.7	Critical path delay breakdown (gate-dependent delay and wire-dependent delay) of benchmark wb_conmax for (1) ‘Original’ and (2) ‘No Large Cells’ approach across four technology nodes.	103
7.8	Critical path comparison between ‘Original’ and ‘No Large Cells’ configuration for benchmarks (a) wb_dma at 65nm technology node and (b) systemcaes at 45nm technology node is shown (dots with circles represent inserted buffers).	105

LIST OF TABLES

Table

2.1	Flop distance and total wire length settings for conventional and ZL edge encoding scheme.	22
2.2	Multi-cycle interconnect results (peak energy, leakage and area) for the ZL edge encoding scheme. Results for both master-slave flip-flops (MS-FF) and time-borrowing pulsed flip-flops (TB-FF) are shown.	23
2.3	Performance, energy, and latency comparison for identical repeater sizes in OCL edge encoding scheme.	25
3.1	Average (2~8mm) alternating repeater bus advantages in 65nm technology.	38
3.2	Comparison of multi-cycle 64b 10mm bus in 65nm technology.	39
4.1	Leakage power measurement (units: μW).	57
5.1	STR power and performance comparison	68
5.2	Area and peak current comparison (iso-delay)	73
5.3	Repeater and STR leakage comparison	74
5.4	Clock distribution network comparison ($W1/W2=2$)	74
7.1	Detailed comparison of the benchmarks for ‘Original’ and ‘No Large Cells’ scheme on critical path, average wire length (=total routed wire length/wire count), inserted buffer count, total standard cell count, wire capacitance, and total standard cell area is shown for (a) 65nm and (b) 45nm technology.	104
7.2	Dynamic and leakage power comparison between ‘Original’ and ‘No Large Cells’ scheme for (a) 65nm and (b) 45nm technology.	107

CHAPTER 1

Introduction

Technology scaling has been the main driving force for the entire semiconductor industry for the past several decades by enabling higher levels of integration and improved chip performance. As depicted in Figure 1.1, however, technology scaling sharply reduces transistor delays along with local wire delays, while global wire delays quadratically increase [1].

Due to this ever growing gap, on-chip global interconnects are becoming an increasingly serious concern in current microprocessor designs in terms of latency, bandwidth, and energy. Repeaters have been a simple yet effective method to improve the latency of global wires as shown in Figure 1.1, but the number of repeaters increases dramatically with each technology node. Moreover, continuous technology scaling results in less wiring pitch with higher coupling capacitance and crosstalk noise, which directly impacts maximum clock frequency and chip power consumption of the conventionally repeated interconnects.

Especially, the increased integration of multiple-cores and large shared caches in microprocessors requires improved energy efficiency for core-to-core, core-to-cache and even intra-core communication to sustain the required performance benefits within shrinking power envelopes. Improvements in performance and power of on-chip busses can be achieved through process improvements, circuit design and novel communication architectures. For multi-core chips, improved bus designs are required to satisfy the constraints of robust operation and performance/energy gains across process corners and design space. Furthermore, the requirement of high clock frequency leads to careful consideration of inductance of the lines, dispersion, and other transmission line effects.

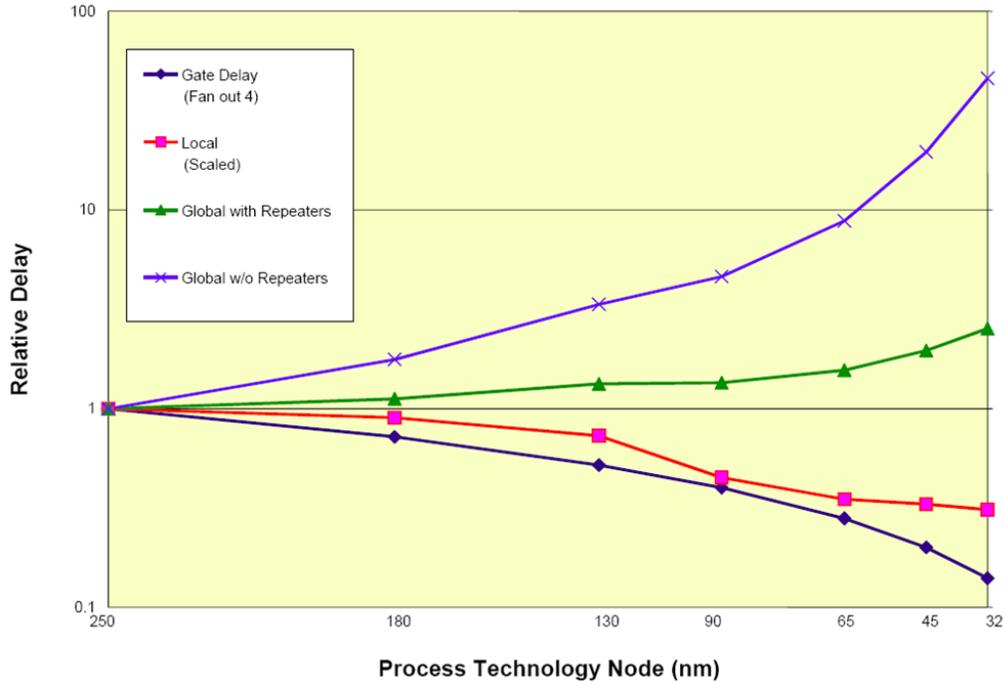


Figure 1.1: ITRS Trends of transistor delay, local interconnect delay, and global interconnect delay with technology scaling.

1.1 Implications of Scaling on Wire Resistance, Capacitance, and Inductance

First, we break down the effect of scaling into different wire parasitics, namely resistance, capacitance, and inductance. Figure 1.2 shows the scaling trends of unit-length wire resistance, ground capacitance, and lateral coupling capacitance for a semi-global metal layer with minimum width and spacing at 90nm, 65nm, and 45nm technology nodes. Despite the use of copper wires, it is observed that the unit-length wire resistance increases by $\sim 2.5X$ for each technology step, which forces excessive repeater insertion for long wires. The ground capacitance between vertical metal layers decreases due to the shrink of geometries such as wire width and height. On the other hand, however, since wire spacing reduces accordingly, the coupling capacitance between adjacent wires increases. As the overall result, the coupling capacitance is pronounced much more among the total wire capacitance as technology scaling continues. The increase in coupling from adjacent wires also informs us that the noise level from aggressor to victim is being boosted, thus posing a limit on the repeater distance or the wire width and spacing for signal integrity in long wires.

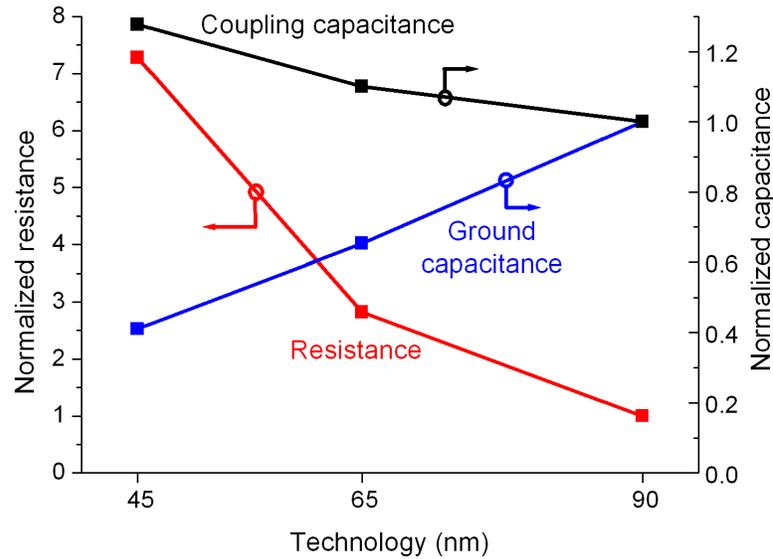


Figure 1.2: Scaling trends of unit-length wire resistance, ground capacitance, and lateral coupling capacitance for a semi-global metal layer with minimum width and spacing at advanced technology nodes.

All wires have inductance, but in most thin wires, the line resistance dominates the inductance, so the inductance hardly has an effect on the wire performance. Inductance has larger effect on wires which approach transmission lines, and a return path always has to be well defined in such interconnect designs. On the other hand, however, the highly increased number of wires in multi-bit busses often inhibits us from utilizing wide wires such as transmission lines due to routing area limitation.

In most of work where close to minimum pitch wires are considered, RC effects dominate any transmission line effects and we safely ignore inductance for them. However, in one of our work which we also target wide wires such as those in clock networks, we consider inductance effects as well.

1.2 Performance and Energy of Wires

Understanding the trends of each parasitic component in a wire from the previous section, deriving the performance and energy consumption of wires becomes more meaningful. First, the

delay of a gate driving a distributed RC wire is shown in the following simple equation. This formulation does not include the inductance wires, but wires with small pitches, which most of our proposed approaches are based on, are dominated by RC effects.

$$Delay \propto R_{gate}(C_{wire} + C_{load}) + R_{wire}\left(\frac{1}{2}C_{wire} + C_{load}\right) \quad (1.1)$$

One can use wide wires to reduce the resistance, but the delay improvement comes with the routing area penalty, and the designer has to note that the bandwidth density per unit area may not be favorable. The delay of a lossless transmission line where the inductance dominates the resistance is proportional to \sqrt{LC} , and one of our proposed approaches in Chapter 5 considers this in the bus system delay.

$$Energy = \alpha(C_{wire} + C_{load})V_{swing}V_{DD} \quad (1.2)$$

Wire energy consumption can be expressed as the above equation, where α is the data activity, and V_{swing} is the voltage swing of the wire. To improve the energy efficiency, we have to reduce either wire capacitance, wire swing, or the data activity factor on the wire. There have been approaches to deal with reducing the activity factor [4,5], but often it introduces complicated encoding circuits, so this work will not focus on it. In this work, we will explore techniques which can effectively decrease the capacitance or the voltage swing while improving or maintaining performance.

$$C_{wire} = C_g + 2 \times MCF \times C_c \quad (1.3)$$

Note that the total capacitance of wires comes from the given relation in Equation 1.3 between ground capacitance (C_g) and coupling capacitance (C_c) between adjacent wires. MCF stands for Miller coupling factor, which represents the switching possibility of adjacent wires. For example, worst-case MCF of 2 occurs when adjacent wires switch in the opposite direction, MCF of 1 is the case of shielded wires, and best-case MCF of 0 occurs when adjacent wires switch in the same direction. Continuous scaling reduces the width of on-chip wires and thus the ground capacitance of a minimum length wire for each technology. On the other hand, the ever-increasing level of integration produces densely packed wires for both intra-module communication and inter-

module communication, making coupling capacitance dominate over ground capacitance. Combining these basic equations with the 45nm node data shown in Figure 1.2 reveals that the coupling capacitance covers 79% of the total wire capacitance in case of MCF=1 and 88% of the total wire capacitance in case of MCF=2.

Therefore, an effective way to attack both the performance and energy of on-chip interconnect is to bring down the wire capacitance through reducing coupling capacitance between adjacent lines [15, 17, 36]. A number of new circuit techniques to enable this have been explored in this work. We want to clarify that our work did not consider any of bus encoding schemes which involves increasing the number of bits for data transmission [20, 21].

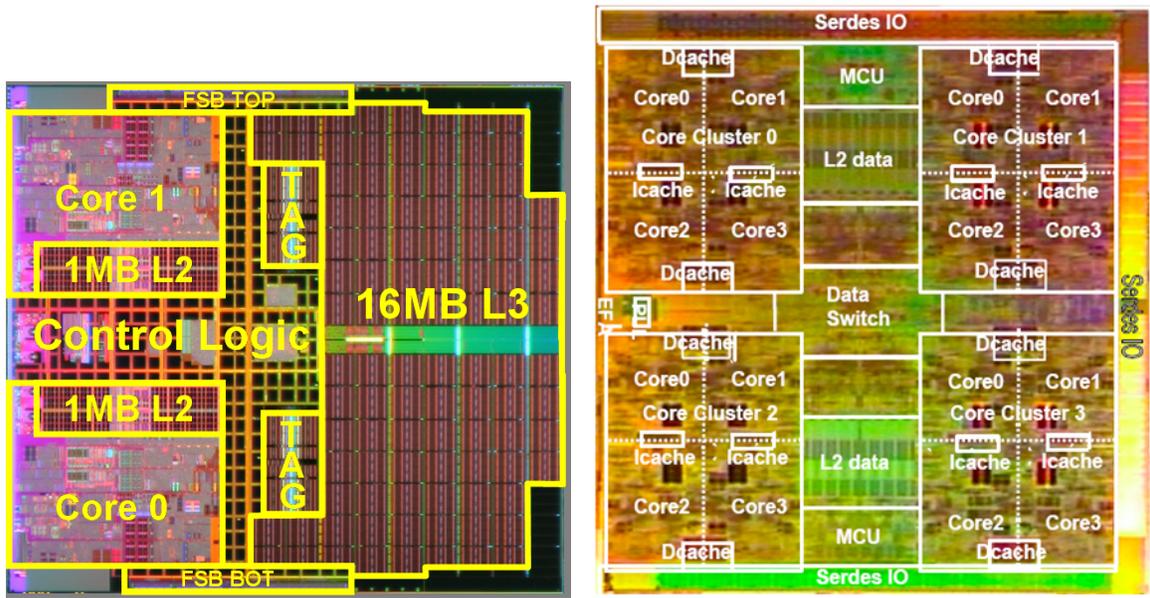
As seen in Equation 1.3, decreasing the swing of the wire is also preferable to reducing energy consumption linearly. Low-swing signaling has been pursued previously [55, 82, 85, 86] for aggressive energy reduction while sacrificing noise margins, but this often resulted in performance reduction. Chapter 6 proposed new circuit techniques to improve the latency and bandwidth of on-chip wires while maintaining high energy-efficiency over conventional schemes.

Previous literature and related work will be discussed in the according chapters to follow, where we describe the shortcomings and effectiveness of each approach that motivated this work and the resulted new circuit techniques to address or further improve those.

1.3 Clock consideration for global interconnects

As seen in the previous section, long on-chip wires do not scale. Although technology scaling continues, since the die size stays the same or even increases, the length of global wires for inter-module communication increases as well. Figure 1.3 shows two recent multi-core microprocessor designs where one side of each chip is as long as 20mm due to high level of integration, and this poses serious concerns for on-chip communication. On the other hand, clock frequency kept on increasing for better performance aided by the reduction in intrinsic gate delay. These trends naturally leads us to multi-cycle interconnect schemes.

In order to prevent spending a large number of cycles for transmitting data from one end of a chip to the other end of a chip, better circuit techniques are required for reduce the number of clock cycles dedicated purely for on-chip communication. Another obvious concern is energy



(a) Intel 65nm dual-core microprocessor [2].

(b) Sun 65nm 16-core microprocessor [3].

Figure 1.3: Die micrograph of two recent multi-core microprocessors.

constraint on the overall intra-chip communication, which cannot be easily satisfied with conventional repeater schemes. Therefore, high-speed and energy-efficient schemes for improved on-chip communication are required for recent and future microprocessor designs.

Throughout this work, we assumed identical synchronous clock in the source and destination of a point-to-point interconnect, following the clock designs in a number of recent microprocessors [6–8]. It is true that clock gating is heavily used nowadays to reduce power of idle components, and different cores could have different clock domains, but we focus on long wires in a single clock domain at a time without losing generality. Some previous works [85, 86, 88] in high-speed interconnects do not necessarily assume this, and require arbitrary transmitter clocks and receiver clocks to achieve higher data rates. This is one of the reasons that those types of designs cannot easily penetrate into industry. Therefore, throughout our work, synchronous clocking scheme is assumed, where the same clock arrives at both the input flip-flop and the output flip-flop, and the long interconnect structure between the two flip-flops will be optimized either with repeaters or without repeaters. The possible skew between the input and output clock would be reserved as the margin while performing the interconnect optimization.

1.4 Contribution of This Work and Organization

This work proposes a number of new circuit techniques for global on-chip communication to improve the energy, performance, and robustness over process variations. A number of related works exist, but our goal is pursue those techniques which are more readily applicable in practical microprocessor or ASIC designs. Specifically, we do not consider previous approaches such as bus encoding work which actually increases the required number of wires, techniques which need more than one supply voltage, and current sensing approaches which constantly dissipates static current.

The proposed circuit techniques in this work can be divided into two categories: (1) techniques to reduce the effective wire capacitance from Chapter 2, 3, and 4 which build upon repeater insertion and (2) techniques from Chapter 5 and 6 which enable and facilitate repeater-less signaling. These are followed by a CAD framework studying the effect of long wires on technology mapping.

Chapter 2, 3, and 4 pursues circuit techniques which build upon using regular repeaters to break down the long wire, reducing the effective wire capacitance by attacking the MCF between adjacent wires. A novel 'Edge Encoding' technique is introduced in Chapter 2, which effectively eliminates the worst-case MCF of 2 while maintaining the best-case MCF of 0. This is performed by desynchronizing the rising and falling transitions, achieving both average and peak energy reduction.

In Chapter 3, we propose another circuit design technique to reduce the worst-case MCF over the length of the bus, hence improving the worst-case delay and energy for on-chip static busses. This technique is applicable for drop-in replacement for minimal change in design methodology and also possesses the ability to improve power-performance of shared busses with multiple driver and receiver points along the bus.

Although Chapter 2 and 3 reduced the effective coupling capacitance on long wires, the wire resistance was identical to the conventional scheme since the wire geometry and the number of wires are same. Chapter 4 proposes a signaling technique to send two bits of information are sent on one wire, thereby halving the number of wires in wide multi-bit busses. This leads to wider wiring pitch for the same footprint with the conventional scheme, resulting in less wire resistance, less coupling capacitance, less delay, and less energy consumption. The proposed technique is

based on pulse width modulation (PWM) and exploits the controllability of pulse widths in pulsed signaling to improve both performance and energy consumption in global wires.

While Chapter 2, 3, and 4 seeked techniques to improve global wires in the presence of repeaters, it is true that the number of repeaters are skyrocketing with each technology step. Addressing aforementioned interconnect issues without using repeaters would be certainly beneficial. Chapter 5 and 6 introduces on-chip circuits which enable and facilitate repeater-less signaling. Self-timed regenerator circuits are discussed in Chapter 5, which are placed along repeater-less global wires and expedite the signal transition. The proposed circuits provide additional charge shortly when it detects a transition and then the self-timed part makes the regenerator ready for the next incoming transition. This circuit technique compensates the loss in resistive wires while consuming minimal power.

To reduce the wire energy consumption further, limiting the voltage swing of the long interconnect is considered. In Chapter 6, a new low-swing signaling technique is proposed for high bandwidth and low energy consumption. The transmitter generates pre-emphasized bipolar signals through series capacitors and the receiver efficiently recovers NRZ data from fast RZ pulses. Employing double data rate (DDR) signaling can further improve the data rate of the on-chip bus system.

While the previous chapters proposed new circuits for better performance/energy of point-to-point interconnects, Chapter 7 explores a methodology to improve the overall performance of general ASIC designs while considering the interaction between long wires and technology mapping. We point out that long wires are often generated by the use of large standard cells, resulting in excessive buffer insertion. As technology scales, wire delay increasingly dominates the critical path delay, and using simpler cells leads to less buffer insertion, shorter wires, and overall better performance.

In Chapter 8, the proposed works in the thesis is summarized and various aspects of future work are discussed.

CHAPTER 2

Edge Encoding

In this chapter, we propose a new circuit technique for on-chip communication, the edge encoding technique, to reduce the energy consumption in multi-cycle interconnects. Both average and worst-case energy are reduced by desynchronizing the edges of rising and falling transitions. In a 1.2V 65nm CMOS technology, the proposed approach achieves up to 34% energy reduction with no latency overhead over optimally designed conventional busses due to coupling capacitance reductions. The technique further reduces energy consumption by 39% with iso-throughput of the conventional scheme at the expense of one-cycle latency. Energy savings are shown to be both larger and more robust to process, voltage, and temperature variations than previous techniques.

2.1 Motivation and Previous Work

Due to higher integration of multiple cores in current microprocessors, the number of wires used for inter-module communication has skyrocketed [9]. Furthermore, the increased complexity and high level of integration requires higher wire densities, and coupling capacitance has dominated total wire capacitance for several technologies already. A high coupling capacitance ratio is not favorable in conventional busses due to the possibility of adjacent wires switching in the opposite direction, yielding a worst-case Miller capacitance factor (MCF) of 2. For example, when $MCF=2$ the coupling capacitance ratio over the total interconnect capacitance is over 80% for a minimum pitch intermediate metal layer in 65nm [16]. It is possible to reduce coupling capacitance by increasing spacing or introducing shielding, but this comes at the cost of significant area

penalties [12]. Hence, a key challenge in interconnect design is to reduce the worst-case MCF while maintaining the same physical footprint of the interconnect, thereby reducing the effective wire capacitance and interconnect energy consumption.

There have been several attempts to reduce worst-case MCF to 1 for delay improvement and power reduction. In [36], the authors introduced a delay element on alternating wires, thereby avoiding the MCF=2 switching case through temporal separation. In this approach, however, fine-tuning of the optimal insertion delay is non-trivial, and due to very small inverter delays in sub-90nm technologies, many inverters are needed to sufficiently separate the switching of adjacent wires, increasing power. Also, this technique is sensitive to process variation since variability in the inserted delay can lead to a lack of sufficient separation for adjacent wires.

Separating the timing of transitions on adjacent wires was also proposed in [15] by assigning different clocks to flip-flops driving adjacent wires. Rather than assigning clocks with different phases, [16] implemented a technique that alternatively used positive-edge triggered and negative-edge triggered flops in every other wire. In this case, however, the wire length associated with the final flop must be short to align to the positive edge at the far end of the wire. In [16], the authors proposed a method to skew alternating wires in the opposite direction using different width, length, V_t and body bias. In this way the worst-case switching is separated without hurting the best-case switching. However, this technique is also very sensitive to process variations, which can lead to less separation than needed to achieve an MCF of nearly 1. A method using careful staggering of repeater locations is introduced in [17]. This method results in alternating MCF=0 and MCF=2 in neighbor wire segments. However, in terms of physical design, this method incurs significant overhead considering that the repeater location cannot always be arbitrarily selected in industrial designs. Without modifying the repeater locations, techniques to use both inverting repeaters and non-inverting repeaters were proposed in [18, 19], such that one half of the wire segments experience MCF=0 and the other half experience MCF=2.

A number of coding techniques including [20, 21] encode the conventional bus such that adjacent bits never switch in the opposite direction. However, in addition to the special encoder and decoder circuit overhead, these techniques require additional wires for bus encoding which increase routing area. At the same footprint of using additional wires, the conventional bus could use extra spacing without any encoding (no additional wires) to improve speed and energy con-

sumption, and it is not clear which will perform better.

Pulsed bus techniques [22] also achieve a worst-case MCF of 1. In these pulsed bus techniques, however, the energy dissipation is increased per transition compared to conventional busses due to the pulse encoding. Techniques in [23] reduced this overhead by selectively using low V_{dd} with nominal V_{dd} to drive the interconnect, but this is done at the expense of design complexity since two power supplies are required.

References [15, 17–19, 36] reduce the overall worst-case MCF of an interconnect to 1, but also eliminate the best-case MCF of 0 (all adjacent wires switching in the same direction), leading to less advantage in average energy consumption. Using the technique in [16], best-case switching is maintained, but at the expense of smaller noise margin in the repeaters and more sensitivity to process variations as mentioned above. Furthermore, the amount of skewing required to effectively separate transitions of adjacent wires is heavily dependent on technology.

This chapter presents a new encoding technique [10, 11] that achieves a worst-case MCF of 1, while preserving the best-case MCF of 0. This is done by controlling the edges of rising and falling transition in time, namely always performing rising transitions on the negative edge of the clock and falling transition on the positive edge of the clock (or vice versa). Since the worst-case switching is separated by as much as one phase (half clock cycle), this technique remains robust against process variation. Hence, both the average and worst-case energy can be reduced without impacting the sensitivity to process variation. Average energy savings will aid battery life and typical energy costs, but worst-case energy is also a meaningful metric in terms of thermal management and peak demand for power grids and decoupling capacitance [24]. These savings are accomplished at the expense of minimal encoder logic with half cycle latency and additional clocking. However, we find that the logic and clocking overhead is small in long interconnects where interconnect power consumption is dominant, and also show that the potential latency overhead can be eliminated or minimized in multi-cycle interconnects. This chapter includes fundamental theoretical analysis, extensive results investigating different types of flip-flops for the proposed technique, and in-depth process, voltage, and temperature (PVT) variation experiments.

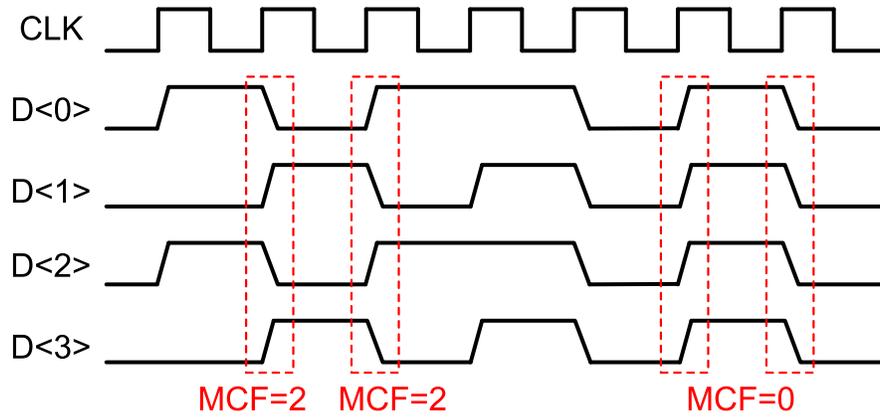
2.2 Edge Encoding Approach

In this section, we first describe the basic concept of the proposed encoding of eliminating worst-case MCF of 2 while maintaining best-case MCF of 0. Theoretical energy savings are shown using analytical models, and the circuitry and operation of two new encoding techniques for multi-cycle interconnects will be explained in detail.

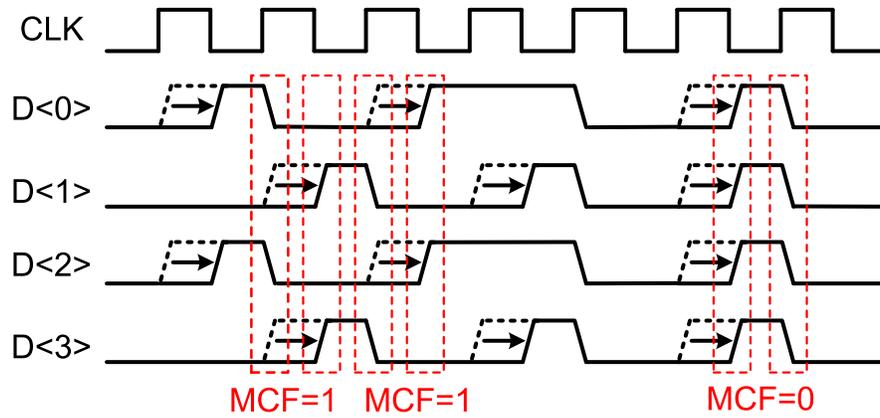
2.2.1 Basic Idea

In a multi-cycle bus structure, the transitions between neighboring wires are synchronized at every flip-flop as the signal propagates down the bus. This often generates simultaneous switching of adjacent wires in the opposite or same direction. In Figure 2.1(a), the worst-case (MCF=2) and best-case (MCF=0) switching of the conventional bus are shown. The MCF=2 case, where every other wire switches in the opposite direction, generates the worst-case delay, which defines the clock frequency and also consumes the worst-case energy due to maximum coupling capacitance. To avoid this, we propose to selectively shift rising and falling edges and separate them by as much as half cycle. For example, as seen in Figure 6.1(b), if we selectively delay only the rising transitions by a half cycle and keep the falling transitions unaltered, the worst-case MCF is reduced from 2 to 1. We refer to this selective edge shifting as *edge encoding*. Since edge encoding shifts the same transitions together, the advantage of best-case switching (MCF=0) is still maintained, which is unachievable in most other approaches [15, 17–19, 36].

Since the edge-encoded signal transitions at both positive and negative edges of the clock, we use dual-edge triggered flip-flops to propagate the signal along long multi-cycle interconnects. Since the signals must be synchronized back to positive-edge triggered flip-flops after long interconnects, the number of dual-edge flip-flops within the multi-cycle interconnect should be even. The methodology for the placement of dual-edge flip-flops in the edge encoding technique to maximize energy-efficiency will be described in Section 2.2.3.



(a) Conventional wire switching.



(b) Proposed wire switching.

Figure 2.1: Conventional and proposed wire switching scenario in adjacent wires.

2.2.2 Theoretical Energy Savings

The total interconnect capacitance is the sum of ground capacitance (C_g) and coupling capacitance (C_c). The effective interwire coupling capacitance depends on the switching behavior of adjacent wires, which is characterized by MCF in Equation 2.1 below. MCF is 0 when all adjacent wires switch in the same direction where the total wire capacitance is only C_g , and MCF is 2 when every alternating wire switches in the opposite direction resulting in total capacitance of $C_g + 4C_c$. Note that MCF is an approximate value since transitions in adjacent wires can occur arbitrarily. Actually, [25] reports the true worst case MCF of 3 if the slew rate of the aggressor is twice as fast as that of the victim, but MCF of 2 is used as a rule of thumb for worst-case switching here to compute theoretical energy savings. In reporting results later in Section 7.3, we use SPICE to reflect the actual interwire coupling in multi-bit busses.

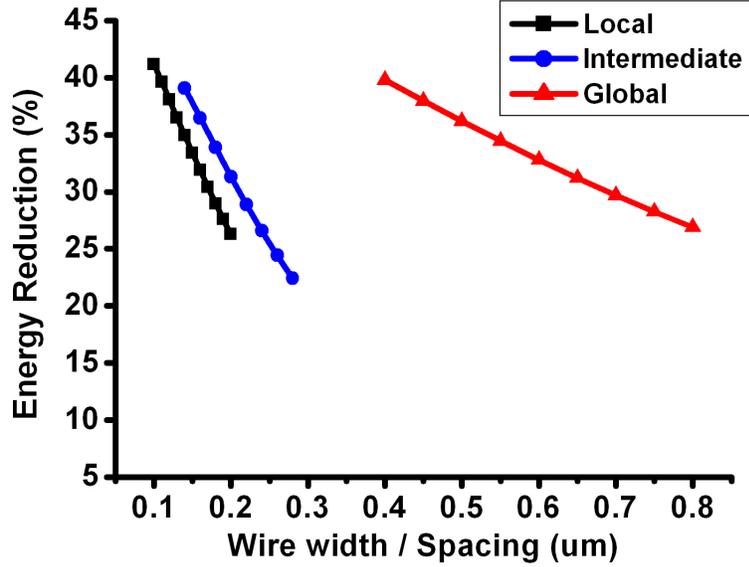


Figure 2.2: Ideal wire energy savings due to MCF reduction based on 65nm interconnect dimensions [14].

If we can control the transitions as shown in Figure 6.1(b), the worst-case MCF is reduced to 1, and reduction of wire energy consumption is achievable, as expressed in Equation 2.2. The maximum wire energy savings we can ideally achieve is dependent on the ratio of ground capacitance and the coupling capacitance in the interconnect.

$$C_{wire} = C_g + 2 \times MCF \times C_c \quad (2.1)$$

$$E_{wire,saving} = 1 - \frac{C_{wire(MCF=1)} V_{dd}^2}{C_{wire(MCF=2)} V_{dd}^2} = 1 - \frac{C_g + 2C_c}{C_g + 4C_c} \quad (2.2)$$

Closed-form equations from [13] compute capacitance values for a given wire geometry, namely wire width, spacing, thickness and dielectric thickness. With typical wire dimensions given in [14] for local, intermediate and global wires in the 65nm technology node, the expected energy savings are calculated using Equation 2.2. A range of wire pitches are shown, with W=S being swept from minimum to double pitch. The ideal energy savings are shown in Figure 2.2.

Besides the wire energy, the energy dissipation due to the capacitance of repeaters should be also included in the total energy consumption of optimally repeated interconnects, as shown in

Equation 2.3. C_{tr} is the sum of gate and drain capacitance of a unit-sized repeater. The total capacitance of repeaters is proportional to the number of repeaters inserted (N_R) and the size of the repeaters (H_R), where these two parameters heavily depend on the resistance and capacitance of the interconnect.

As the pitch increases, the achievable energy savings due to MCF reduction decreases as expected since the interwire coupling capacitance diminishes. In general, even for less favorable non-minimum pitches, Figure 2.2 shows that manipulation of MCF can lead to appreciable (25-40%) energy savings.

$$E_{total} = (C_{wire} + N_R H_R C_{tr}) V_{dd}^2 \quad (2.3)$$

Equations 2.4 and 2.5 from [26] show the interaction between the repeater parameters and wire parasitics for energy-delay optimal repeater insertion. R_{tr} is the average transistor resistance unit-sized repeater, R_{wire} is the total interconnect resistance, and C_{wire} is the total interconnect capacitance from Equation 2.1.

$$N_R = \frac{1}{\sqrt{3}} \sqrt{\frac{0.4 R_{wire} C_{wire}}{R_{tr} C_{tr}}} \quad (2.4)$$

$$H_R = \frac{0.6}{\sqrt{3(0.4)(0.7)}} \sqrt{\frac{R_{tr} C_{wire}}{0.7 R_{wire} C_{tr}}} \quad (2.5)$$

Given a 40% reduction in C_{wire} (minimum pitch wires in Figure 2.2), both N_R and H_R are reduced by 23% ($1 - \sqrt{1 - 0.4} = 0.23$). In this simple energy model, peak MCF reduction decreases C_{wire} and $N_R \times H_R$ by the same ratio. Therefore, the total energy reduction (Equation 2.3) will be identical to the wire energy reduction from Equation 2.2, regardless of the absolute values of wire capacitance and repeater capacitance. The total energy savings including repeaters for local, intermediate, and global interconnects is equivalent to Figure 2.2.

In our proposed scheme, the aforementioned analytical energy reduction will be degraded by additional clock and encoder energy, however for long intermediate and global interconnects, this additional energy will be small compared to the total wire energy consumption. Detailed results will be shown in Section 7.3.

2.2.3 Edge Encoding Technique

As described in Section 2.2.1, the objective of the edge encoder is to selectively shift the rising and falling transition by different amounts. This encoding is done simply by performing an AND operation between the original signal and the half-cycle delayed version of itself. In this way, only the rising edge is delayed by a half cycle, separating simultaneous rising and falling transition by a half cycle. Since the encoder logic is very simple, the encoding overhead in terms of power and area is very small. This makes the edge encoding technique a highly practical approach.

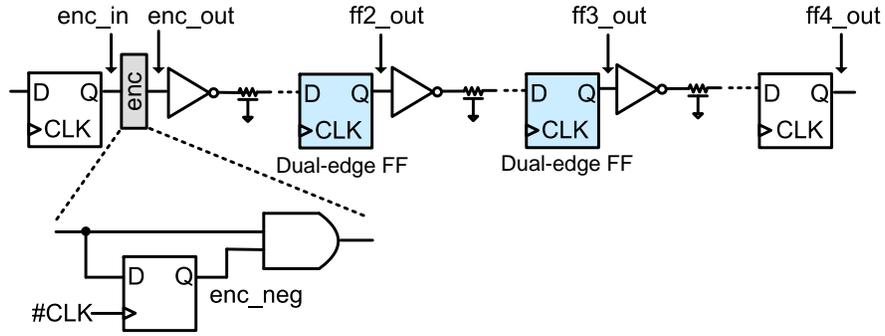
We propose two schemes to effectively use the edge encoding technique in multi-cycle interconnect. The two methods differ in the procedure to cope with the initial half cycle latency required for edge encoding and to address the issue of aligning back to the positive-edge triggered signal at the far end of the wire.

Zero Latency (ZL) Scheme

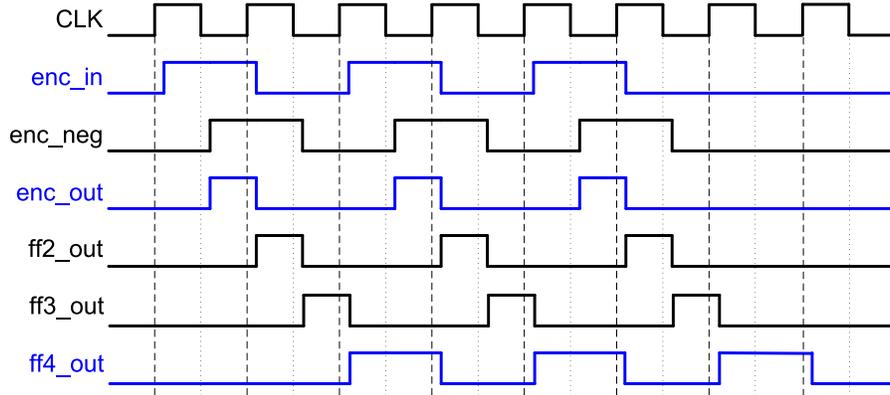
The zero latency (ZL) scheme reduces energy consumption in multi-cycle interconnects without any latency overhead although encoding requires a half-cycle delay at the near end of the wire. This scheme exploits the fact the signal propagation in the edge-encoded bus is faster than that in the conventional bus due to reduced coupling capacitance as described in Section 2.2.2.

The block diagram of a multi-cycle interconnect with simple encoder logic is shown in Figure 2.3(a). The encoding procedure and the propagation of the encoded signal are shown in Figure 2.3(b). When data toggles every cycle, the encoder generates a half-cycle pulse (enc_out). As this half-cycle pulse propagates through an even number of dual-edge flip-flops, it automatically aligns back to a positive-edge triggered signal (ff4_out) at the far end. Therefore, there is no need for any decoder circuit.

To achieve overall zero-latency, the interconnect system is set up as shown in Figure 2.4. If the conventional scheme requires n cycles to propagate through the entire interconnect, the edge-encoded bus must propagate through in $(2n - 1)$ half cycles, considering that the encoding takes one half cycle to synchronize at the far end of wire. In Figure 2.4, L1 is the distance between positive-edge triggered flip-flops in the conventional bus, and L2 is the distance between dual-edge triggered flip-flops in the edge-encoded bus. If L2 is defined by L1 and n in Equation 2.6,



(a) Encoder logic and block diagram of ZL scheme.



(b) Timing diagram of ZL scheme.

Figure 2.3: Block and timing diagrams of ZL scheme.

overall zero latency is achievable. For example, in a 9mm interconnect, when $n=3$ and $L1=3$, the edge-encoded signal will propagate 1.8mm every half cycle while the conventional signal will propagate 3mm every cycle. Effectively, the edge-encoded signal is traveling 20% longer (1.8mm vs. 1.5mm) during the same time period, which is possible when at least a 17% ($1 - 1/1.2$) speedup is achieved in the edge-encoded bus due to coupling capacitance reduction.

$$L2 = L1 \times \frac{n}{2n - 1} \tag{2.6}$$

One Cycle Latency (OCL) Scheme

In multi-cycle interconnects, multiple cycles are required to propagate across the entire wire. In these cases, one additional cycle latency may be acceptable if a clock frequency increase or aggressive energy reduction is a higher design priority. The one cycle latency (OCL) scheme is

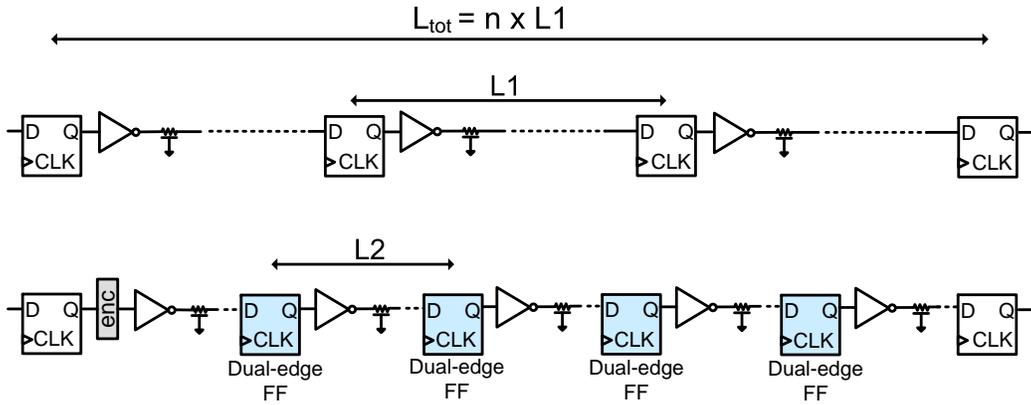


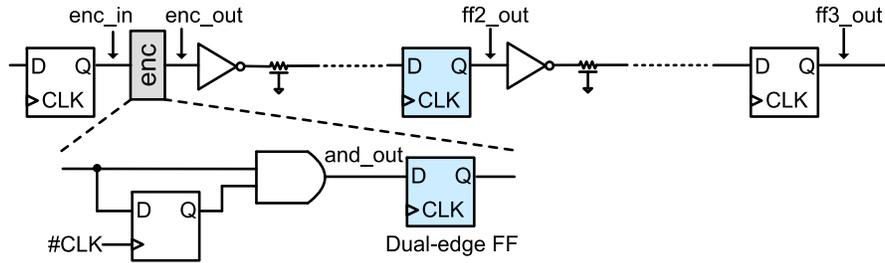
Figure 2.4: Flip-flop placement in conventional and ZL edge encoding scheme.

intended to achieve further performance improvement and energy reduction for a fixed throughput at the expense of one-cycle latency. After the encoding, the data must eventually align to the positive edge of the clock at the far end of the wire. To achieve this, we can align the transition at the near end to the positive edge of the clock by encoding with a full one cycle delay, and then allow for normal signal propagation along the wire. The one-cycle latency is therefore introduced once at the beginning of the wire and the throughput is not hampered.

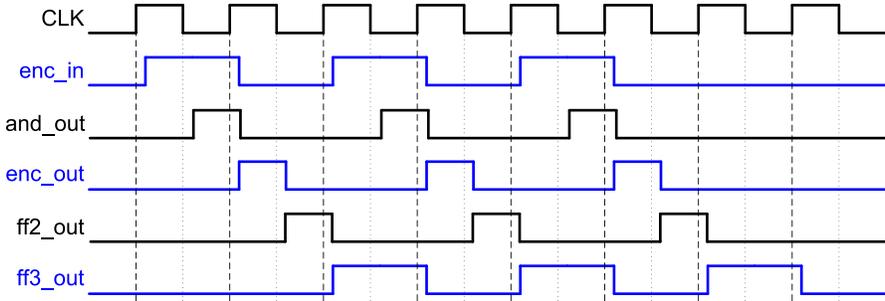
The block and timing diagrams of the OCL edge encoding scheme are shown in Figure 2.5. The difference in the encoder in Figure 2.5(a) compared to ZL is that a dual-edge flip-flop is added at the output to intentionally delay enc_in by one cycle and align the rising edge of enc_out at the positive edge of the clock as shown in Figure 2.5(b). The corresponding flip-flop placement in the OCL edge encoding scheme is shown in Figure 2.6. Dual-edge flip-flops are placed at intervals equal to half the flop distance of the conventional bus. In the OCL edge-encoded bus, since the worst-case wire delay is reduced due to MCF reduction, we can either increase the clock frequency for high-performance busses or downsize the repeaters for iso-performance to the conventional bus for aggressive energy reduction.

2.2.4 Dual-edge flip-flops and timing concerns in edge encoding

Since the edge encoding technique requires dual-edge flip-flops, the number of flip-flops placed in multi-cycle interconnects is inevitably increased compared to conventional multi-cycle interconnects with single-edge flip-flops. First, due to shorter distance between dual-edge flip-flops, the



(a) Encoder logic and block diagram of OCL scheme.



(b) Timing diagram of OCL scheme.

Figure 2.5: Block and timing diagrams of OCL scheme.

slew rate constraint in the edge encoding scheme should be more stringent. Both in the conventional and edge encoding schemes, the repeaters between flip-flops are sized such that 10%-90% slew rate is $\sim 10\%$ of the respective signal propagation delay between the flip-flops.

In addition, along the critical path of multi-cycle interconnects, both the total setup time and CLK-Q delay in flip-flops increase as well. Hold time is not a concern even when using dual-edge flip-flops because the interconnect paths are well-defined with several repeaters and large wire load, and thus short paths between flip-flops do not exist.

It is well known that time-borrowing flip-flops have zero or negative setup time, providing performance benefits over master-slave flip-flops [27, 28]. Also, particularly for multi-cycle interconnects, time-borrowing single-edge flip-flops were suggested [29] for better tolerance against within-die variation and higher maximum frequency. Therefore, in the proposed edge encoding techniques, time-borrowing flip-flops are considered for the inserted flip-flops to mitigate the increase in total setup time and variation. However, the benefits of time-borrowing flip-flops come at the expense of additional transistors and higher energy consumption. In the proposed edge encoding techniques, since additional dual-edge flip-flops are placed in the multi-cycle intercon-

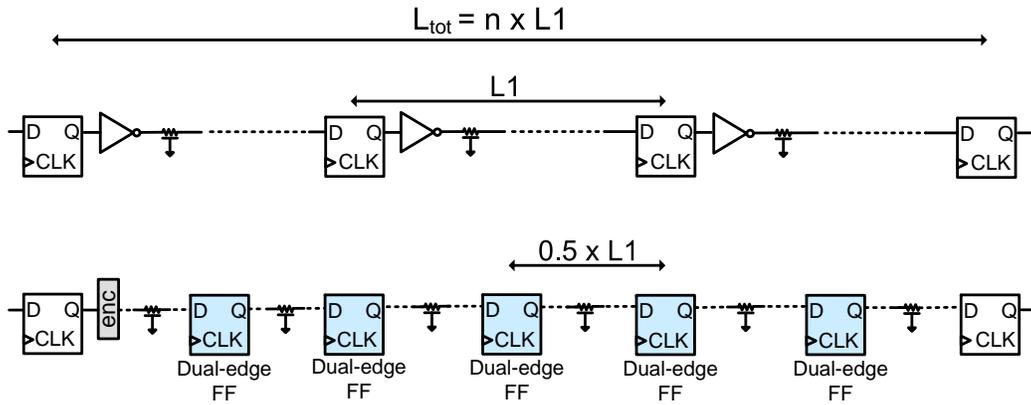


Figure 2.6: Flip-flop placement in conventional and OCL edge encoding scheme.

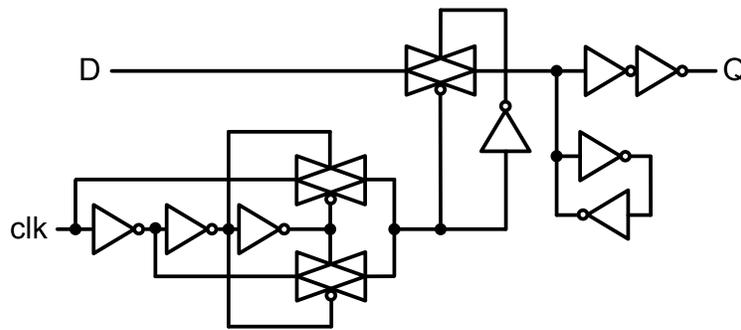


Figure 2.7: Schematic of time-borrowing pulsed dual-edge flip-flop.

next path, the energy-delay tradeoff of the time-borrowing dual-edge flip-flops has to be carefully investigated.

In our experiments, we performed each edge encoding technique (ZL and OCL) with both conventional master-slave dual-edge flip-flops and time-borrowing dual-edge flip-flops. The master-slave dual-edge flip-flops have positive setup time, and the additional setup time and D-Q delay will be considered in the overhead involved with the proposed techniques. For the time-borrowing dual-edge flip-flops, several types of flip-flops from [28] were considered to absorb the setup time and mismatch between interconnect paths. We selected the pulsed triggered dual-edge flip-flop in Figure 2.7 for the proposed edge encoding because it is more area-efficient and the D-Q path is shorter. Throughout this chapter, we designed the transparent window to be ~ 1.6 FO4 delay. Since the transparent window will be applied on both the rising and falling edge of a high frequency clock in our schemes, the size of the transparent window cannot be arbitrarily large. The results of using

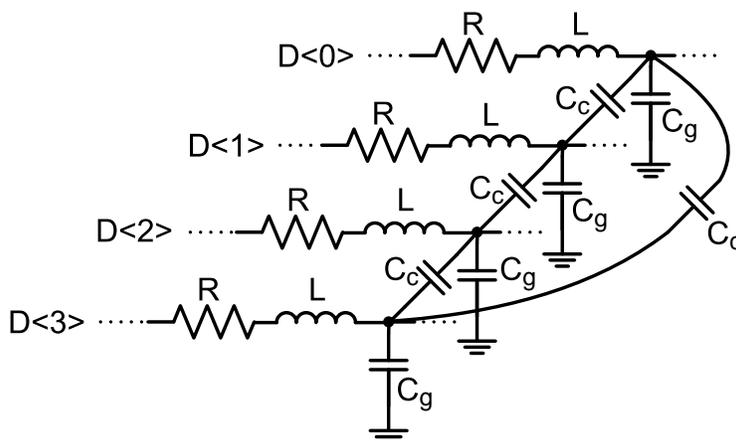


Figure 2.8: 4-bit cyclic bus model.

both master-slave flip-flops and time-borrowing pulsed flip-flops for the proposed techniques are shown in the following section.

2.3 Experimental Results

To accurately capture the effect of coupling capacitance in adjacent wires, we use the 4-bit RLC cyclic model [16] for the interconnect shown in Figure 2.8. Interconnect parasitic values are extracted for a minimum pitch intermediate layer metal 4 in 65nm technology, and the corresponding values are $R = 1400\Omega/\text{mm}$, $L = 0.6\text{nH}/\text{mm}$, $C_g = 64\text{fF}/\text{mm}$, and $C_c = 72\text{fF}/\text{mm}$. All experimental results are obtained from SPICE simulations with a 1.2V supply.

For various flop distances, the conventional repeater bus is optimized by sweeping both the number and sizes of repeaters. Energy, delay, clock frequency, and leakage power are measured for the optimally designed conventional busses, with this serving as the baseline for comparison with edge encoded busses. Unless mentioned otherwise, activity factor of 50% (data switches on every positive edge of clock) is assumed. We now show results for the two edge encoding schemes as proposed in Section 2.2.3.

Table 2.1: Flop distance and total wire length settings for conventional and ZL edge encoding scheme.

n (number of cycles)	L1	L2	Total wire length ($n \times L1$)
3	2mm	1.2mm	6mm
3	3mm	1.8mm	9mm
3	4mm	2.4mm	12mm
3	5mm	3mm	15mm

2.3.1 Zero Latency (ZL) Scheme

As described in Section 2.2.3, both the conventional and ZL edge encoding schemes operate at the same clock frequency, however the flop to flop distance in the ZL scheme is effectively larger. From Figure 2.4, L2 in the ZL scheme depends on L1 in the conventional scheme as defined by Equation 2.6. The optimized set of flop distances and interconnect lengths using the ZL scheme is summarized in Table 2.1. The number of cycles is set to 3 in all cases for simplicity. A flop distance of 1mm in a conventional bus was found to be too short for the edge encoding technique to gain enough speedup for the ZL scheme to be applicable, thereby 2-5mm are selected for L1. This gives a range of applicability for the proposed technique in this particular technology - note that more advanced processes should allow for benefits at even shorter wire lengths.

For each configuration in Table 2.1, we found the maximum clock frequency at which we compared the total energy consumption in the conventional and ZL edge encoding schemes. The resulting energy reduction obtained in the ZL scheme and the clock frequency achievable at each flop distance (L1) are shown in Figure 2.9 for both master-slave dual-edge flip-flops (MS-FF) and time-borrowing pulsed dual-edge flip-flops (TB-FF). Both peak energy and average energy are shown. For average energy, we generated random data over 100 cycles with activity factor of 25% for each of the 4-bit input. As L1 increases, more energy reduction can be achieved using edge encoding, while the maximum clock frequency degrades. Using time-borrowing flip-flops allows negative setup time in the dual-edge flip-flops and improves performance. In the case using time-borrowing flip-flops, the repeaters and flip-flops can be sized down for iso-performance at each flop distance (L1), which leads to additional energy reduction as shown in Figure 2.9.

A detailed comparison for a flop distance (L1) of 3mm is shown in Table 2.2. Both schemes operate at 2GHz, and we can see that considerable energy savings are achieved for various activity

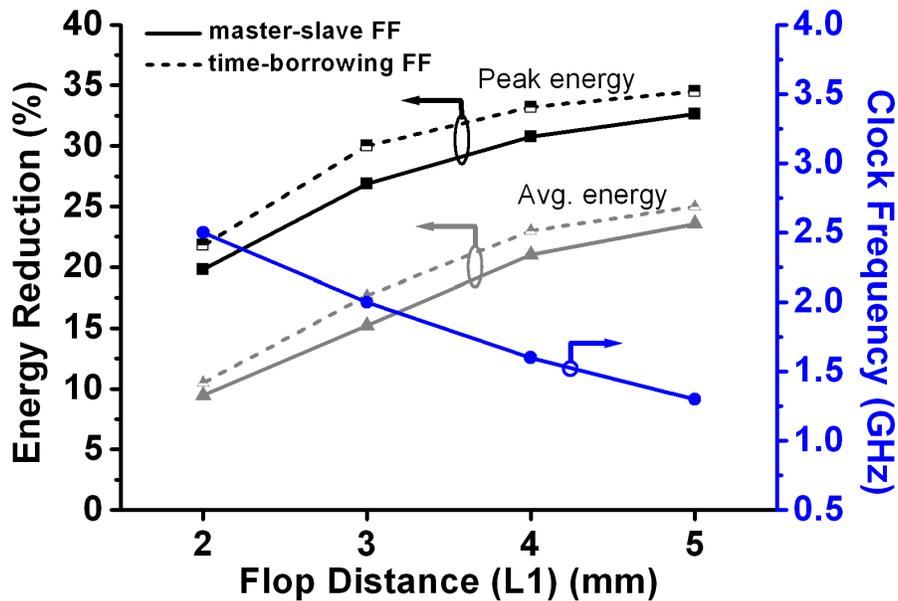


Figure 2.9: Comparison of the ZL edge encoding scheme to conventional busses in worst-case/average energy and clock frequency for flop distances (L1) of 2-5mm.

Table 2.2: Multi-cycle interconnect results (peak energy, leakage and area) for the ZL edge encoding scheme. Results for both master-slave flip-flops (MS-FF) and time-borrowing pulsed flip-flops (TB-FF) are shown.

Scheme	Frequency	Energy/cycle @25% activity	Energy/cycle @15% activity	Energy/cycle @10% activity	Leakage power	Total transistor width
Conventional	2GHz	1.83pJ	1.26pJ	0.77pJ	16.9 μ W	492.4 μ m
Proposed (ZL) MS-FF	2GHz	1.39pJ (-24.2%)	0.96pJ (-23.6%)	0.64pJ (-16.5%)	14.2 μ W (-16.2%)	424.7 μ m (-13.7%)
Proposed (ZL) TB-FF	2GHz	1.34pJ (-27.3%)	0.94pJ (-25.4%)	0.67pJ (-13.0%)	11.5 μ W (-26.0%)	375.1 μ m (-23.8%)

factors. The amount of energy saving decreases at lower activity factors, because the edge encoding scheme consumes additional clock energy and the portion of clock energy increases as the data activity rate is lowered (this may be ameliorated by clock gating or other similar techniques). Using time borrowing flip-flops accelerates this trend since time-borrowing flip-flops have additional transistors and higher clock energy consumption than normal master-slave flip-flops. Due to the reduction of effective capacitance on the wire fewer repeaters are required in the ZL scheme than the conventional scheme for optimal performance and energy, yielding less leakage power and total transistor width, as seen in Table 2.2.

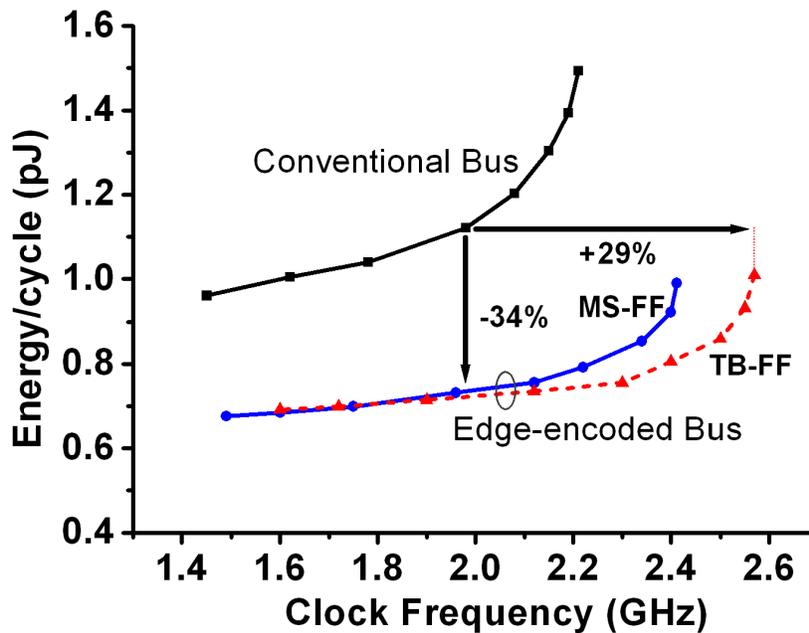


Figure 2.10: Energy-clock frequency comparison for a 3mm OCL edge-encoded bus.

2.3.2 One Cycle Latency (OCL) Scheme

As we saw in Section 2.2.3, the OCL edge encoding scheme can either reduce energy at iso-performance or improve the performance at iso-energy at the expense of one-cycle latency. To quantify the performance gain or energy reduction, we optimized both the conventional and OCL edge encoding schemes (MS-FF and TB-FF) for a minimum pitch 3mm wire. Figure 6.11 shows the energy per cycle versus clock frequency for each scheme. Edge encoding shows a potential 29% performance improvement at iso-energy or a 34% energy reduction at iso-performance (2GHz).

Figure 2.11 shows the breakdown of energy for a 5mm wire in both an optimally-designed conventional repeater bus and OCL edge-encoded bus at iso-throughput. The wire energy, which is the dominant source of total energy consumption, is reduced considerably using edge encoding at the expense of minimal encoder logic and additional clocking energy. Overall, the OCL approach can achieve 36% energy reduction with MS-FF and 39% energy reduction with TB-FF. Similar to the ZL scheme, larger energy reductions are achieved with TB-FF due to smaller repeaters and flip-flops at a fixed performance.

Also, in the OCL scheme the placement and number of repeaters can be unaltered, allowing the

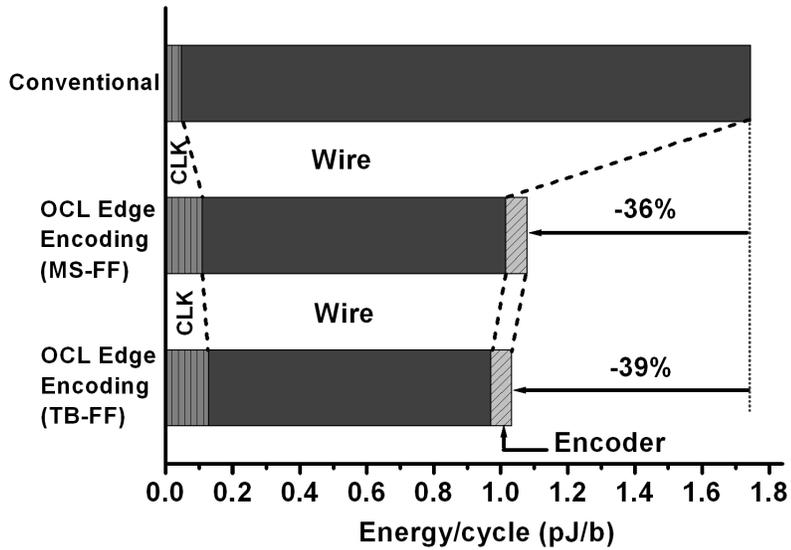


Figure 2.11: Energy breakdown of a 5mm wire for conventional and OCL edge encoding scheme.

Table 2.3: Performance, energy, and latency comparison for identical repeater sizes in OCL edge encoding scheme.

Total Length	Flop Distance	Performance Gain	Energy Reduction	Latency Overhead
10mm	1mm	-8.2%	-4.9%	10%
10mm	2mm	4.6%	10.0%	20%
12mm	3mm	11.7%	15.3%	25%
12mm	4mm	17.1%	18.1%	33%
10mm	5mm	22.2%	21.1%	50%

designer to simply drop-in the encoder and additional flip-flop to enable edge encoding. Results of this approach using identical repeater placement and sizes to the conventional repeater scheme are summarized in Table 2.3. Total wire length of 10-12mm is assumed for flop distances of 1-5mm, and the latency overhead is calculated as the relative overhead of encoding (1 cycle) to the number of cycles needed to propagate through the entire interconnect for each flop distance. As flop distance increases, the wire energy sufficiently dominates, allowing OCL edge encoding to achieve larger performance improvements and energy reductions, at the expense of larger relative latency overhead.

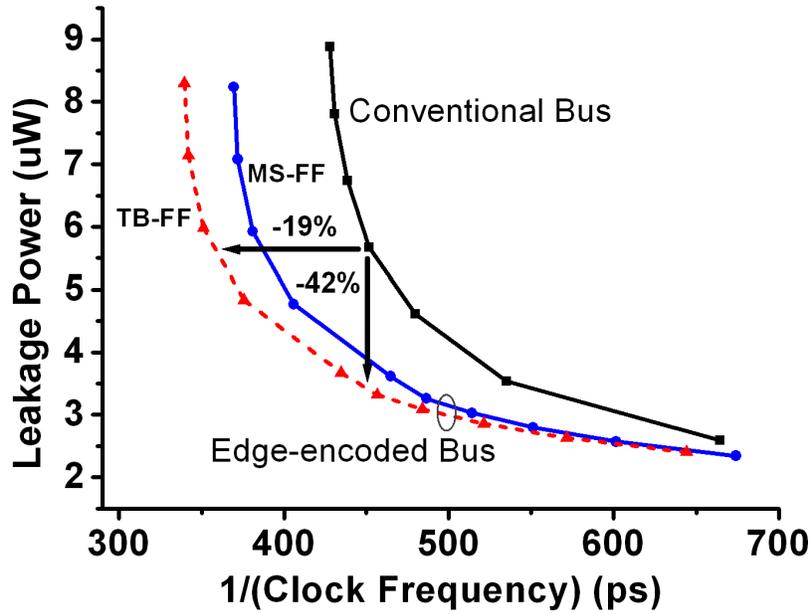


Figure 2.12: Leakage power comparison between a conventional and OCL edge-encoded bus. Ten repeaters are used in both cases and wirelength is 3mm.

2.3.3 Leakage Power Comparison

In sub-90nm technologies, leakage power in repeaters has become problematic [30]. Given that the proposed schemes includes more flip-flops than the conventional case, it is worthwhile to investigate the impact on static power. Table 2.2 shows a 26% reduction in leakage power for the TB-FF ZL edge encoding scheme over the conventional approach. This is achieved due to the use of smaller repeaters in the ZL edge encoding scheme.

Since the signal propagation in the ZL scheme needs to be completed a half cycle earlier than the OCL scheme, the repeaters in the OCL edge-encoded bus can be smaller than those in the ZL edge-encoded bus. Therefore, the leakage power, which is proportional to the size of repeaters, is expected to be less in the OCL scheme.

Given an identical number of repeaters in both conventional and OCL edge encoding schemes, Figure 2.12 shows the leakage power and delay characteristics for a 3mm wire. Due to the performance benefit of the edge-encoded bus, the repeaters in the conventional bus must be upsized to achieve the same performance as the edge-encoded bus, leading to higher leakage power. On the other hand, for identical repeater sizes and hence similar leakage power, the edge-encoded bus can operate at a higher clock frequency than the conventional bus. The OCL edge encoding

performance is further improved by using TB-FF rather than MS-FF as expected. Therefore, for iso-performance, we can see that the reduction of coupling capacitance allows fewer and smaller repeaters, resulting in leakage power reduction of 42%.

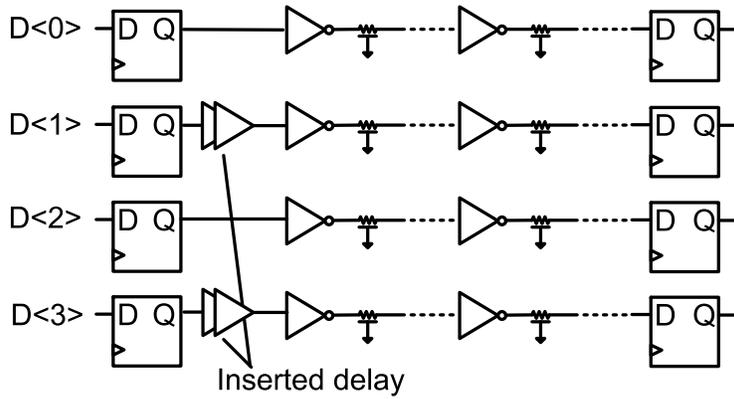
2.3.4 Sensitivity to Variation

Variability has become a major concern in modern CMOS technologies. Previous techniques [16, 36] rely on inserted delays and p/n skewing to separate the worst-case switching scenario, which are sensitive to process, voltage and temperature variation. We implemented the techniques proposed in [36] and [16] in the 65nm technology used in this chapter to compare the overall robustness of the achievable gains in the presence of different sources of variation. For on-chip communication circuits with long interconnects, local variation is less of a concern than global variation, because the inserted repeaters are large and the wire delay or energy is not that sensitive due to device mismatch. For 1,000 Monte Carlo runs, the spread in delay for the conventional bus was only 2%. Since global variation has a more dominant effect on the area of concern, we focused on global variation of process, supply, and temperature to analyze the robustness of each technique.

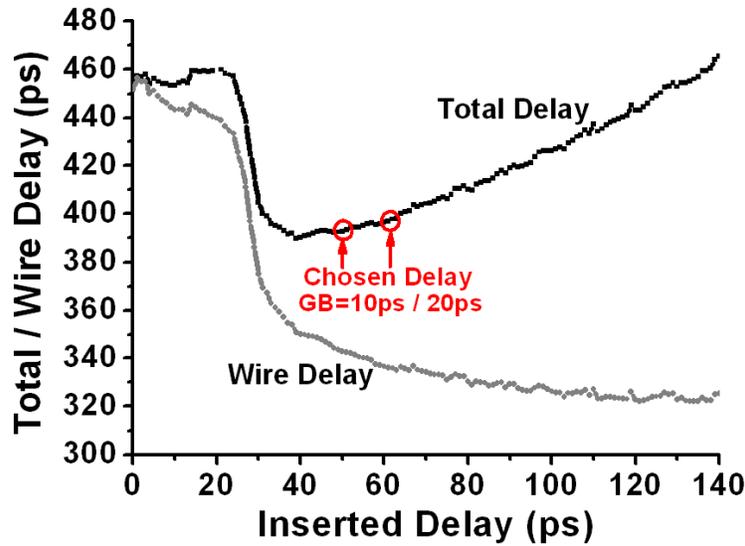
The technique in [36] inserts additional delay elements at the beginning of alternating wires. As more delay is added to adjacent wires, the worst-case switching is further separated, however this additional delay is included in the total delay leading to an optimal inserted delay as shown in Figure 2.13. Based on the total delay curve in Figure 2.13 we selected an inserted delay of 50ps and 60ps, which is guard-banded by 10ps and 20ps, respectively, to avoid the steep slope of the total delay curve for inserted delay of less than ~ 40 ps.

The technique in [16] propagates a specific transition (either 0-1 or 1-0) faster down the wire, such that the worst-case MCF pattern naturally separates as the signal propagates. However, this benefit is achieved at the expense of a greatly slowed transition in the non-preferred direction. In this experiment, alternating repeaters are skewed by modifying the width and length of repeaters and flip-flops.

Simulation results for the OCL edge-encoded bus with both MS-FF and TB-FF, staggered firing bus with 10ps and 20ps guard-banding, and skewed repeater bus at different process, voltage,



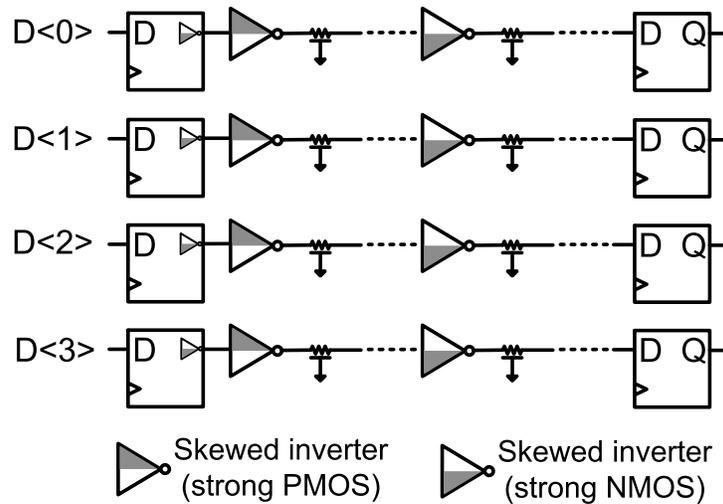
(a) Interconnect system of staggered firing bus.



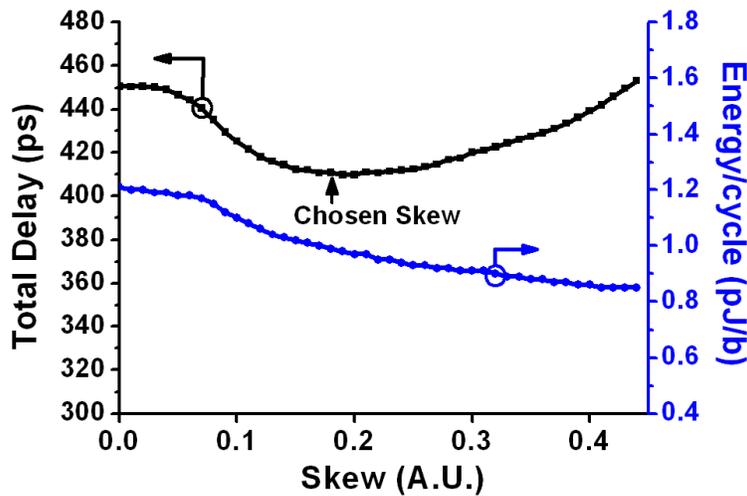
(b) Total/wire delay versus inserted delay.

Figure 2.13: Delay selection in staggered firing bus [36].

and temperature (PVT) corners are shown in Figure 2.15. We consider 10% supply voltage variation and 0-100C temperature variation. We first see that edge encoding provides the best overall performance improvement and energy savings. Edge-encoded bus with TB-FF provides additional performance improvement at the cost of small increase in energy consumption compared to edge-encoded bus with MS-FF across PVT corners. The added delay on alternate wires (staggered firing bus) or the intentional skew added to the repeater bus do not retain the delay and energy characteristics of the MCF=1 case when they separate worst-case switching events. This can be seen in Figure 2.13(b), where the wire delay at the chosen delay points is still larger than the minimum wire delay, and in Figure 2.14(b), where the energy at the chosen skew is still larger than the energy



(a) Interconnect system of skewed repeater bus.



(b) Total delay and energy versus skew.

Figure 2.14: Skew selection in skewed repeater bus [16].

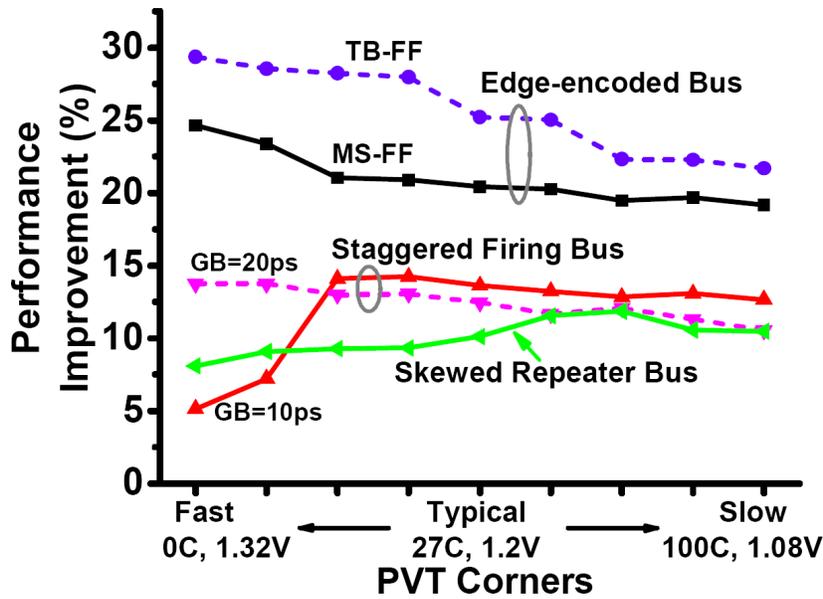
consumption with further separation of adjacent switching.

Furthermore, edge encoding is more robust across all PVT corners, achieving consistent energy savings. In both [36] and [16] improvements in delay and energy are achieved at the expense of susceptibility to PVT variation. Additional guard-banding to improve robustness to variation will result in less delay savings, as shown in the staggered firing bus curve with 20ps guard-banding in Figure 2.15(a). In Figure 2.15(b), the energy consumption of 20ps guard-banding is slightly better than that of 10ps guard-banding due to further separation of opposite switching, but the amount of energy reduction will saturate as we increase the guard-banding while the performance will

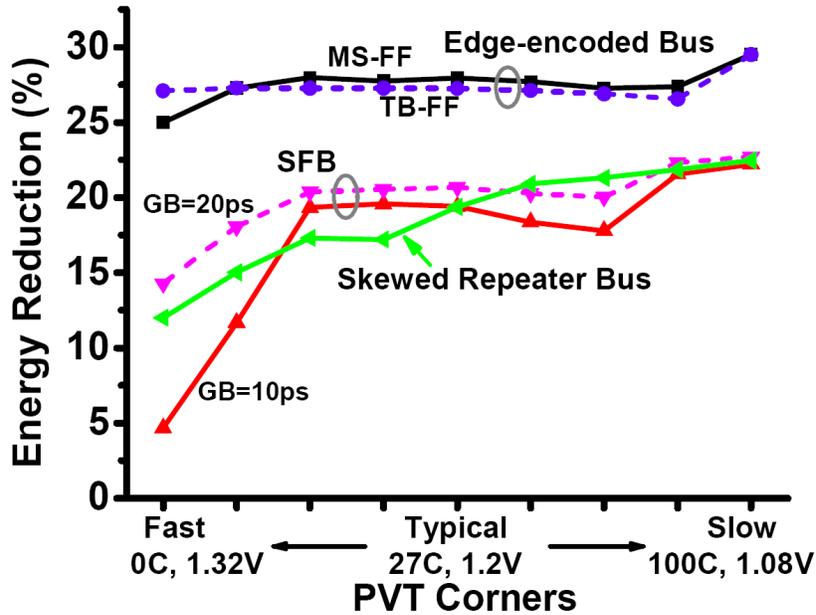
continue to degrade.

2.4 Summary

In this chapter, we proposed two new edge encoding techniques to improve energy efficiency and performance for multi-cycle on-chip interconnects. Both master-slave flip-flops and time-borrowing flip-flops were studied for optimal delay and energy. Compared to previously proposed techniques, edge encoding reduces both peak and average energy with improved robustness to process, supply, and temperature variation. For typical flip-flop distances of 2-5mm (corresponding to clock speeds of 1.3-2.5GHz in 65nm CMOS), the new techniques achieve 20-34% energy reduction without any overall latency, and 26-39% at the same throughput when one-cycle latency is introduced, comparing to conventional static bus in a multi-cycle interconnect.



(a) Performance improvement across PVT corners



(b) Energy savings across PVT corners

Figure 2.15: Sensitivity of improvements against process, supply, and temperature (PVT) variation for OCL edge-encoded bus (with MS-FF and TB-FF), staggered firing bus with 10ps and 20ps guard-banding (GB) [36], and skewed repeater bus [16] (flop distance: 3mm).

CHAPTER 3

Alternating Repeater Technique

This chapter describes an alternating repeater insertion technique that uses correct-by-construction polarities to reduce worst-case miller coupling factor (MCF) across any multiple segmented portion of a repeated bus. Simple static CMOS circuits with nominal p-n skews allow drop-in replacement while maintaining robust operation. For the same repeater area, number and position of repeaters of conventional busses, this technique simultaneously reduces delay by 15%, energy by 29% and peak current by 12% for 2~8mm on-chip busses in 1.2V, 65nm CMOS. Under equal delay constraints, the proposed technique reduces worst-case energy and peak current by 39% and 36%, respectively. The technique easily extends to shared busses for multi-core designs and shows a 41% improvement in energy-efficiency for a 10mm 5GHz multi-cycle on-chip core-to-core bus.

3.1 Related Work

Techniques to improve bus performance through MCF reduction have been proposed before [32–38]. In [32], a transition encoded dynamic bus is used to maintain the energy profile of a static bus with the performance of a dynamic bus. This technique requires two transitions in a single cycle which reduces the potential for energy reduction, especially at relaxed delay targets. [33, 34] propose bus encoding techniques to minimize worst-case MCF switch patterns for bus power reduction. These encoding approaches result in area, delay and power overhead that make them practically infeasible for on-chip busses. The active shielding work in [35] simultaneously switches the shield wires to enable MCF=0. This approach is not practical for busses as it results

in a considerable wiring overhead. In [36], the authors proposed delaying every alternate wire to decouple adjacent transitions and getting an overall bus delay that is faster than $MCF=2$ but slower than $MCF=1$. This technique requires finely tuned delay elements, which are sensitive to process variations. Guard banding for variations may negate much of the performance gains with the technique in [36]. The staggered repeater approach in [37] reduces average wire MCF at the cost of doubling the repeater blocks. This complicates the physical design and reduces the feasibility of the approach.

Recently in [38] the authors proposed flipping the switching relationship between adjacent wires only at the midpoint of the bus to reduce bus delay as well as the delay uncertainty. However, this technique does not benefit shared busses in multi-core designs with multiple driver and receiver points along the length of the bus. Therefore, using the technique in [38], multi-drop busses can still experience worst-case MCF of 2 at an intermediate point, which is not favorable.

The alternating repeater approach proposed in this chapter [18] is a simple drop-in replacement for conventional repeaters that does not alter the physical placement or size of repeater blocks and achieves average MCF over any multi-segmented length of the wire close to 1 in the worst-case. This technique also enables consistent gains for shared bus architectures with multiple source and driver points on the bus as well as across process skews and for irregular repeater placement.

3.2 Alternate Repeater Concept

Worst-case switching of opposite transitions on adjacent wires results in $MCF=2$ for every wire segment of a conventional bus (Figure 3.1). The alternating repeater technique (Figure 3.2) uses inverting (single inverter) and non-inverting (2 inverters) repeaters alternately along the wire. The ordering of the repeaters is flipped for the adjacent wires. The signal polarities with this design ensure that consecutive wire segments for any wire on the bus experience $MCF=2$ and $MCF=0$ when a neighbor wire switches simultaneously. As a result, the effective MCF for the length of any wire is closer to 1. Consequently, this design technique also guarantees MCF reduction for shared bus architectures with multiple driver sources and receivers along the length of the shared bus, as long as the driver and receiver are separated by more than 1 segment. Figure 3.3 shows the effect on delay (relative to the conventional bus) of replacing conventional repeaters with the

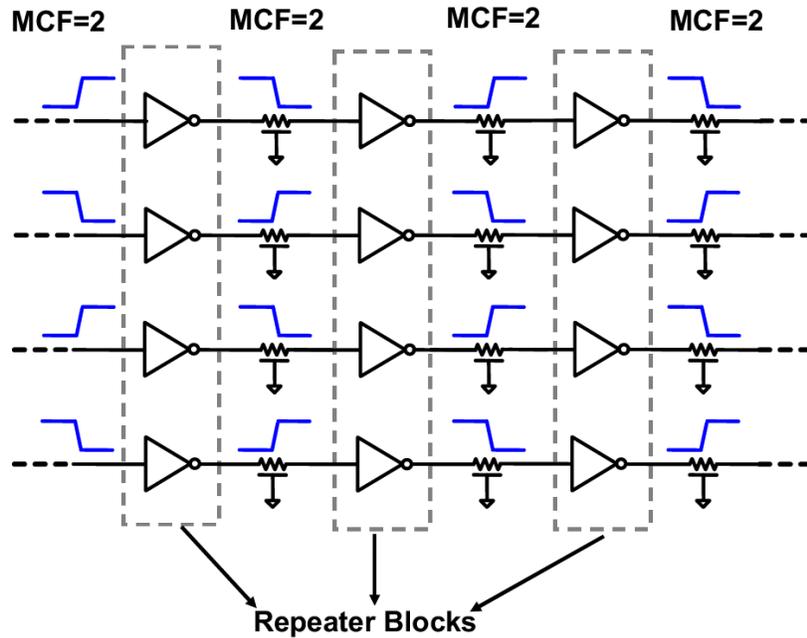


Figure 3.1: Worst-case switching pattern for conventional bus design.

proposed alternating repeaters. The contributors to change in wire delay are - (i) speed up from MCF reduction, (ii) speed up of wire segments loaded by the smaller input capacitance of non-inverting repeaters and (iii) increase in repeater delay for non-inverting repeaters due to the extra inverter. The significant wire MCF reduction enables overall wire delay to reduce considerably (25%).

Busses with even number of segments result in exactly 50% of the wire length with MCF=2 and MCF=0. For busses with an odd number of segments, >50% of the bus wire results in MCF=2 for worst-case delay. This discrepancy from the ideal of 50% for busses with odd number of segments reduces with increased wire lengths and scaled technologies as the number of repeated wire segments increase to compensate for the increased wire resistance. Increased wire segments result in smaller differences between odd and even segmented busses. All quantitative results in this chapter are conservative as only busses with an odd number of wire segments have been used. Gains with the proposed technique improve for busses with an even number of segments.

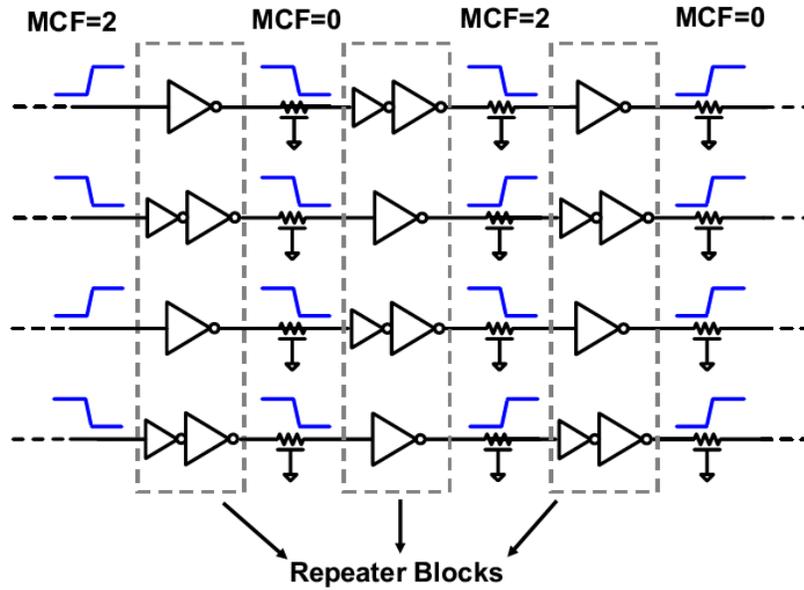


Figure 3.2: Worst-case switching pattern for proposed alternating repeater bus.

3.3 Bus Comparisons

Comparisons between conventional and alternating repeater based busses are shown under equal delay and repeater area constraints. The feasibility of the gains with alternating repeaters across process skews and with irregular repeater placement are shown to emphasize the robustness of this technique. The impact of replacing conventional repeater blocks with alternating repeaters in a high-performance multicore multi-cycle on-chip bus is also shown.

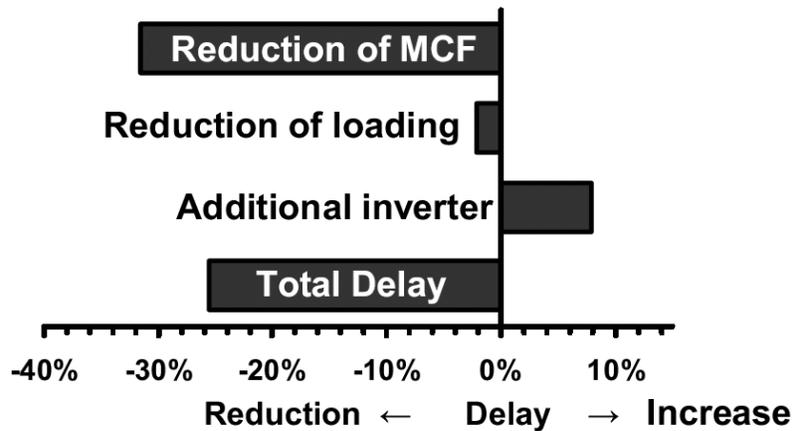
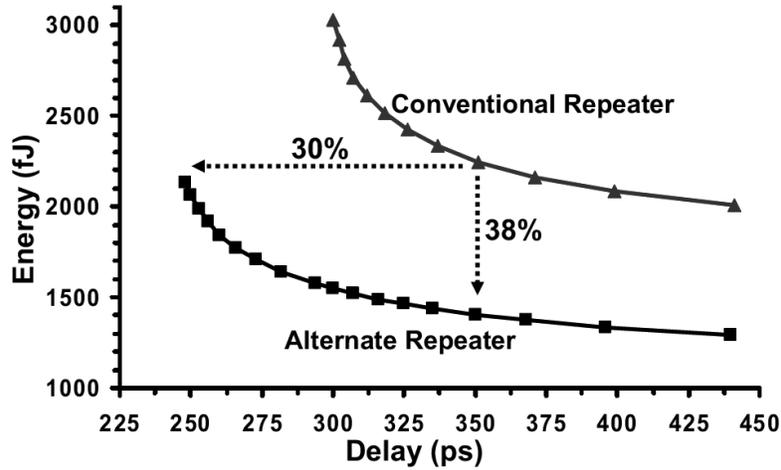
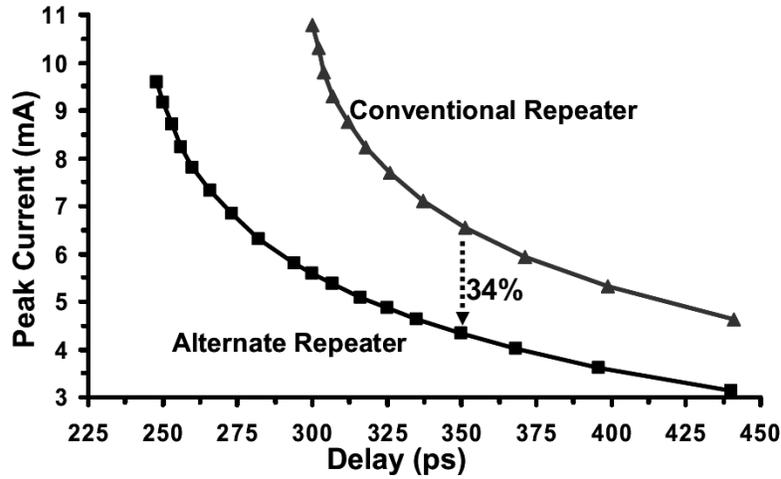


Figure 3.3: Impact on delay with alternating repeater technique.



(a)



(b)

Figure 3.4: (a) Energy-delay and (b) Peak Current-Delay comparisons for 5mm bus in 65nm technology.

3.3.1 Equal Delay and Area

Static busses of various lengths were optimized using conventional repeaters. The optimization involved tuning the number and size of repeaters for minimum worst-case energy for a given delay target in 65nm CMOS technology [31].

For comparisons at equal delay targets, alternating repeater blocks replace the conventional repeater blocks and only the transistor sizing of alternating repeaters is optimized without affecting p-n skew or number/placement of repeater blocks. Figure 3.4 shows comparisons of a 5mm bus (Metal 6, 1.2V, 110C, 65nm CMOS) for worst-case energy/bit and peak current/bit over a

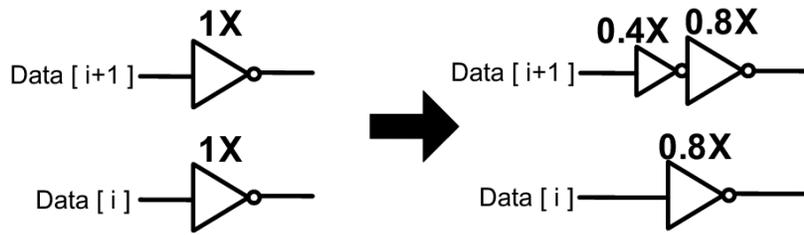


Figure 3.5: Sizing optimization for same repeater area.

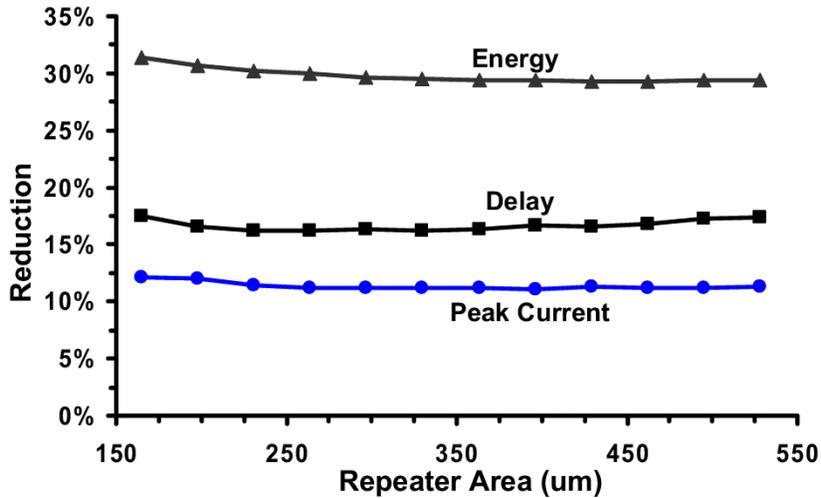


Figure 3.6: Energy, delay and peak current reductions for the same repeater area with alternating repeaters for 5mm bus in 65nm technology.

wide range of delay targets. The MCF reduction lowers the worst-case energy consumption with alternating repeaters. Improved performance from reduced MCF allows downsized repeaters for the same delay, which further reduces the energy consumption and peak current demands. Peak currents in repeater blocks can be significant due to large driver sizes which makes it an important metric for reducing localized voltage droops and the area dedicated to local decoupling capacitors.

For comparisons at equal transistor widths, every pair (of adjacent bits) of conventional repeaters in a repeater block is converted to alternating repeaters using the sizing methodology shown in Figure 3.5. These fan-out and driver size ratios under the same total transistor width constraint (2X per 2 bits) result in a balanced delay redistribution across the inverting and noninverting repeater stages. The number of repeater blocks and p-n skews of the repeaters are unchanged. Figure 3.6 shows the improvements in worst-case energy, delay and peak current for a 5mm bus (Metal 6, 1.2V, 110C, 65nm CMOS) for a range of total repeater sizes (per bit).

Table 3.1: Average (2~8mm) alternating repeater bus advantages in 65nm technology.

	Equal delay	Equal transistor width
Delay reduction	0%	15%
Total transistor width reduction	34%	0%
Energy reduction	39%	29%
Peak current reduction	36%	12%
Leakage current reduction	35%	0%

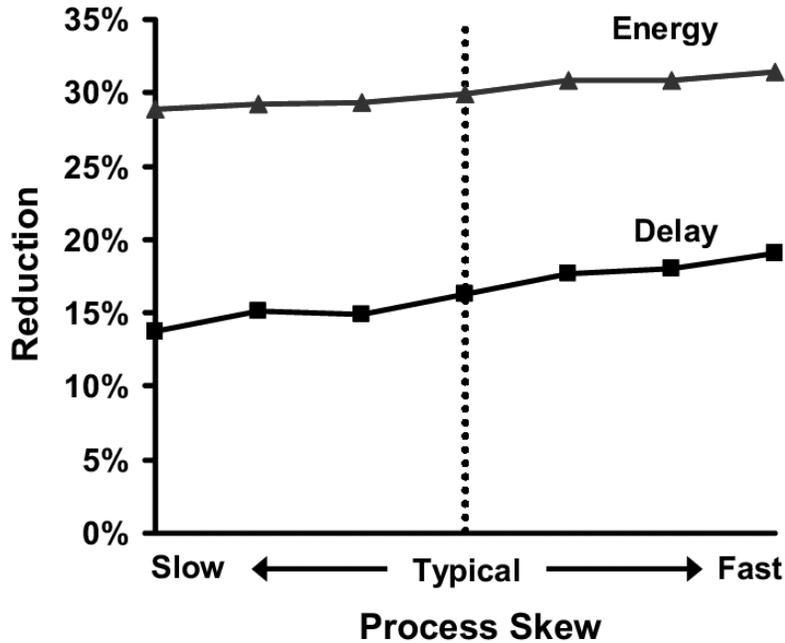


Figure 3.7: Energy / delay gains across process skews for 5mm alternating repeater bus in 65nm technology.

The advantages with the alternating repeater technique for Metal 6 busses ranging in length from 2~8mm under delay and area constraints are summarized in Table 3.1 for a 65nm CMOS technology at 1.2V, 110C. As described in Section 3.2, the reported gains are conservative as the busses have an odd number of segments.

3.3.2 Process Skew and Repeater Placement Sensitivity

Since the benefits of the alternating repeater technique are obtained primarily due to the correct by construction polarities, the technique is robust, allowing consistent gains across a range of process corners. Figure 4.14 shows the delay and worst-case energy reductions across a range of

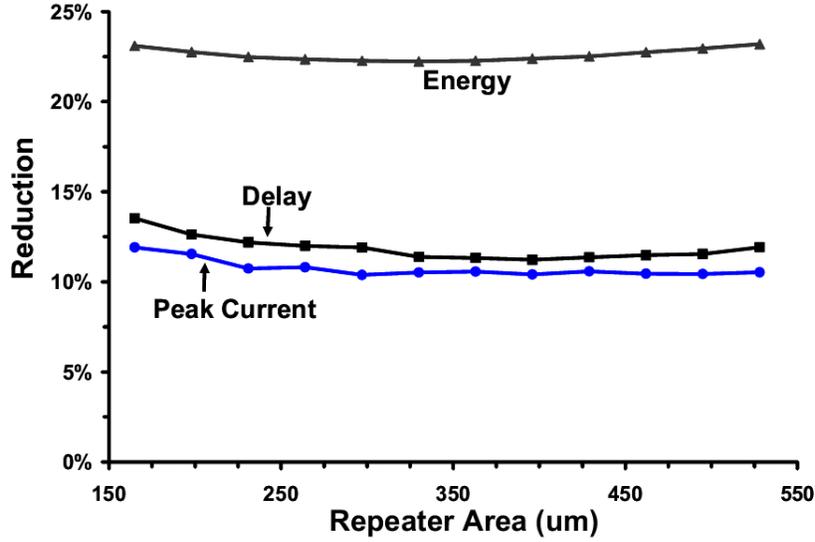


Figure 3.8: Energy, delay and peak current reductions for the same repeater area with alternating repeaters for 5mm bus with nonequidistant repeater placement.

Table 3.2: Comparison of multi-cycle 64b 10mm bus in 65nm technology.

	Frequency	Power @ 100% activity	Energy efficiency	Power @ 10% activity	Peak current	Total transistor width
Conventional	5GHz	1.413W	226Gbps/W	0.178W	333mA	35040 μ m
Proposed	5GHz	1.004W (-28.9%)	318Gbps/W (+40.7%)	0.134W (-24.7%)	253mA (-24.0%)	28952 μ m (-17.4%)

process corners for the same alternating repeater bus that has been optimized at the typical corner for the same repeater area as the conventional bus. The delay and energy gains vary about the typical corner by $\pm 2.8\%$ and $\pm 1.5\%$, respectively.

Irregular or non-equidistant repeater placement can affect the achievable gains with this technique since the proportion of total wire segment with MCF=2 can increase beyond the ideal 50%. However, the reduction in MCF to 0 for the remaining wire length with the alternating repeater technique enables considerable gains even with non-equidistant repeater placement. Figure 3.8 shows the resulting gains for non-equidistant repeater placement for a 5mm bus under a constant repeater area constraint (similar to the analysis of Figure 3.6). The repeaters have been displaced by 20% from their ideal position, resulting in alternating short and long segments, with the long segments 20% longer than ideal inter-repeater distance. Along with the fact that this bus already uses an odd number of segments, these conservative gains still result in reduction in delay, energy and peak currents of 12%, 23% and 11%, respectively, over a wide range of the design space.

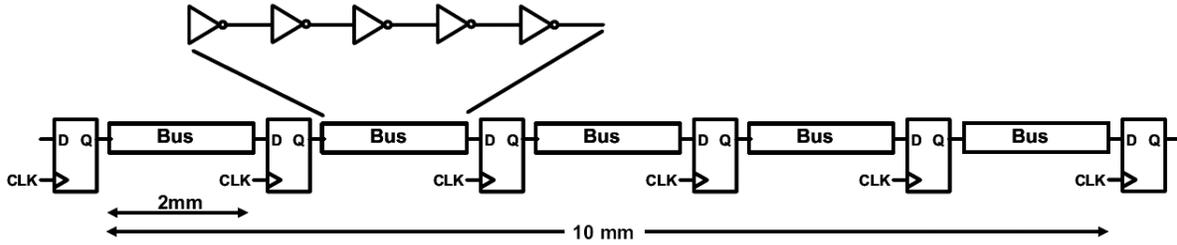


Figure 3.9: A core-to-core/core-to-cache 64b bus for a high-performance multi-core microprocessor.

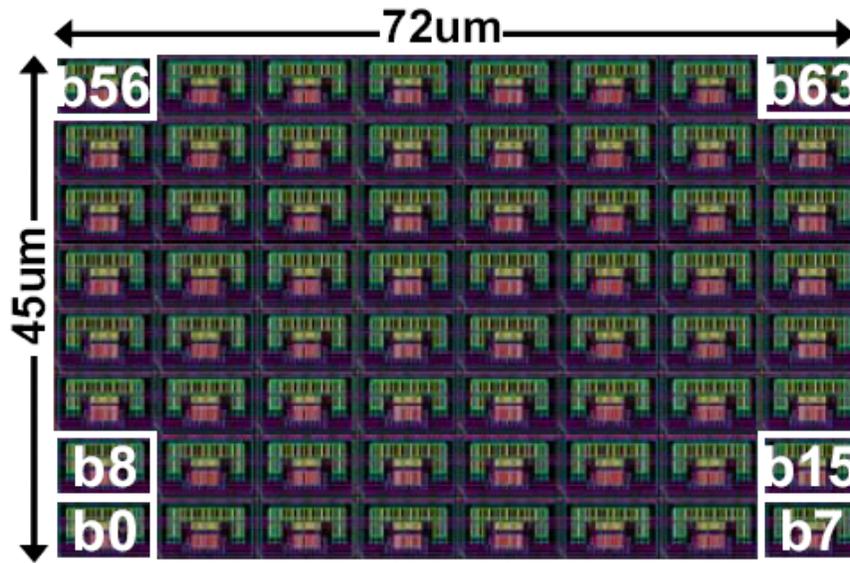


Figure 3.10: 64b driver/repeater block layout

3.3.3 Multi-Cycle Bus

Reduced cycle times and increased wire propagation delays have resulted in the wide use of multi-cycle on-chip busses. Figure 3.9 shows a performance-critical shared bus design used for core-to-core or core-to-cache communication in a multi-core high-performance microprocessor. The conventional 64b bus supports a 5GHz clock frequency with a 5 cycle latency over a 10mm distance in 65nm CMOS (1.2V, 110C). This bus has been re-optimized with alternating repeaters for the same clock frequency without disturbing the placement of flops and repeaters (Figure 3.10). The results (Table 3.2) show a 40% improvement in energy efficiency, 29% reduction in worst-case power (100% data switching activity), 25% reduction in power at lower (10%) data switching activity and 24% lower peak currents.

3.4 Summary

An alternating repeater technique has been proposed to reduce worst-case MCF for on-chip busses. Simultaneous reductions in delay (15%), energy (29%) and peak current (12%) for 2-8mm on-chip busses in 1.2V, 65nm CMOS have been shown under equal repeater area constraint over a wide range of the design space. The energy and peak current gains increase to 39% and 36%, respectively, for equal delay constraints. The technique is robust across process corners and for bus designs with nonequidistant repeater placement. The proposed technique can be integrated in shared bus designs for multi-core processors and enables 40% higher energy-efficiency for a multi-cycle 64b bus on a multi-core processor at the same 5GHz clock frequency.

CHAPTER 4

Crosstalk-Aware Pulse Width Modulation based Signaling

Continuous technology scaling results in tighter wiring pitch with higher coupling capacitance and crosstalk noise, which directly impacts maximum clock frequency and chip power consumption. Typically the wiring pitch is kept small in microprocessor designs to accommodate a number of wires in a given routing area, and this brings up a number of concerns as we scale beyond 65nm. To address these problems, we propose to convey two bits of information in one wire through PWM-based signaling, where the pulse width carries the 2-bit information. Exploiting the controllability of pulse widths in pulsed signaling [22, 23], two bits of information are sent on one wire. By halving the number of wires in multi-bit busses, the wire width and spacing can be doubled for the same overall footprint, leading to a decrease in wire resistance and coupling capacitance, and thus improved delay and energy consumption.

4.1 Related work

A number of techniques, including encoding [10, 20], pulsed signaling [22, 23, 40, 41], and alternating repeaters [17–19], were proposed to improve the performance and energy of global wires. Most of these techniques reduced the effective coupling capacitance to lower energy consumption, without changing the number of wires. Previous work on pulsed signaling showed good improvements in energy and/or performance over conventionally repeated interconnect. Static pulsed bus [22] reduces worst-case Miller coupling factor (MCF) to 1 and gains further performance by skewing the repeaters, but peak energy consumption penalties existed. Short pulses are combined with

the voltage doubling property of properly terminated transmission lines in [40], but the bandwidth density is poor due to wide wires. Other pulse signaling techniques incurred design complexity of dual supplies [23], or additional delay elements and signals [41].

On the other hand, there have been approaches to reduce the number of wires in multi-bit busses by sending more than one bit of information over one wire in a clock cycle. The authors in [42] proposed wave-pipelined multiplexed (WPM) interconnect for time-domain multiplexing. However, the WPM scheme requires an additional clock signal, used relatively wide wires ($1.05\mu\text{m}$ width), and did not consider the effect of crosstalk noise. Current-mode quaternary signaling [43] was proposed to encode two data bits into four different current levels, but the energy consumption does not scale with data activity. Well-known double-data-rate (DDR) signaling sends one bit of data per each clock phase. References [44–46] considered serializing data streams for on-chip links, and the incoming data are multiplexed with different phases of the clock or strobe. These serialization techniques have concerns of dynamic power consumption at low data activities, Even if the incoming data bits are idle, in case the following bit has the opposite polarity with the previous bit (50% probability), dynamic power will be consumed. For data traces with low activity, the average energy consumption will certainly be dependent on the data pattern.

4.2 PWM-based signaling concept

This chapter presents new encoding techniques that combine the benefits of pulsed signaling and 2-bit signaling. We propose two pulse width modulation based encoding techniques to improve both performance and energy consumption in on-chip global interconnect. This is done by exploiting the controllability of pulse widths in the pulse generators and using the pulse width to contain 2 bits of information (i.e., pulse width modulation). The encoder will compress the 2-bit data into pulse width and transition type, and after propagating through repeaters over long wires, the decoder reconstructs the original 2-bit information.

By halving the number of wires in multi-bit busses, the wire width and spacing can be doubled for the same overall footprint, leading to a decrease in wire resistance and coupling capacitance, and thus improved delay and energy consumption. These savings are accomplished at the expense of encoder and decoder logic, but there is no additional clocking overhead. However, we find

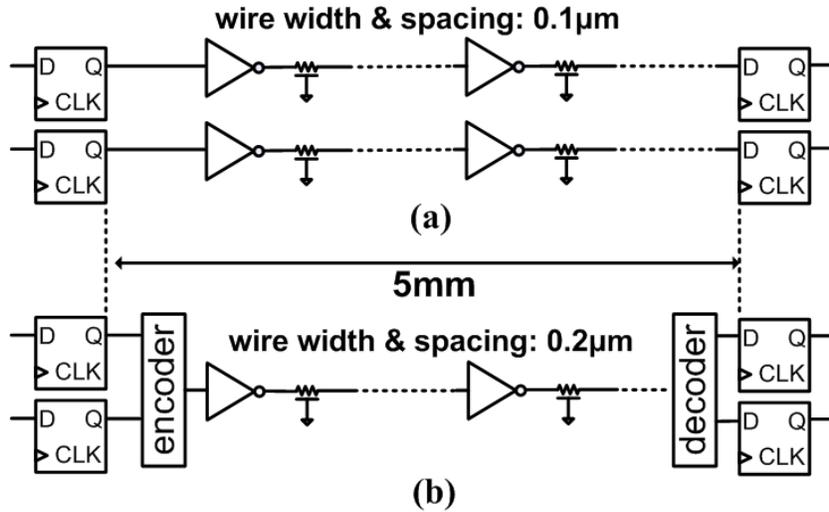


Figure 4.1: (a) Conventional bus with repeaters (b) Proposed PWM-based bus with repeaters (both schemes have same footprint).

that the encoding and decoding overhead is small in long interconnects where interconnect power consumption is dominant. The effect of crosstalk noise on generated pulses is dynamically pre-corrected in the encoder, and the encoder and decoder circuits are able to self-calibrate against die-to-die and within-die variation. Moreover, the proposed techniques are transition-encoded, hence the on-chip link does not consume any switching energy when the bus remains idle, further reducing the average energy consumption of the overall system. The preliminary version of this paper appeared in [15].

Figure 4.1(a) shows the configuration of the conventional repeated bus, where long interconnect between flip-flops is optimally repeated. Typically the wiring pitch is small to accommodate a number of wires in a given routing area, and this brings up a number of concerns as we scale beyond 65nm. To address these problems, we propose to convey two bits of information in one wire through pulse signaling, where the pulse width carries the 2-bit information. We point out that typically in pulsed signaling [22], the pulse width is controllable with variable delay elements, and depending on the 2-bit switching, different pulse widths are sent along the interconnect. The according proposed bus system is shown in Figure 4.1(b), where one wire replaces two wires of the conventional bus through an encoder and decoder. This way, for the same footprint with the conventional scheme, the wire width can be increased to reduce the wire resistance or the wire spacing can be increased to reduce the coupling capacitance. Both will lead to better performance

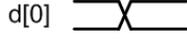
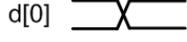
Conv.	Mono-PWM	Hybrid-PWM
d[0] _____ d[1] _____	—	— (—)
d[0]  d[1] _____		 ()
d[0] _____ d[1] 		 ()
d[0]  d[1] 		 ()

Figure 4.2: Concept of PWM-based signaling for proposed mono-PWM and hybrid-PWM schemes.

and energy consumption of the multi-bit busses.

We propose two specific schemes for implementing these new encoding techniques. The concept of the two PWM-based signaling approaches is illustrated in Figure 4.2. Both proposed schemes use transition-based encoding to suppress power consumption with non-switching inputs. When the data is idle, the encoders of both proposed schemes will stay idle as well, resulting in no transition. The first scheme, namely monotonic PWM scheme (stands for monotonic pulses), realizes pulsed signaling based on pulse width modulation for on-chip data communication. Figure 4.5 shows the encoder and decoder circuits of the mono-PWM scheme. When at least one bit switches, the encoder generates a pulse with three different pulse widths depending on the bit switching pattern, and this pulse propagates through the repeated interconnect. The decoder evaluates the pulse width and converts this information back to the original 2-bit data.

The second scheme, named hybrid PWM scheme, combines single transition signaling with PWM-based signaling to further reduce energy consumption. In one of three switching cases, the hybrid-PWM encoder generates a single transition instead of a pulse, while the two remaining cases lead to two different pulse widths. Similar to the mono-PWM system, the decoder interprets the pulse width information and generates the original 2-bit data. The hybrid-PWM encoder and decoder circuits (Figure 4.6) are slightly more complex than the mono-PWM scheme since the quiescent state can be either high or low. However, the average energy consumption for random

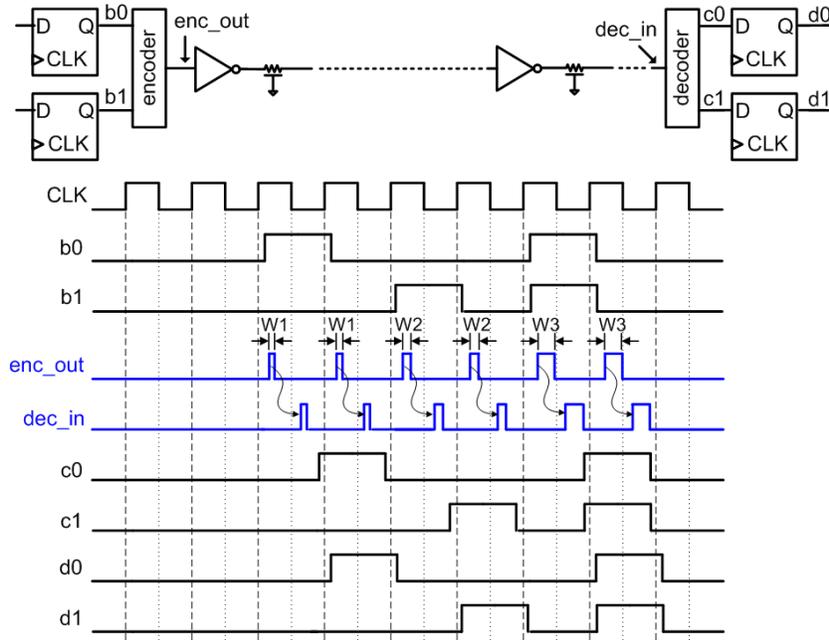


Figure 4.3: Timing diagram of monotonic PWM scheme.

data is reduced due to the use of single transitions. Note that there is no clock overhead in the encoder and decoder in both schemes.

The detailed timing diagram of the two proposed schemes based on pulse width modulation is shown in Figure 4.3 and Figure 4.4. When the data from previous flip-flops (b0, b1) are idle, the encoders of both mono-PWM and hybrid-PWM will stay idle as well, maintaining the same state at the output flip-flops.

In Figure 4.3, the timing waveforms of mono-PWM system is described. For different switching patterns of b0 and b1, appropriately different pulse widths are generated from the encoder, and these pulses with distinct widths propagates through the long wires. The same delay chains in the encoder are utilized in the decoder to recover the variable width pulses to 2-bit non-return-to-zero (NRZ) data. Figure 4.4 shows the timing details of hybrid-PWM system. The main difference between mono-PWM system is that we propose to signal single transition in one of the three switching cases, and only two different pulses for the remaining two scenarios. The encoder and decoder circuitry becomes slightly more complicated, but the average energy consumption on random data could be reduced due to single transitions. Similar to mono-PWM system, the decoder interprets the pulse width information and delivers the original 2-bit information to the next flip-flops.

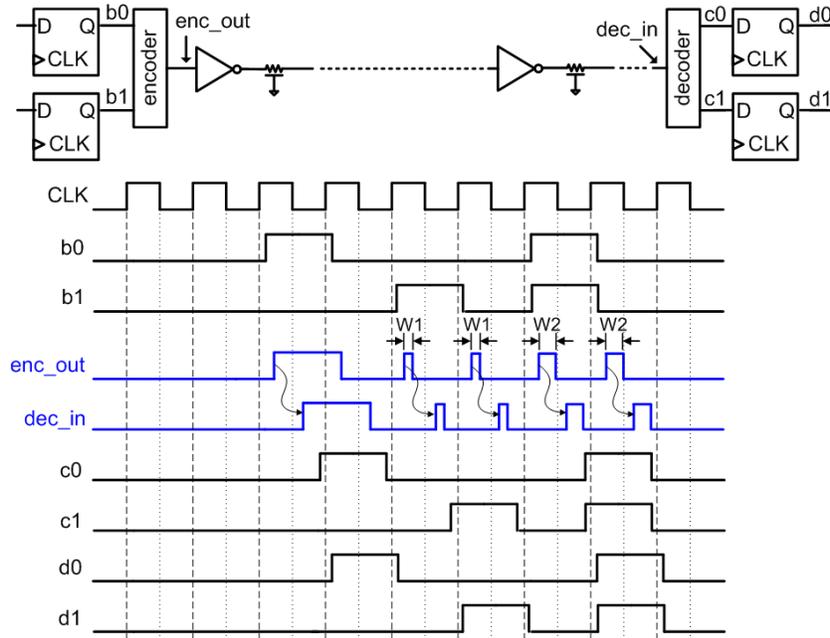


Figure 4.4: Timing diagram of hybrid PWM system.

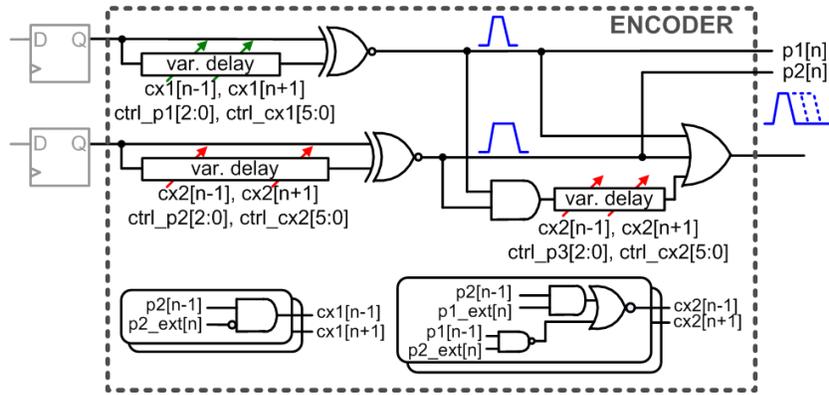
4.3 Proposed Encoder and Decoder Circuits for PWM-based signaling

In this section, the encoder and decoder circuits for both mono-PWM and hybrid-PWM are shown and the operations upon data transitions are described.

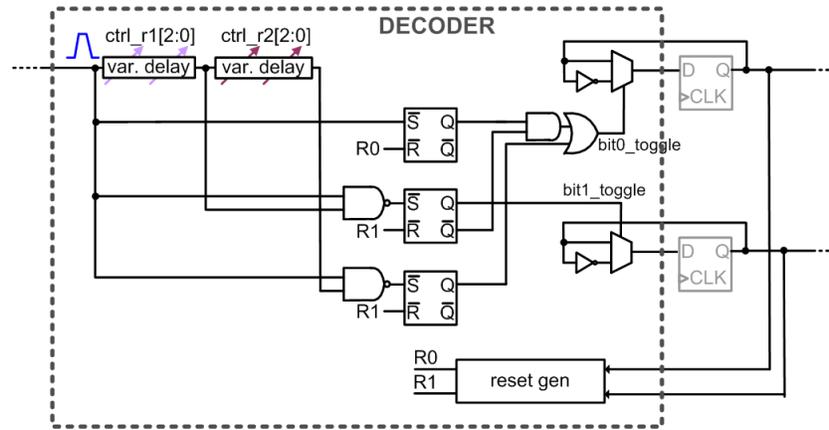
4.3.1 Mono-PWM encoder and decoder

In Figure 4.5(a), the encoder circuit of mono-PWM is shown. Each bit from the previous flip-flop stages is sent to separate pulse generators with variable pulse widths. The default pulse width of the pulse generator that the second bit goes through is larger than that of the pulse generator the first bit goes through. Since pulse generators only detect transitions, when the incoming bits are idle, there is no circuit activity inside the encoder. When only one bit out of the two transitions, the output of either one of those internal pulse generators is chosen and sent to the encoder output. In case both bits transition, the AND gate and the additional delay chain at the bottom right part extends the pulse width and the widest pulse is generated at the encoder.

The decoder circuit of mono-PWM is shown in Figure 4.5(b). When the incoming bits from



(a) Encoder circuit.



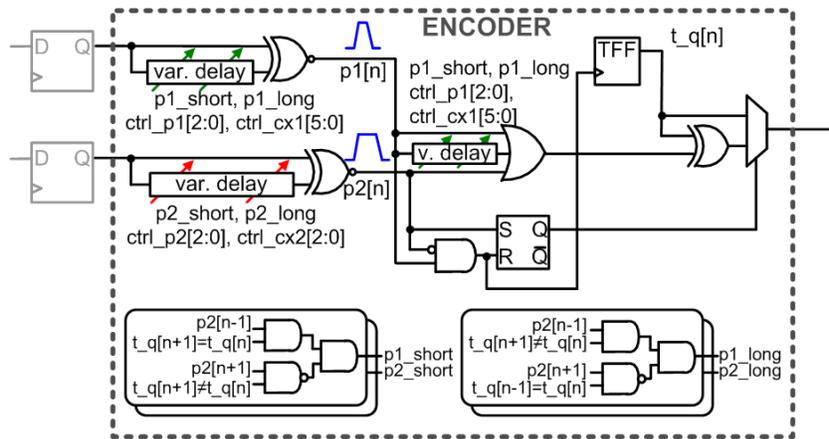
(b) Decoder circuit.

Figure 4.5: Encoder and decoder circuits for mono-PWM signaling.

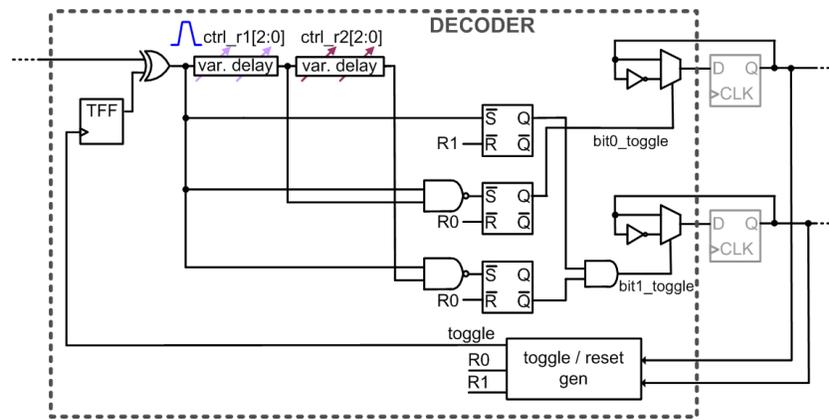
the wire is idle, obviously the decoder stays idle and the next flip-flop outputs are unaltered. When a pulse with a certain width arrives at the input of the decoder, the according delay chains compare its delay with the pulse width and determines the range of the incoming pulse width. After going a set of SR latches and simple combinational logic, the output of the decoder is decided and directly sent to the next flip-flops. To prevent any kind of clocked reset, we only reset the SR latches when there was a transition in the according bit in the previous cycle. This way minimal energy is consumed on the resetting logic.

4.3.2 Hybrid-PWM encoder and decoder

In Figure 4.6(a), the encoder circuit of hybrid-PWM is shown. Each bit from the previous flip-flop stages is sent to separate pulse generators with variable pulse widths, similarly to mono-PWM.



(a) Encoder circuit.



(b) Decoder circuit.

Figure 4.6: Encoder and decoder circuits for hybrid-PWM signaling.

Since pulse generators only detects transitions, when the incoming bits are idle, there is no circuit activity inside the encoder. Now, when only the first bit on top switches, the SR latch on bottom is triggered and the mux at the output of the encoder selects single transitioning. Based on the value stored in the toggle flip-flop, the next single transition of low-to-high or high-to-low is sent to the long wire. For the remaining two transitions, the pulse generator detects the bit transitions, and encodes the appropriate pulse width to the encoder output.

The decoder circuit of hybrid-PWM is shown in Figure 4.6(b). The basic structure is equivalent to the mono-PWM decoder. At the first part of the decoder, the toggle flip-flop is required to store the default low or high state of the wire and then allowing correct detection of low-high-low pulse or high-low-high pulse, respectively. The combinational logic is different after the SR latches are different from that in mono-PWM decoder due to the single transition. When the incoming bits

from the wire is idle, obviously the decoder stays idle and the next flip-flop outputs are unaltered. When a certain transition arrives at the input of the hybrid-PWM decoder, the according delay chains compare its delay with the pulse width and determines the range of the incoming pulse width as well as whether the transition is single transition or pulse transition. After going a set of SR latches and simple combinational logic, the output of the decoder is decided and directly sent to the next flip-flops. The resetting logic shown is very similar to the one showned in Figure 4.5(b), but in this case it also resets the toggle flip-flop in case single transition occurred.

4.4 Crosstalk and Variability Considerations

Since the pulse width contains data information in our proposed schemes, if the pulse width is altered by crosstalk noise or if it is interpreted differently by the encoder and decoder due to variability, the functionality of the system could be affected. In scaled technologies, crosstalk and variability will be worse effects on wires and transistors, hence they are analyzed carefully and we introduce techniques to mitigate their effect.

4.4.1 Crosstalk aware signaling

One of the main challenges in the proposed PWM-based signaling schemes is to mitigate the effect of crosstalk from adjacent wires on the transmitted pulse width. As illustrated in Figure ??, this challenge arises because the first edge (always rising in mono-PWM) of the pulses are aligned between adjacent signals, but the second edge (always falling in mono- PWM) of the pulses can be separated depending on the switching pattern, thereby modulating the pulse width. Ignoring this effect would result in excessive guard-banding of pulse width margins, negating the speed and energy improvements. Crosstalk effects could be avoided by shielding each wire, at the cost of degraded delay and energy consumption compared to double pitch wires.

To address this challenge without shielding, encoding circuits in Figure 4.5(a) and Figure 4.6(a) are designed to enable crosstalk-aware signaling. Since pulse width dependency on data switching is deterministic, shorter pulses are generated through variable delay chains in the encoders if the pulse will be lengthened by crosstalk, and vice versa. The pulse shortening (or lengthening) for

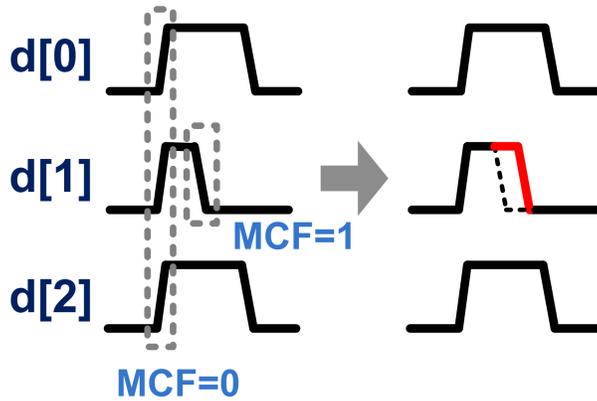


Figure 4.7: Due to crosstalk from adjacent bits, certain data patterns can result in changing pulse width over long wires.

each data pattern is analyzed through extracted SPICE simulations and can be controlled digitally in our design for testability purposes. Also, the pulse width for each input case was chosen such that the shortened (or lengthened) pulse due to crosstalk will not overlap with the next pulse width, resulting in a pulse step of 80-100ps. The crosstalk-awareness feature is implemented in both encoders of mono-PWM and hybrid-PWM scheme through the signals named as $cx1^*$ and $cx2^*$, which are generated by the additional NAND, AND, and NOR logic gates that are shown at the bottom part. These logic gates have inputs that are coming from encoders of adjacent bits, which means that the adjacent bit switching dynamically controls the current bit pulse width every cycle.

In this work, we addressed the effect of crosstalk using a pre-correction circuitry in the encoder and have the decoder receive the same pulses and signals regardless of the switching in the adjacent bits. To further suppress crosstalk in future scaled technologies, the decoder could include post-correction circuitry with inputs from adjacent far-end signals, adjusting pulse width changes due to crosstalk upon the signal reception.

4.4.2 Mitigating variability and self-calibration

Since the pulse width generators in the encoder and the pulse width detector in the decoder can behave differently with global and local variation, it may result in functionality failures of the proposed schemes. To minimize variability between the encoder and decoder in both schemes, identical variable delay chains are used and sized up to minimize RDF. Hence, the encoder pulse

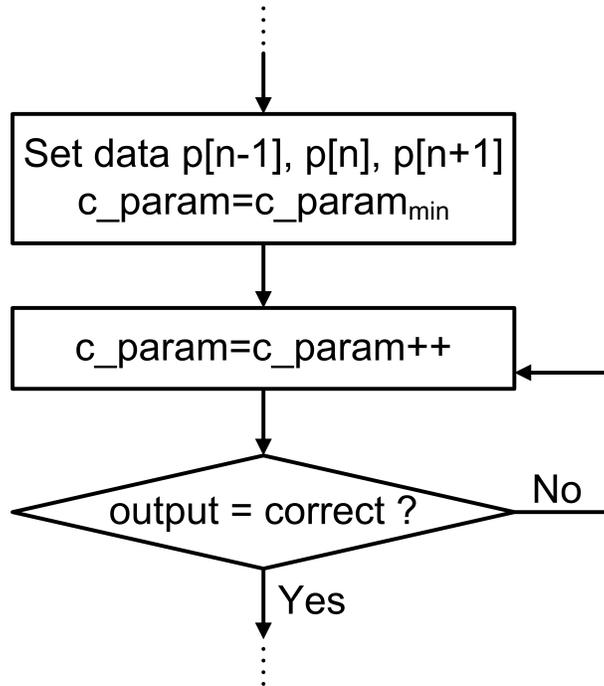


Figure 4.8: Self-calibration methodology for each control signal.

width and decoder delay track global PVT variation, but for calibration against local mismatch they are separately controllable by scan chain. Signals named as *ctrl_** in Figure 4.5(a) and 4.6(a) control the variable delay in the encoder and decoder by either turning on an additional MOSFET capacitance or adjusting the strength of the inverter that comprises the variable delay chain. All of these signals can be digitally controlled through scan chain for testability purposes, but it would be impractical to manually adjust all control signals for every wire independently by running data patterns. The simple and easy way to address the delay spread due to variation is to reserve enough margin and guard-band the pulse width sufficiently between different 2-bit switching. However, this could result in excessive performance penalty of the overall link system since the margin is required on top of the worst-case corner or worst-performance chip. On the other hand, it is practically impossible to tune or adjust the control signals of each encoder differently if the proposed scheme were applied to a wide multi-bit bus (i.e., 512 or 1,024 bit bus).

To avoid the difficulty of these methods, a self-calibration module is introduced. The purpose of the self calibration module is to automatically find the optimal values the control signals of encoder and decoder which govern the variable inverter chain delay and the crosstalk-aware fine-tuning, and adjust them accordingly.

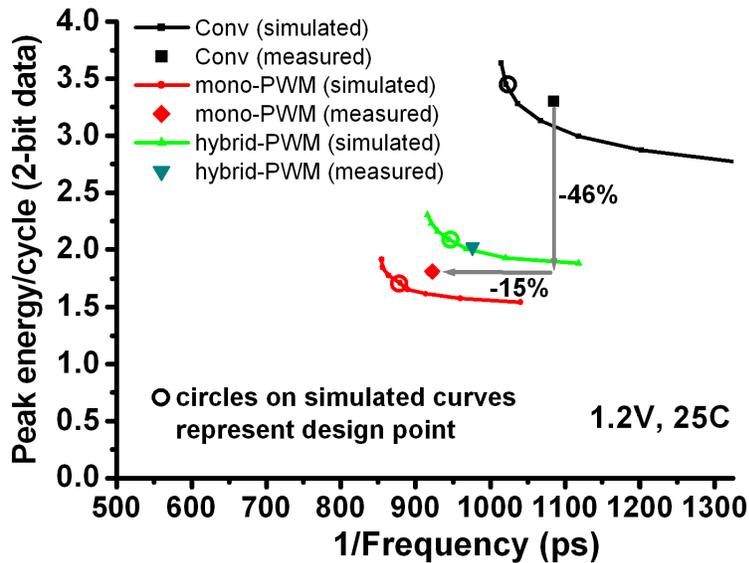


Figure 4.9: Energy vs delay comparison of conventional, mono-PWM, and hybrid-PWM bus system.

Upon initialization and every once in a while (which can be defined by the processor), the self-calibration circuitry will be enabled, and depending on the variation environment (process, voltage, temperature), the control signals will be adjusted automatically. First, variation due to local mismatch and process corners can be addressed by initial self-calibration, since the local transistor mismatch and process corners are static once the chip is fabricated and they will not change as time goes by. On the other hand, dynamic variations such as supply droop or local hot-spot due to temperature variation can still affect the encoder and decoder in opposite directions. However, these types of variation usually occur in a relatively slow manner, hence as long as the calibration process is enabled frequently enough, mitigating the transient variation is possible. Once enabled, the self calibration module loads pre-defined data patterns for each calibration set, and sweeps the control signals following the steps shown in Figure 4.8 for the encoder as well as the decoder. The total procedure takes ~ 200 cycles. The calibration module is fully synthesized and auto-placed-routed.

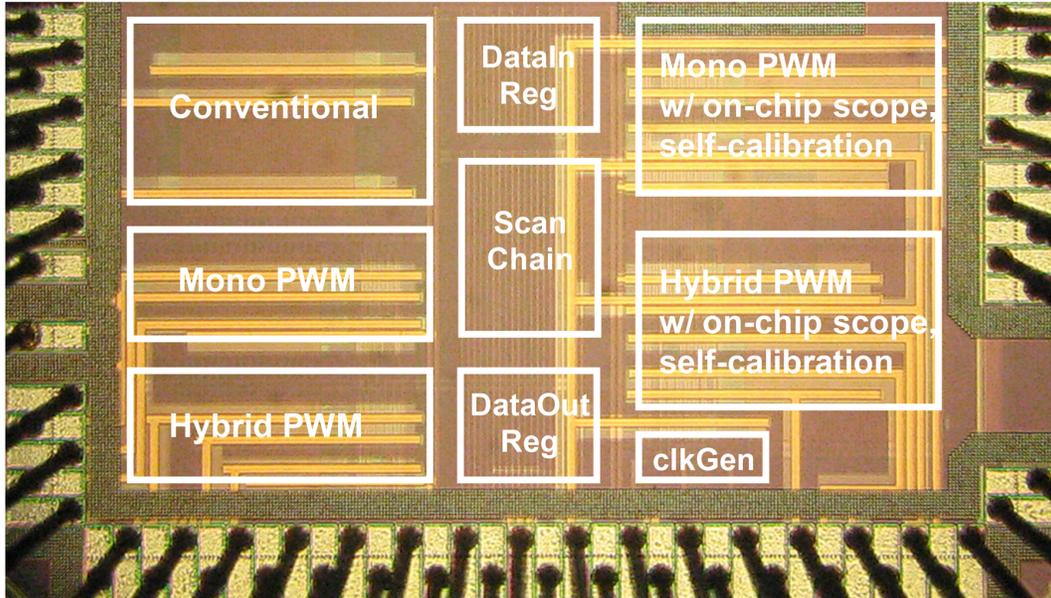
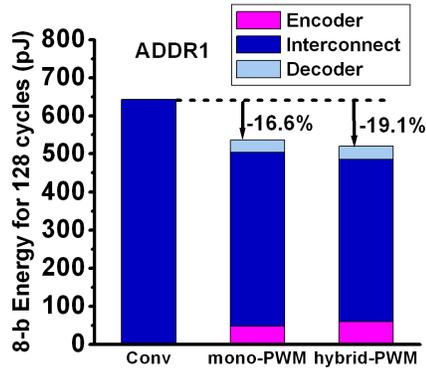


Figure 4.10: Die photograph (1.5mm X 0.7mm).

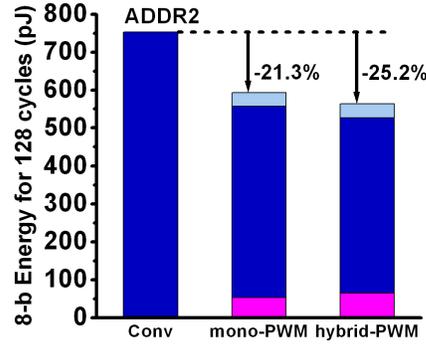
4.5 Measurement Results

8-bit 5mm links using the conventional, mono-PWM, and hybrid-PWM schemes were fabricated in 65nm CMOS (Figure 4.10). The conventional scheme used M5 minimum-pitch wires, and the proposed schemes used M5 2X minimum-pitch wires for identical routing area. The peak energy versus delay characteristics are shown in Figure 4.9. Despite the encoder/decoder overhead, the considerable reduction in wire parasitics results in both a 15% delay improvement and 46% peak energy reduction in the mono-PWM scheme. The improvements are smaller in the hybrid-PWM scheme due to the additional logic to enable single transition and a peak MCF of 2. For both schemes, the bit error rates measured by an on-chip bit error monitor are less than 10^{-13} .

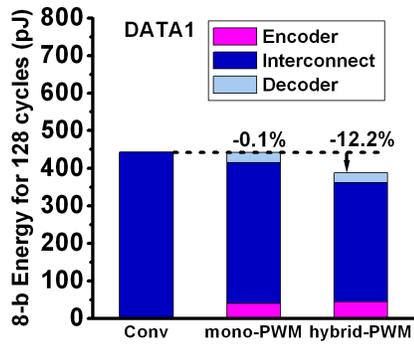
To measure average energy, traces from a microprocessor cache and memory bus were tested continuously through the in/out data registers. The data patterns were generated using the M5 full system simulator [47] for a GZIP workload from the Spec2000 benchmark suite. The results are shown in Figure 4.11, where up to 21% and 25% average energy reduction is achieved for the mono-PWM and hybrid-PWM scheme, respectively, compared to the conventional bus. The proposed schemes are observed to be more energy-efficient in address traces since adjacent address bits tend to switch simultaneously. Overall, the hybrid-PWM scheme exhibits better average energy savings than mono-PWM scheme due to single transition signaling in 33% of the switching



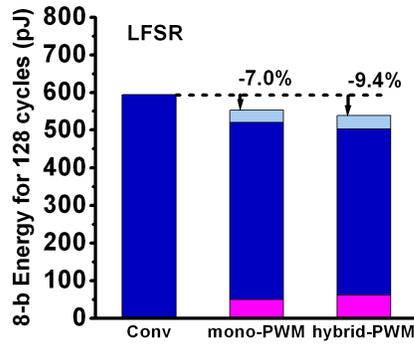
(a) Addr1 trace



(b) Addr2 trace



(c) Data trace



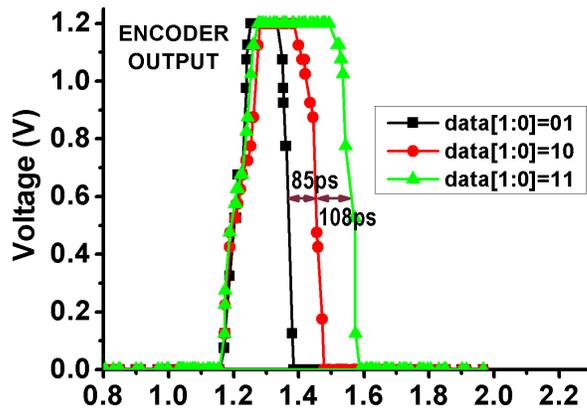
(d) LFSR trace

Figure 4.11: Average energy comparison of conventional, mono-PWM, and hybrid-PWM scheme for microprocessor address, data and LFSR traces.

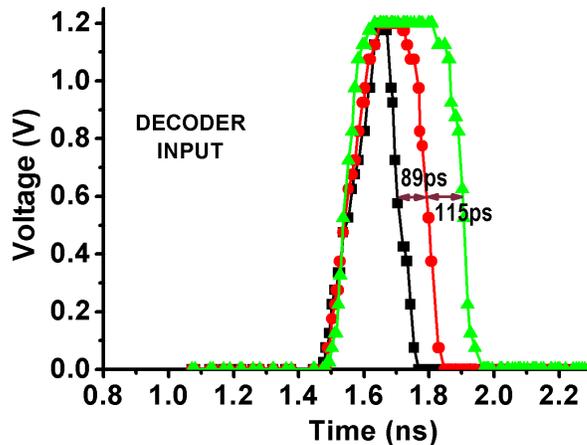
cases, which is most evident in the DATA1 trace of Figure 4.11.

An on-chip oscilloscope similar to the one proposed in [48] is implemented to capture waveforms at internal nodes. While sending a repetitive data pattern, the sample clock and reference voltage is finely swept to acquire each data point of the on-chip signal of interest. Figure 6.13 shows measured timing waveforms of the encoder output and decoder input in mono-PWM for different incoming data patterns. It can be seen that the pulse widths are well preserved across the 5mm link.

The effect of crosstalk-aware signaling is shown in Figure 4.13. The y-axis depicts the pulse



(a) Encoder output waveform.

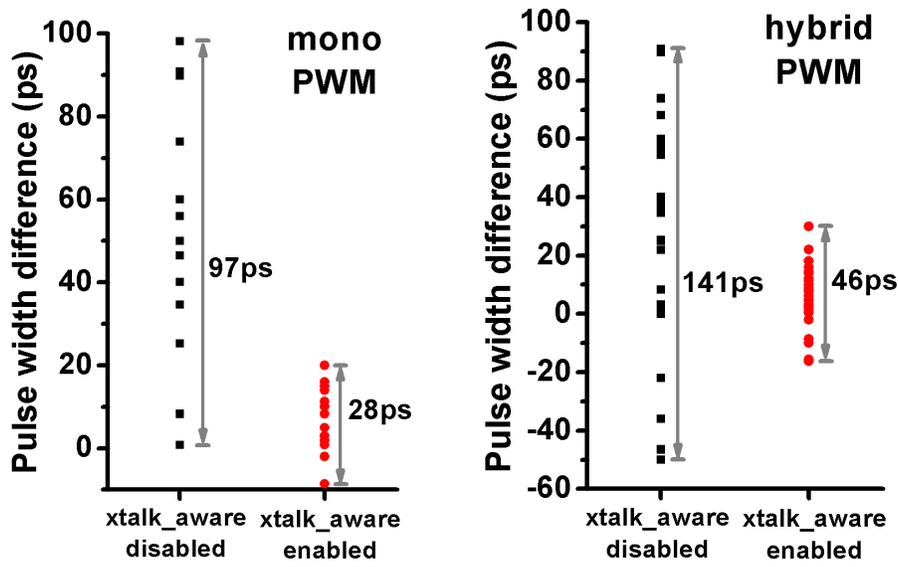


(b) Decoder input waveform.

Figure 4.12: Measured waveforms of mono-PWM using on-chip oscilloscope.

width change at the decoder input between the case when adjacent wires are idle and the case when adjacent wires are switching with different data. Crosstalk-aware signaling suppresses the effect of crosstalk on pulse width by 71%, enhancing the performance improvement of the proposed techniques. Leakage power is compared in Table 5.3, where the proposed schemes achieve $>2X$ reduction over conventional repeaters primarily because the number of optimal repeaters is significantly reduced due to both fewer wires and lower resistance and capacitance of the double pitch wires.

The low sensitivity of the proposed system to global supply and temperature variation is shown



(a) mono-PWM scheme.

(b) hybrid-PWM scheme.

Figure 4.13: Comparison of crosstalk-aware signaling on decoder input pulse with spread for all possible data patterns.

Table 4.1: Leakage power measurement (units: μW).

	Encoder	Wire	Decoder	Total
Conventional	-	35.2	-	35.2
Mono-PWM	3.4	8.5	2.3	14.2 (-60%)
Hybrid-PWM	4.6	8.3	2.7	15.6 (-56%)

in Figure 4.14(a), where ‘P’ represents pass with default control signals and ‘C’ represents pass with control signal adjustment. For variation experiments of mono-PWM, functionality is measured over 10^{11} bits. Functionality down to 700mV also demonstrates good robustness to process-induced mismatch, which is emphasized at low Vdd. To consider local supply voltage variation, different voltages are supplied to the encoder and decoder, and the contour plot of measured functionality at 40°C is shown in Figure 4.14(b). Calibrating the timing margin by altering the pulse widths trades off performance gains for robustness under supply variation.

The effectiveness of self-calibration is illustrated in Figure 4.15, where measurements from 21 chips are included. Before calibration, measurements show considerable spread of the mono-PWM scheme in the energy-delay space due to variation, and guard-banding on the worst chip would result in excessive performance penalty. Also, the proposed scheme has relatively larger spread comparing to the conventional scheme since the repeaters in the conventional scheme are significantly

larger than the variable inverter chain in the proposed scheme thus the transistor mismatch is fairly small. However, after self-calibration is completed, the energy-delay spread is tightened significantly, leading to $\sim 11\%$ performance improvement. Furthermore, the overall spread of proposed scheme in 21 chips is now less than that of the conventional scheme.

Figure 4.16 and 4.17 shows the delay spread of the mono-PWM and hybrid-PWM scheme, respectively, before and after self-calibration across 21 chips. It can be seen that self-calibration effectively mitigates variability and tightens the delay spread of the proposed schemes in 21 chips. The σ/μ of both schemes is reduced by more than 2.5X through self-calibration.

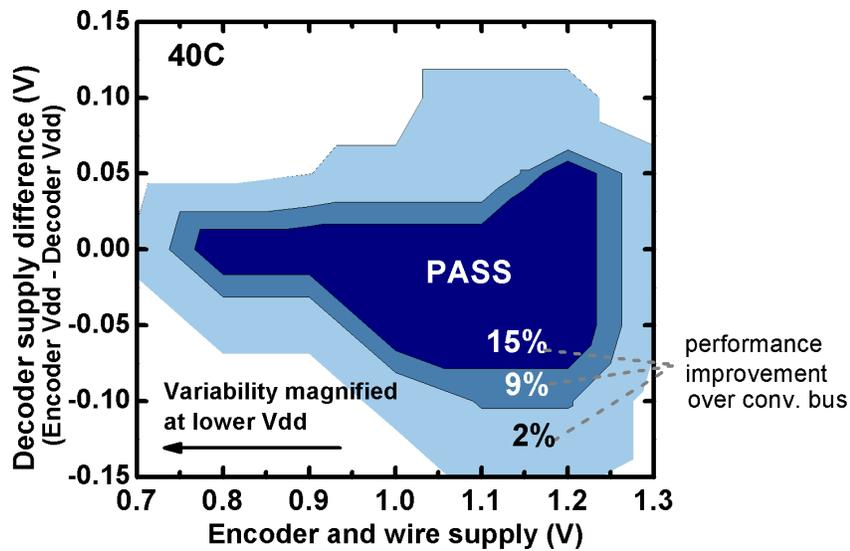
4.6 Summary

This chapter explored two crosstalk-aware signaling techniques based on pulse width modulation for energy- efficient on-chip global busses. Exploiting the fact that pulse widths are controllable in pulsed signaling, two bits of information are encoded into transition type and pulse width for transmission over one wire. For the same footprint, the wire width and spacing can be increased, leading to less wire resistance and capacitance, and thereby improving the overall performance and energy consumption of the bus system. The proposed encoder circuits dynamically pre-corrected the effect of crosstalk noise over long wires. Measurements from 5mm on-chip links in 65nm CMOS technology showed that the proposed schemes simultaneously achieve 15% performance improvement, 46% peak energy reduction, up to 25% average energy reduction, and $>2X$ leakage reduction compared to conventional repeaters. Comparing to previous works, these improvements relate to 25% additional energy improvement of [44] in similar technology, 2.5X higher bandwidth density than a transmission line based serialization approach [46], and 11% better performance than a current mode multi-bit signaling scheme [43] in the presence of crosstalk. Furthermore, self-calibration of the variable delay chains in the encoder and decoder reduced the delay spread of 21 chips by $>2.5X$.

80°C	C	P	P	P	P	P	P
70°C	C	P	P	P	P	P	P
60°C	C	P	P	P	P	P	P
50°C	C	P	P	P	P	P	P
40°C	C	P	P	P	P	P	P
30°C	C	P	P	P	P	P	P
20°C	C	P	P	P	P	P	P
10°C	C	P	P	P	P	P	P
0°C	C	P	P	P	P	P	P
	0.7V	0.8V	0.9V	1.0V	1.1V	1.2V	1.3V

Global (encoder/wire/decoder) supply

(a)



(b)

Figure 4.14: Global and local variation experiments are shown. (a) Sensitivity of mono-PWM system to global voltage and temperature variation. (b) Contour plot showing functionality and performance with local supply variation at 40°C. Additional guardbanding improves robustness at the expense of performance gains.

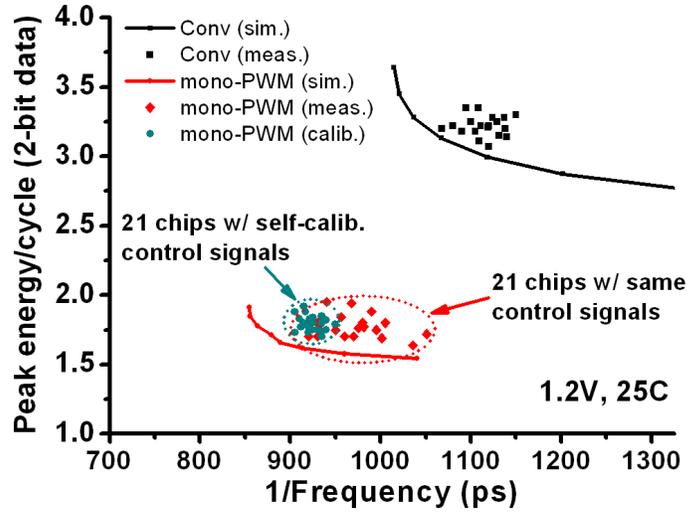


Figure 4.15: The performance and energy spread of 20 chips before and after self-calibration is shown.

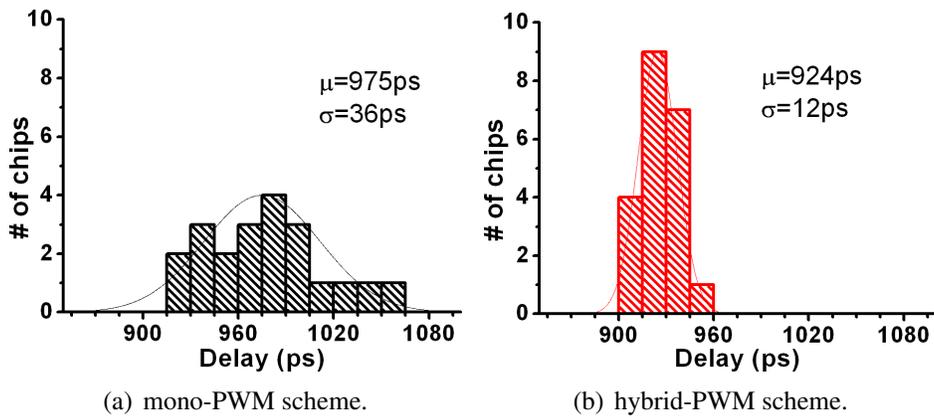


Figure 4.16: Comparison of mono-PWM scheme delay distribution with self-calibration. Self-calibration reduces σ/μ of 21 chips by 2.7X.

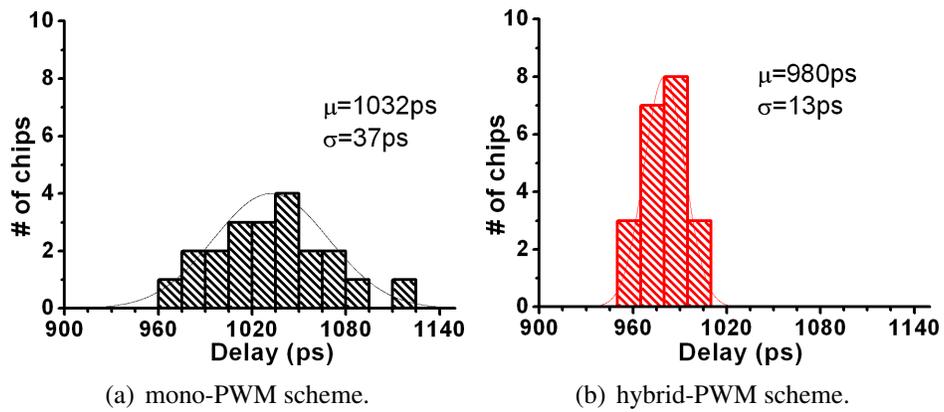


Figure 4.17: Comparison of hybrid-PWM scheme delay distribution with self-calibration. Self-calibration reduces σ/μ of 21 chips by 2.5X.

CHAPTER 5

Self-Timed Regenerators

In this chapter, we propose a new circuit technique called self-timed regenerator (STR) to improve both speed and power for on-chip global interconnects. The proposed circuits are placed along global wires to compensate the loss in resistive wires and to amplify the effect of wire inductance in the wires to enable transmission line like behavior. For different wire widths, the number of STR and sizing of the transistors are optimized to accelerate the signal propagation while consuming minimum power. In 90nm CMOS technology, STR design achieved a delay improvement of 14% over the conventional repeater design. Furthermore, 20% power reduction is achieved for iso-delay, and 8% delay improvement for iso-power compared with the repeater design. The proposed technique has also been applied to a clock distribution network, reducing clock power by 26%.

5.1 Consideration on Repeater-less Signaling

As CMOS technology scaling continues, the number of repeaters increases dramatically. Deep submicron projections in [65] show that global interconnect distribution (repeaters + wires) will consume $\sim 40\%$ of the total power in 50nm technology. Recently, a microprocessor design [51] reported using as many as 12,900 repeaters showing that power and area overhead due to repeaters is becoming a serious concern.

A number of methods to address interconnect issues without using repeaters has been proposed [52–54]. The first design uses so-called boosters [52] where extra current is supplied when a

transition is detected. However, the fact that it has a stack of two transistors in the charge path limits the speed improvement. In [53] a method is proposed where the receiver biases the voltage at which a transition is detected based on the expected transition direction. [54] proposed a capacitive coupling accelerator, similar to a booster, to reduce RC delay, but the improvement over repeater design was not significant. Reducing the voltage swing has been used [55] to improve power, but is problematic in that another power rail has to be present and the delay tradeoff is not highly favorable. Finally, alternative approaches include modulated signaling [57] and pulsed current-mode signaling [58]. These methods achieved near speed-of-light latency, but they require wide wire topologies with low loss characteristics and the complexity of these designs makes them difficult to adopt in the industry.

In this work, we present a new circuit technique [59,60] to achieve high performance, repeater-less propagation for global interconnects. The proposed design was implemented and tested for a number of interconnect structures, and applied to a clock network design.

5.2 Self-timed Regenerator Design

5.2.1 Transmission Line Configuration

We first consider a lossless transmission line where the driver is perfectly matched with the line impedance and the receiver is sufficiently small to present a negligible load. When the driver input transitions, a wave of $V_{DD}/2$ will be propagated along the transmission line, due to the matching of the driver impedance. When the propagated wave reaches the receiver, this voltage will be doubled to the full rail due to the light loading of the receiver, as shown in Figure 1(a). Note that while a reflected wave is sent back through the transmission line, this reflected wave will be absorbed completely by the driver since it is perfectly matched. This configuration is particularly advantageous for point-to-point signaling in VLSI designs, for instance between the processor and the cache. Taking advantage of reflection at the receiver termination to obtain a full swing signal allows the new design to be easily incorporated in existing design methodologies.

For a sufficiently wide wire, the resistance is insignificant and we obtain behavior similar to that of an ideal transmission line, as shown in Figure 1(b). However, when the wire becomes

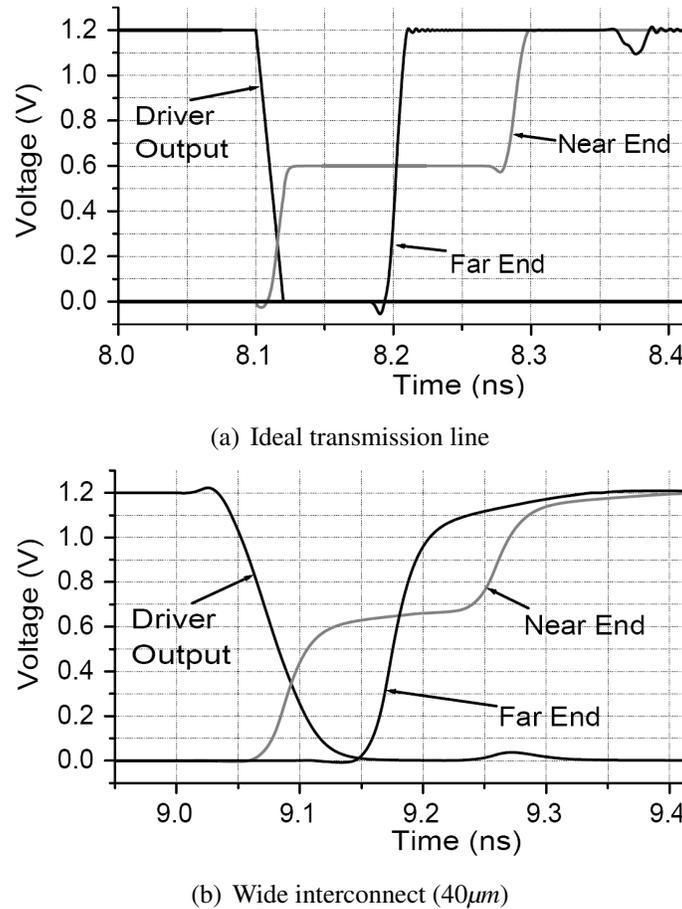


Figure 5.1: Interconnect with transmission line behavior.

thinner, the resistance becomes significant and the signal attenuates as it propagates through the wire. In this case, the signal swing at the receiver may not be sufficient to detect the transition reliably and signal propagation speed is also degraded. To compensate for this signal degradation, our design enhances the transition by properly supplying additional current, while still utilizing the impedance matching and the receiver reflection.

5.2.2 Self-Timed Regenerator (STR) Circuit Operation

The self-timed regenerator (STR) is designed to quickly detect and accelerate the transition for a certain amount of time with a certain amount of current from the rail. Figure 5.2 shows the self-timed regenerator (STR) circuit on both sides of the interconnect. The upper and lower part accelerates the rising and falling transition, respectively, each one being complementary to the other.

The main idea is to generate a pulse at node B and C which would turn on P3 and N6 for a time equal to the width of the pulse. When transistors P3 and N6 are turned on, additional current is supplied from the power rail to the propagating signal to expedite the transition. Transistors N1 and P4 are low threshold transistors which turn on quickly according to the polarity of the signal. P1 and N4 are weak transistors which are present only to establish and maintain initial conditions at nodes B and C. The delay set by the odd-number inverter chain determines the width of the pulse. The number and size of inverters in the chain can be optimized for different wires and constraints. This enables self-timing of the pulse width.

The initial state of the internal nodes of the circuit should be known. When the signal line is at low-voltage steady state, transistor P1 is driving node B to V_{DD} . Node D is also set at V_{DD} , making the pull-up circuit just ready to detect low-to-high transition while lower circuit remains insensitive to any rising transition. If any noise pulls node D down to GND , P4 and P5 charges node C and after going through a chain of inverters, P6 actively drives node D back to V_{DD} , which is the desired initial condition. Similarly, when the signal line is at high-voltage steady state, node C and D is set at GND .

Upon a transition, the circuit works as follows and the timing diagram for each transition is shown at Figure 5.3. Note that node A is the interconnect line itself. When the wire is initially at GND , the next transition will be a rising transition. Since transistor P5 is off, the pull-down circuit would be insensitive to this transition. When a rising transition is detected by N1, node B is pulled down to GND immediately as can be seen in Figure 3(a). P3 turns on and it enhances the transition of the signal. After some delay, N3 turns on and node D is grounded, also shown in Figure 3(a). P2 charges node B to V_{DD} , so P3 turns off. After some time, N3 turns off but node D is maintained at GND by the cross-coupled inverters. Now, the pull-up circuit becomes insensitive to the high-to-low transition as N2 is turned off. In case of a falling transition, the timing waveforms of the internal nodes are shown in Figure 3(b). Figure 5.3 also compares the waveforms between the case when we use low V_t transistors for N1 and P4 and when we do not. Considerable performance improvement is observed by using 2 low V_t transistors in the STR.

A key feature of our design is that transistors (P2, N1, N2) and (P5, P4, N5) are never turned on simultaneously. This eliminates the fight during a transition which degrades the performance improvement and results in additional short circuit current, resulting in power reduction. In addi-

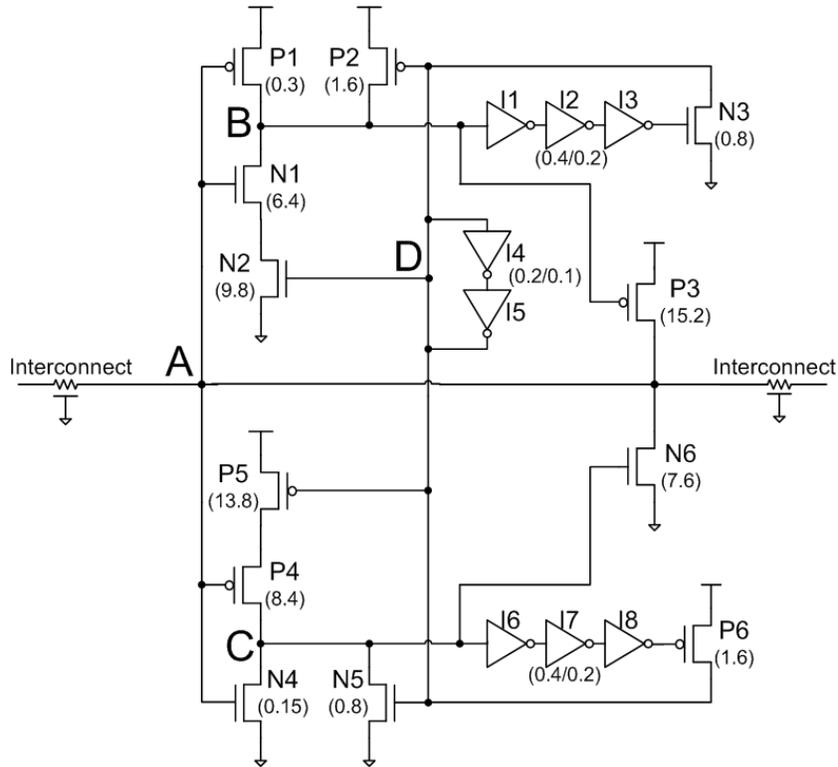
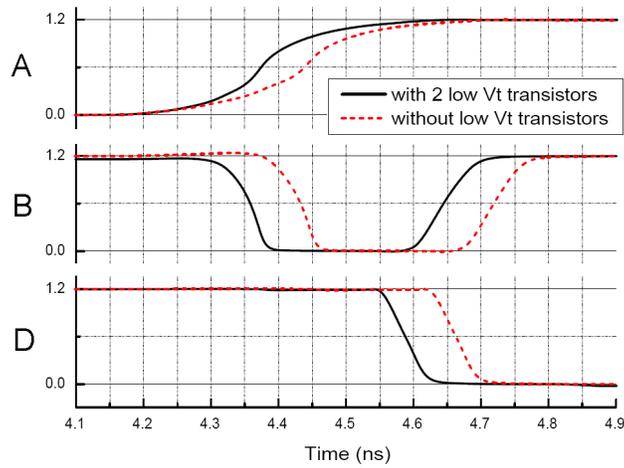


Figure 5.2: Self-timed regenerator circuit. Optimal sizing(unit: μm) for power reduction when 5 STRs are placed for a $0.45\mu m$ wire is shown.

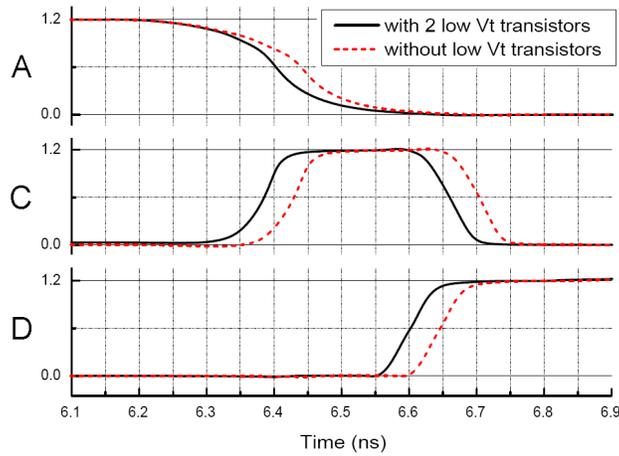
tion, the signal line is accelerated by a single transistor between the signal line and the supply rail allowing for a high drive current.

5.2.3 Sizing of the Circuit

Sizing of the STR circuit should be done carefully to facilitate the desired operation. An example of optimal sizing of STR is shown in Figure 5.2. Transistors N1, N2, P3 and P4, P5, N6 are larger transistors than the others. As the number of STRs placed on a wire increases, their sizes get reduced. The size of transistors P3 and N6 determines the amount of current supplied to the signal line. Sizes of N1, N2 and P4, P5 determine the response time of the circuit to the propagating wave. The faster these transistors are, the more quickly transistors P3, N6 get triggered and the better output waveform is at the far end of the wire. The rest of the transistors are not critical in terms of speed and are sized relatively smaller than these six transistors to minimize power consumption. N3 and P6 should be strong enough to switch the state of cross-coupled inverters.



(a) Low-to-high transition



(b) High-to-low transition

Figure 5.3: Timing diagrams of STR at rising and falling transition. Speedup due to 2 low V_t transistors is shown.

Sizes of P2 and N5 determine the slope of the trailing edge of the pulse. The sizing of STR is optimized using HSPICE optimizer and manual sweep for every combination of wire width and number of STRs placed on the line, resulting in different sizing for each situation.

Presence of handful transistors seems to pose STR as a timing critical circuit for sizing. However, other transistors besides the main 6 transistors are relatively not sensitive to sizing. In addition, depending on the technology, several optimal ratios can be imposed among the 6 critical transistors (i.e., $P3/N6=2$, $N2/N1=P5/P4=1.5$), which accelerates sizing.

Table 5.1: STR power and performance comparison

Opt. Repeaters: optimal number of repeaters, Opt. Regen: optimal number of STRs

Wire width	Opt. Repeaters	Power reduction (Iso-delay)	Opt. Regen	Delay improvement (Iso-power)	Opt. Regen	Delay improvement (Best case)	Opt. Regen
$0.3\mu m$	10	19.8%	6	7.7%	9	14.0%	13
$0.45\mu m$	8	17.6%	5	6.8%	6	13.8%	12
$1\mu m$	5	16.9%	3	7.4%	5	13.3%	10
$4\mu m$	2	10.9%	1	8.1%	1	11.8%	9

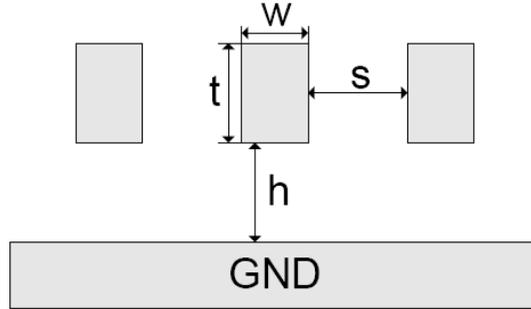


Figure 5.4: Structure of global interconnect.

5.2.4 Effect of STR on Signal Integrity

STRs create discontinuities in the transmission line as additional capacitive and resistive load. This causes the reflection of the original signal at positions where STRs are placed. In narrow wires, the dominant resistance of the line effectively suppresses these reflections. Although the resistance is less in wide wires, STRs are placed further apart, resulting in less reflections. Therefore, the extra ISI generated by STRs are largely insignificant.

Another issue which needs to be addressed is inductive coupling. As STR is intended for long wires it increases the current loop thereby resulting in larger inductance, but this is compensated by the reduction in peak current in the STR scheme, which is shown in Section 3. Furthermore, global wires or clock nets, which are the applications of STRs, are typically shielded with dedicated wires which provide well-defined return paths, thereby reducing inductive coupling.

5.3 Experimental Results

The $90nm$ technology results are obtained from SPICE simulations for a broad range of wire widths using industrial device models. For simulation of different wire widths, width w in Figure

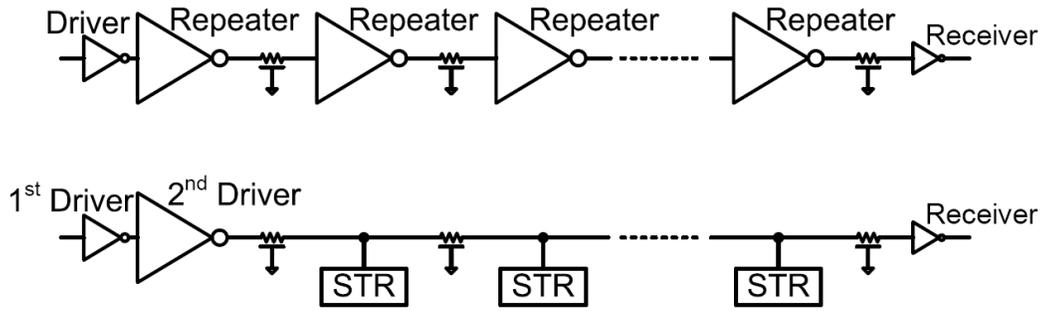


Figure 5.5: Repeater/STR implementation scheme.

5.4 is changed and all the other parameters such as spacing, thickness, and distance from the ground plane are fixed. We modeled the interconnect as a top global layer metal with shielding wires on either side. Using a distributed RLC pi model, and each $25\mu\text{m}$ of the wire is modeled as a lumped RLC circuit, and the RLC values have been extracted using FastHenry [61] and Predictive Technology Model [62] for a 10mm line, and the line length is fixed throughout this chapter.

We also compared the proposed STR and repeater technique against a traditional booster design [52]. However, the booster design was not able to improve on the performance of the repeater design even after extensive optimization of the transistor sizing of the booster topology. Hence, no comparison of our proposed approach against the booster design is given since the gains will be more than that compared to repeater designs. Previous reported measurements for the booster design were performed by measuring delay from the output of the inverter driving the interconnect to the input of the receiver gate. This method of delay measurement ignores the delay due to the loading of the initial driver and hence might not be accurate.

5.3.1 Repeater and STR Design Scheme

The overall scheme for repeater and STR is shown in Figure 5.5. To make the comparison fair, we have identical initial drivers both in the repeater and STR scheme. In the repeater scheme, the first repeater is placed after the initial driver. In the STR scheme, the initial driver is followed by a stronger driver, which is sized properly to match the impedance of the line. Adding STR changes the load of the line which affects the impedance, but the difference is not significant. For example,

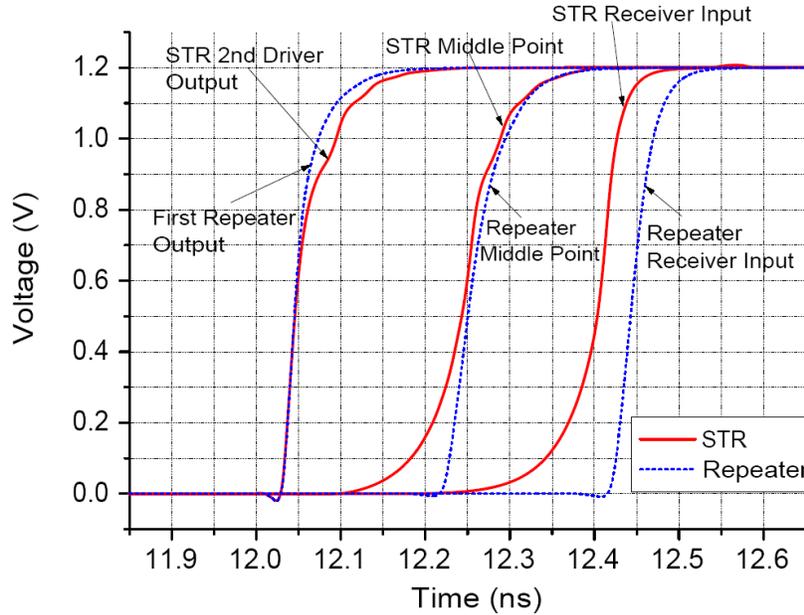


Figure 5.6: STR and repeater simulation waveforms of $0.3\mu\text{m}$ wide interconnect.

the impedance of $1\mu\text{m}$ wide wire when unloaded and loaded with STRs are 46.6Ω and 44.1Ω , respectively. As a result, the optimum size of this 2nd driver increases as the wire becomes wider. Throughout the interconnect, repeaters and STRs are inserted regularly. The delay is measured from the initial driver input to the final receiver output.

At each wire width, both the number and size of repeaters are swept to achieve best performance. Among all the combinations, optimal number and size of repeaters which result in the best case delay is chosen, and this serves as the baseline of comparison for the STR designs. Similarly, for a given wire geometry, both the sizing of STR and number of STRs along the interconnect is varied and optimized to achieve better energy and delay compared to the repeater scheme. For simplicity, all repeaters and STRs are sized identically.

Figure 5.6 shows the waveform of intermediate nodes of $0.3\mu\text{m}$ wide wire simulation for STRs and repeaters. For this lossy wire, we see significant delay improvement using STRs compared to the repeater design. The waveform of a $4\mu\text{m}$ wide interconnect is shown in Figure 5.7. Since the resistance of the wire is now reduced significantly, we start to see the transmission line effects with reflections at the far end. This results in faster transition time at the output of the interconnect than an intermediate point of the wire. The overall power and performance comparison is summarized in Table 7.1.

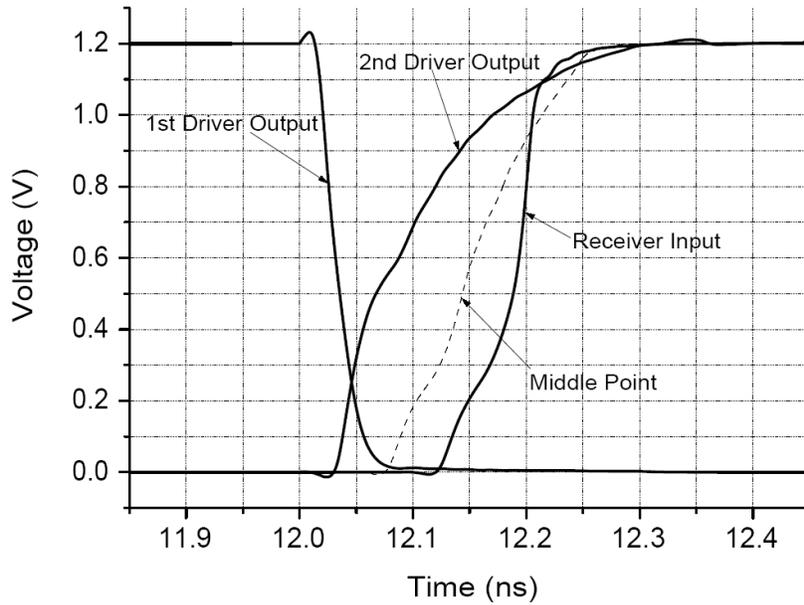
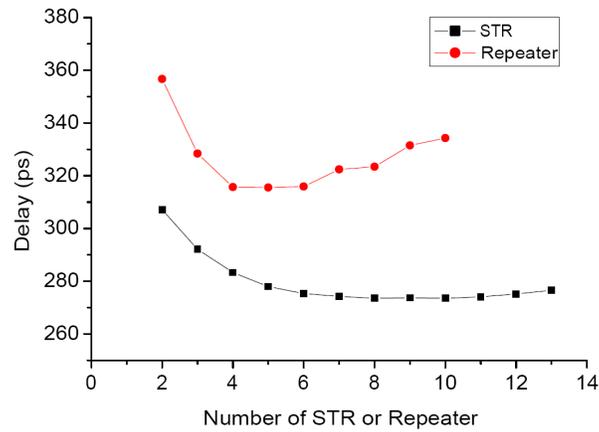


Figure 5.7: STR simulation waveform of $4\mu\text{m}$ wide interconnect.

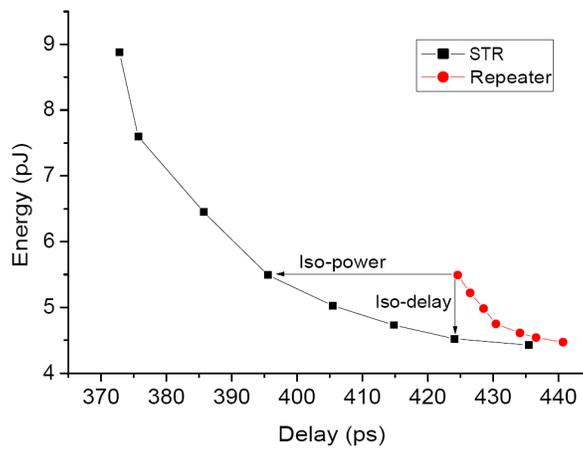
5.3.2 Power, Area, and Peak Current

To measure the amount of power savings we achieve with STRs, the power comparison of the two designs is performed with the same delay. For iso-delay, power reduction up to 19.8% is achieved in the STR design as shown in Table 7.1. This is first due to the fact that smaller numbers of STRs are needed than that of repeaters at the same delay constraint. Also, the short-circuit current is minimized in the STR design because there is no strongly conducting direct path from V_{DD} to GND at any given time. Furthermore, the STR circuit need not be oversized to produce equivalent delay with repeaters. Therefore, the total device width (including drivers, receivers, STRs and repeaters) is reduced significantly as shown in Table 5.2. In Table 7.1, we observe that power savings of STRs comparing to the repeaters decreases as the wire width increases. This is because the capacitance dominates the interconnect parasitics in wide wires, and therefore sizing down the STR cannot reduce the power dissipated by the capacitance by a large amount.

Peak current is an important metric in interconnect design since it determines the value of decoupling capacitance which suppresses noise and voltage droop. The maximum current is measured in each repeater and STR block for peak current. The results are shown in Table 5.2 for iso-delay case, and substantial savings in peak current are achieved. This is first due to fact that a direct path from V_{DD} and GND does not exist at any given time and secondly the transistor widths



(a)



(b)

Figure 5.8: (a) Delay comparison with different numbers of STRs and repeaters (width : $1\mu m$). Sizing is optimized for each different number of STR and repeater. (b) Energy vs. Delay of STR and repeater (width : $0.45\mu m$).

in STR are considerably smaller than those in repeater scheme. Also whenever a transition is detected, the fact that the transition is accelerated for a self-timed window explains almost invariant peak current values across different wire widths.

5.3.3 Performance

For performance comparison, the power consumption of both STR and repeater design is set to be the same, and the delay is measured in each case. Across different wire widths, the maximum delay improvement is 8.1% for iso-power. Furthermore, we obtained the maximum performance

Table 5.2: Area and peak current comparison (iso-delay)

Wire width	Scheme	Total device width	Peak current
$0.3\mu m$	Repeater	$738\mu m$	$7.1mA$
	STR	$451\mu m$ (-39%)	$3.1mA$ (-56%)
$1\mu m$	Repeater	$543\mu m$	$9.9mA$
	STR	$317\mu m$ (-42%)	$3.9mA$ (-61%)
$4\mu m$	Repeater	$282\mu m$	$12.5mA$
	STR	$130\mu m$ (-54%)	$3.5mA$ (-72%)

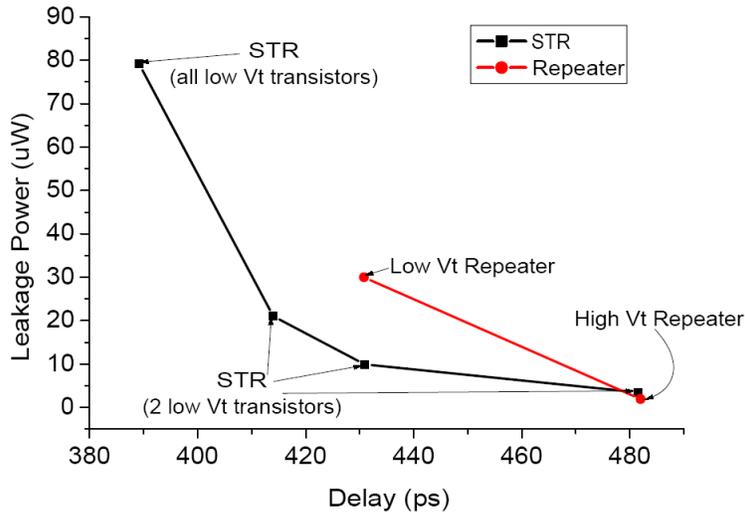


Figure 5.9: Leakage power with different Vt assignments (width : $0.3\mu m$).

with the STR design when performance has a higher priority than power consumption. Figure 8(a) shows that the performance of the STR design dominates that of the repeater design. Performance improvement up to 14% is achieved, and it slightly decreases for wider wires.

Figure 8(b) shows energy vs. delay for STRs and repeaters. The data points in this plot are the minimum energy points obtainable with the given delay for STRs and repeaters. We can see that the STR energy-delay curve exists in the left-bottom side than that of the repeater.

5.3.4 Low Vt Repeaters and Leakage power

Since we are exploiting 2 low Vt transistors in STRs, we also optimized repeaters with low Vt devices to achieve a comprehensive comparison. Figure 5.9 shows the leakage power comparison with STR using two low Vt devices, repeaters with high Vt and low Vt. The data points in Figure

Table 5.3: Repeater and STR leakage comparison

(i) : Iso-delay with low Vt repeater (ii) : Iso-delay with high Vt repeater

Wire width (μm)	Repeater	Leakage (μW)	STR	Leakage (μW)
0.45	all lvt	28.01	(i)	11.70
	all hvt	1.86	(ii)	4.76
1	all lvt	22.81	(i)	14.37
	all hvt	1.78	(ii)	3.92

Table 5.4: Clock distribution network comparison (W1/W2=2)

Scheme	W1	Power	Skew	Delay	Slope
No STR	$8\mu m$	$19.1mW$	$11.8ps$	$130ps$	$51ps$
With STR	$5\mu m$	$14.1mW$	$8.9ps$	$140ps$	$59ps$

5.9 for repeaters are the best case delay for each scheme. It is shown that although low Vt devices are used for all repeaters, it cannot reach the speed of STR, which has only 2 low Vt transistors. In a $0.3\mu m$ wide wire, the leakage power of STRs is 3X lower than that of low Vt repeaters for iso-delay. Since the total area in the STR scheme is only 50~60% of that in the repeater scheme for iso-delay, leakage power is reduced although there are more transistors in STR comparing to a repeater. Also, these results show that our design has a very specific critical path so that we gain considerable speed improvement by using a few low Vt devices without sacrificing leakage power. As the wire becomes wider and the capacitance dominates the interconnect parasitics, leakage power improvement over the repeater design diminishes as shown in Table 5.3.

5.4 Clock Network Application

One application of the STR is clock network design. To minimize skew and delay, global clock wires are typically very wide. This results in large wire capacitance and large clock drivers as well, consuming a considerable portion of the entire chip power. To reduce the width of the wires, we propose to place STRs along the clock wire to improve clock delay and clock skew, while meeting the same constraint as in the conventional case with wide wires.

Figure 5.10 shows the spine clock distribution network without and with STR. The clock speed is set at $1GHz$, and the objective is to keep the maximum clock skew less than 15ps between any

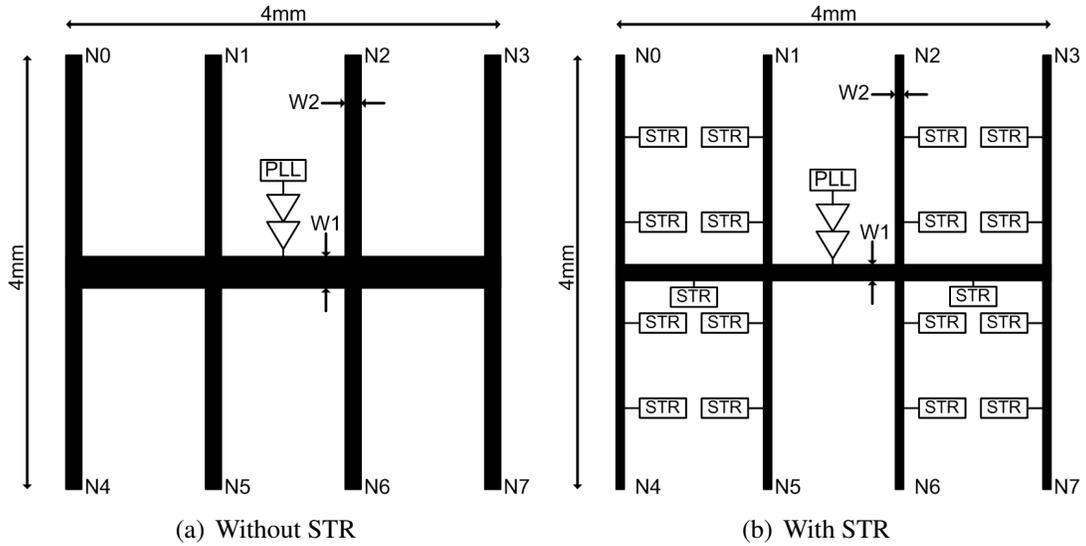


Figure 5.10: Spine clock distribution network configuration ($W1/W2=2$).

two nodes out of the 8 leaf nodes $N0\sim N7$, the clock delay less than 150ps, and the clock slope less than 70ps. The clock delay is defined as the worst delay from the driver input to a leaf node, and the slope is defined as the worst delay between the 20% and 80% point of the clock signal at a leaf node. Without using any STR in the clock network, we had to use $8\mu m$ for $W1$ and $4\mu m$ for $W2$ to meet the given constraint. When STRs are added approximately every $0.8mm$ along the clock path, the $W1$ and $W2$ could be reduced down to $5\mu m$ and $2.5\mu m$, respectively, while maintaining similar clock skew, delay and slope. Considering that the distance of clock net to multiple destinations cannot be perfectly matched, especially in spine clock networks, adding STR per unit length as suggested can compensate for the inherent skew between different nodes. The comparison results of the two clock distribution networks are shown in Table 5.4. By reducing the clock width substantially, the clock driver need not be sized as much as the conventional case to keep the slope similar, and the capacitance of the clock wires decreased significantly as well. In result, clock power consumption is reduced by 26.2%.

Although the transistor count is more in the STR scheme, under PVT variations, the spread of skew, delay and slope were found to be actually less than that in the scheme without STR. This would be due to the averaging effect in the presence of more transistors.

5.5 Summary

In this chapter, we presented a new circuit technique to improve delay and save power for global interconnects. For a $10mm$ wire, we could achieve up to 20% power reduction, 2X peak current reduction, and 3X leakage power reduction than the repeaters. Applying STRs in a spine clock distribution network reduces the width of the clock wires considerably while meeting the same skew and delay constraints as in the conventional case, resulting in 26% power reduction.

CHAPTER 6

High Bandwidth Low Swing Signaling

This chapter presents new circuit techniques for high bandwidth and low energy repeaterless on-chip communication. The transmitter generates pre-emphasized bipolar signals through series capacitors, resulting in significant reduction of inter-symbol interference (ISI) and considerable bandwidth improvement in RC-dominated global wires. The hysteresis receiver efficiently recovers non-return-to-zero (NRZ) data from fast return-to-zero (RZ) pulses. Employing double data rate (DDR) signaling, the slopes of RZ pulses are adaptively controlled such that the bandwidth is effectively doubled only when required. Measurement results from a 90nm CMOS prototype chip show up to 4.9Gb/s operation with 0.34pJ/b energy consumption for a $0.28\mu\text{m}$ wide, 5mm long interconnect.

6.1 Motivation and Previous Work on Low Swing Signaling

Long on-chip wires pose well-known latency, bandwidth, and energy challenges to the designers of high-performance VLSI systems. Repeaters effectively mitigate wire RC effects but do little to improve their energy costs. Moreover, proliferating repeater farms add significant complexity to full-chip integration, motivating circuits to improve wire performance and energy while reducing the number of repeaters.

For repeaterless on-chip communication, low swing signaling schemes have been used to improve the energy efficiency of global on-chip wires while sacrificing noise margin and stability. While maintaining low energy consumption, there have been a number of approaches and methods

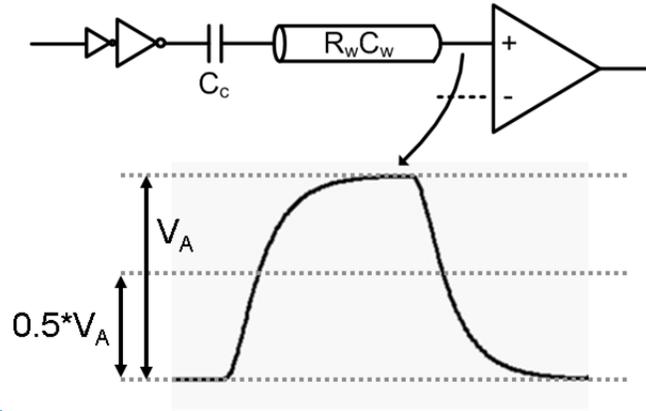
to also enable high-performance for long wires.

Such methods include capacitive-mode signaling, which combines a capacitive driver with a capacitive load [82, 85, 86]; and current-mode signaling, which pairs a resistive driver with a resistive load [87, 88].

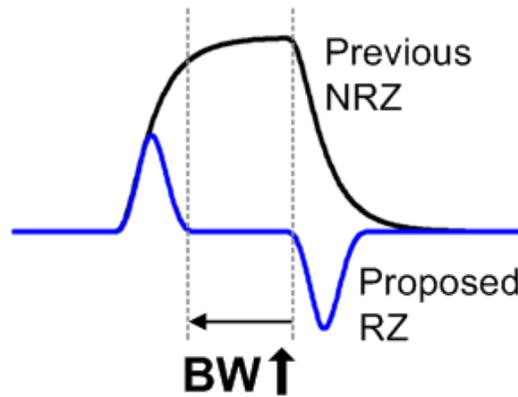
A single-ended voltage-mode pre-emphasis technique was introduced by Zhang [82] offering a performance boost by the transmitter, but it uses a large wiring pitch of $9\mu\text{m}$. Schinkel proposed a pulse-width pre-emphasis technique to reduce intersymbol interference (ISI) and improve data rate with a wiring pitch of $1.6\mu\text{m}$ [83]. However, its energy consumption does not scale with data activity.

Driving a wire capacitively [85,86] increases on-chip wire bandwidth by capacitive pre-emphasis and employs low-swing signaling without a second supply. Also, the improvements are achieved with a relatively small wiring pitch ($1.2\mu\text{m}$ in 180nm [85], $1.72\mu\text{m}$ in 90nm [86]). However, the latency of these approaches do not reach that of optimally repeated wires in scaled technology nodes with narrow wires, and slow slew rates on highly resistive interconnects will still limit wire performance due to ISI. Recently, Bae [89] claimed an additional bandwidth improvement by the combination of a shorting resistor between the differential wires and negative feedback at the receiver end. However, transistor mismatch between the transmitter and receiver will impair reliable RZ signaling, and the continuous negative feedback reduces effective eye-height at the receiver, requiring a larger series capacitor to drive the wire. Furthermore, a clocked comparator without hysteresis, as proposed by Bae, is not a good receiver for RZ signals because the differential inputs are at a identical voltage in case of no data activity, and then it can easily evaluate into the wrong output value. Further improvement can come from equalization circuits on receivers [86] and transmitters [88] that trade off power for bandwidth.

In this work, we extend these ideas to a capacitively driven pulse-mode wire using a transmit-side adaptive FIR filter and a clockless receiver [90]. A 3-tap transmitter generates bipolar signals to eliminate ISI along the wire, and as a result, complete RZ signals arrive at the receiver end, allowing the use of a simple hysteresis receiver. With a global clocking scheme, 2.5Gb/s signaling costs 0.24pJ/b energy consumption, improving the latency and bandwidth of prior work. If a double data rate (DDR) scheme is applied for the identical transceiver, the system works with data rate up to 4.9Gb/s while only consuming 0.34pJ/b . These results are achieved at the expense of additional



(a) NRZ signaling of [85, 86] at receiver input



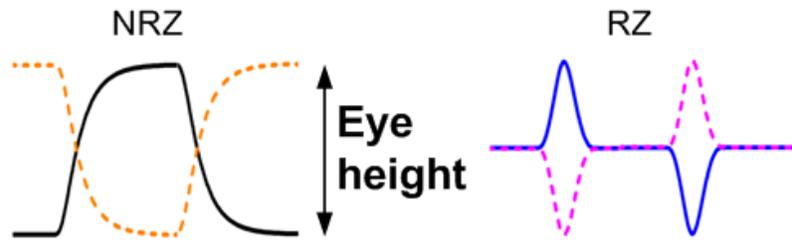
(b) Proposed RZ signaling at receiver input

Figure 6.1: Receiver input signal waveform of [85, 86] and proposed scheme. Fast RZ signaling leads to bandwidth improvement.

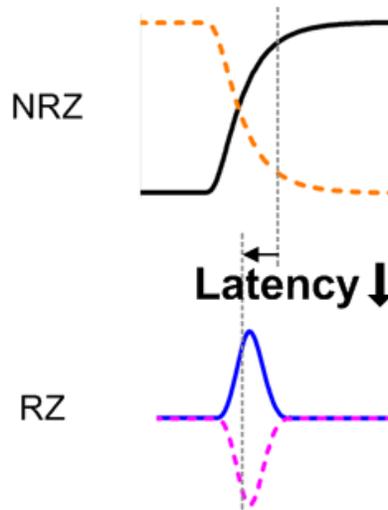
transmitter energy, but the overall energy is still 2X lower than an optimal repeater design with the same footprint. Compared to previous capacitive-mode signaling, the receiver eye height is unchanged and the wire energy is comparable.

6.2 Concept and Advantages of the Proposed RZ signaling

Figure 6.1(a) shows the receiver input waveforms of the previous work [85, 86], where V_a denotes the full swing at the receiver end. We observe that the signal makes a fast transition to $0.5 * V_a$, but then slowly saturates to full V_a . After utilizing the fast first half of the transition, if we can generate a fast RZ signal as seen in Figure 6.1(b), up to 2X bandwidth improvement is



(a) Eye height comparison.



(b) Latency comparison.

Figure 6.2: Eye height and latency comparison of previous NRZ scheme and proposed RZ scheme.

possible. We refer to this basic concept as a single-data-rate (SDR) scheme to distinguish it from a double-data-rate (DDR) scheme described later.

Although we halved the signal swing for one of the differential signals, this does not mean a reduction in eye height. In Figure 6.2(a), we show that RZ signaling has a common value in case of no data activity, and each differential signal expands in the opposite direction, resulting in an equivalent eye height with the one in NRZ signaling. This is important, because pushing towards the minimum detectable swing at the receiver would require more offset compensation circuitry and complicate receiver designs, leading to more energy consumption.

Sending data using differential RZ pulses improves wire latency, as the wires start each transition already equalized. A timing comparison of differential receiver end signals of NRZ and RZ schemes is shown in Figure 6.2(b). The time to reach minimum separation between the differential signals is less in RZ scheme, leading to improved latency. When the NRZ scheme reaches half

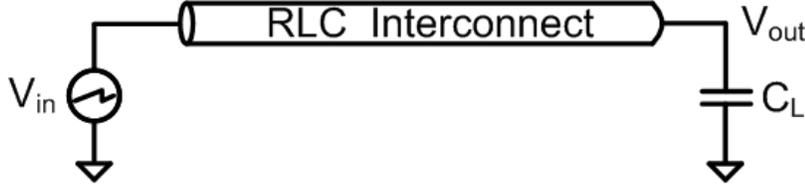
of the swing, the two differential signals are at the same voltage level and are definitely not ready to evaluate the correct output. However, at this point, the differential signals of RZ scheme are already separated by the maximum amount, evaluating new output values. For reliable detection of the correct values, the signals in the NRZ scheme require sufficient separation, which is particularly slow after passing the half swing point. Therefore, the RZ latency of the long wire is improved considerably while providing the same eye-height at the receiver. More detailed receiver operation is discussed in Section 6.4.2.

RZ signaling combined with capacitive coupling has also been introduced in off-chip communication [91–94] for better performance and lower power. In chip-to-chip communication, the communication channel is usually a transmission line, so it is more suitable to use proper termination resistance to bring the capacitively coupled signal back to a common voltage level. For narrow on-chip wires, however, a simple RZ pulse will smear out as it propagates through a resistive wire, limiting the achievable data rate. Capacitive pre-emphasis [85,86] improves the bandwidth to some degree, but slow transition to full swing past the mid point is a bottleneck. Therefore, to extend the bandwidth further, we propose using bipolar signals generated by the transmitter in order to explicitly eliminate ISI and to receive fast RZ pulses at the end of the wire. To verify the precise shape of waveforms required at the transmitter output or the beginning of the wire, we conducted simple Matlab experiments as described in the following section.

6.3 Matlab Analysis

Because a proper sequence of bipolar signals can eliminate ISI, the objective of the Matlab experiments is to (1) find the effect of transmitted amplitude and pulse width on the quality of RZ pulses at the end of the wire and (2) discover the corresponding waveform at the transmitter output. First of all, we used a distributed interconnect model as shown in Figure 6.3 and used a 2nd order approximation of the transfer function to represent the long wire. For the wire resistance and capacitance, we used extracted parasitics from a layout of 5mm long M5 wire with 0.28 μ m width and 0.28 μ m spacing in TSMC's 90nm CMOS technology. M4 and M6 layers were filled with ground planes to mimic densely packed perpendicular interconnects.

In the first experiment, following a fixed positive portion of a pulse, various negative pulses



$$H(s) = \frac{1}{(1 + C_L s) \cosh \theta + (Z_o C_L s) \sinh \theta}$$

$$\approx \frac{1}{1 + b_1 s + b_2 s^2}$$

$$h(t) = \frac{1}{b_2 (s_1 - s_2)} (e^{s_1 t} - e^{s_2 t})$$

$$\theta = \sqrt{(r + sl)sc}$$

$$Z_o = \sqrt{\frac{R + sL}{sC}}$$

$$b_1 = \frac{RC}{2} + RC_L$$

$$b_2 = \frac{(RC)^2}{24} + \frac{R^2 CC_L}{6} + \frac{LC}{2} + LC_L$$

Figure 6.3: Distributed interconnect model and 2nd order approximation of the transfer function.

with different pulse widths and amplitudes were sent through the interconnect as illustrated in Figure 6.4. The negative area was kept identical with positive area in order to bring the output signal back to common level. For given input signals, the output waveform is generated using the following equation.

$$V_{out}(t) = IFFT[FFT(V_{in}(t)) \times FFT(h(t))] \quad (6.1)$$

Figure 6.4 shows that the negative pulse with largest amplitude and shortest pulse width eliminates ISI most effectively. Intuitively, symmetry implies that the negative spike which has the same amplitude and pulse width as the positive spike generates a sharp RZ pulse at the far-end of the wire.

For a given output signal, the second experiment finds the waveform of the required input signal using an inverse FFT in Matlab. To represent a pulse at the far-end with equal rise and fall times, a normal distribution waveform was used, and the desired input waveform was found using the following equation.

$$V_{in}(t) = IFFT[FFT(V_{out}(t))/FFT(h(t))] \quad (6.2)$$

The near-end signal which generates the given far-end signal is shown in Figure 6.5, where we see positive and negative spikes have identical amplitudes as in the first experiment. The small positive dip after the negative spike is required to critically damp the falling part of the pulse to

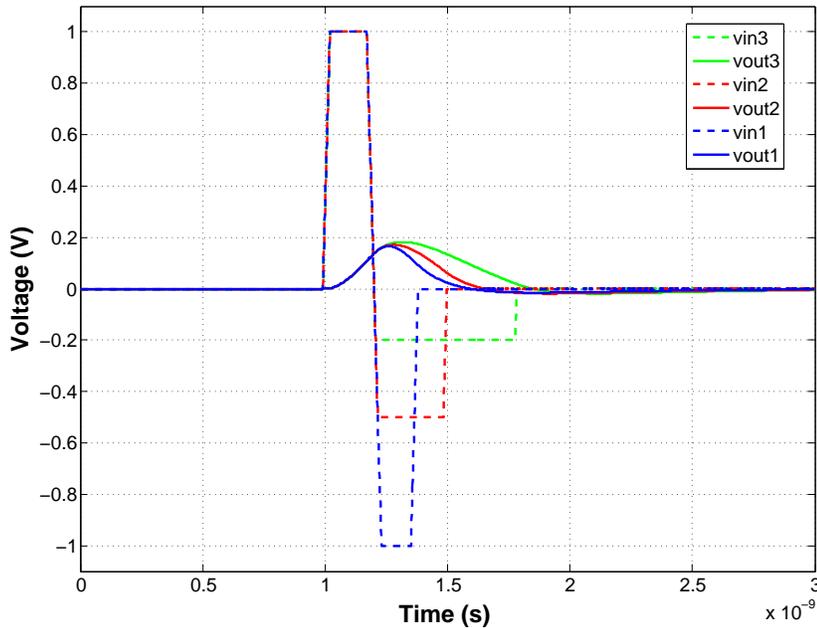


Figure 6.4: Far-end signal for various pulse shapes.

zero. Transmitter circuitry to create these signals and receiver circuitry to effectively recover NRZ signals from fast RZ pulses are described in the next section.

6.4 Transceiver Circuit Design in SDR Scheme

Most modern microprocessors employ global clocking schemes, but most on-chip low-swing designs [83, 86, 88, 89] either use different clocks for the transmitter and receiver side or require phase adjustment of the receiver clock, thereby making it hard to adopt the circuit technique to existing designs. In our work, we assume a global synchronous clock and use the same clock for both the flip-flop before the transmitter and the flip-flop after the receiver. In this case, the clock frequency, which is dictated by the latency of the interconnect system, will have to match the bandwidth. To achieve high bandwidth per cross sectional area, we use dense wires of $0.28\mu\text{m}$ width and $0.28\mu\text{m}$ spacing (twice the minimum pitch). Considering that state-of-the-art microprocessors operate from 1.5 to 5GHz , in order to make the total delay less than one clock period, we chose the flop-to-flop distance to be 5mm . Global wires which are much longer than this will include several of the proposed 5mm interconnect systems in series.

We used differential signaling to reject common-mode noise on low-swing signals, and we

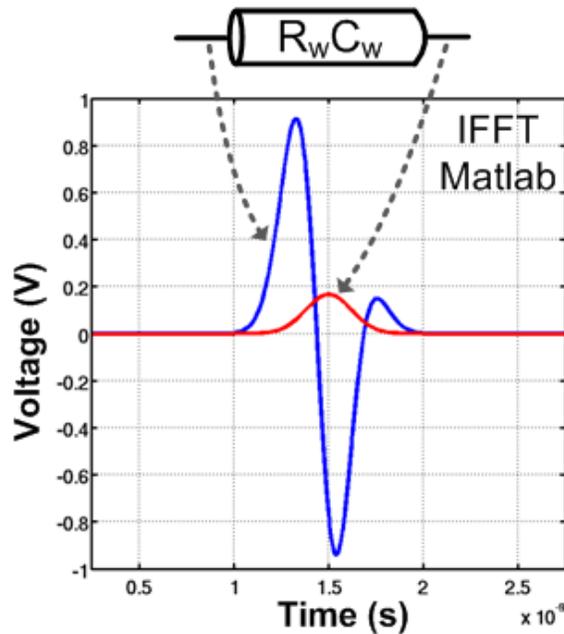


Figure 6.5: For a given signal at the far-end, desired signal at the near-end is found.

optimally twisted the long wires to prevent worst-case crosstalk between adjacent wires. Using a nominal supply voltage of 1V in 90nm CMOS, the transceiver system is designed such that 100mV differential swing is achieved at the receiver input.

6.4.1 Transmitter Design

To eliminate ISI over RC-dominated long wires, the proposed transmitter generates pre-emphasized bipolar pulses that follow the input transition (*i.e.*, negative-positive bipolar signal for falling transitions, positive-negative bipolar signal for rising transitions, as shown in Figure 6.6). Three series capacitors are used in the transmitter, where each capacitor explicitly controls a pre-emphasis spike. The rising (falling) data transition goes through the first capacitor to initiate the intended positive (negative) transition. After the data is delayed and inverted, it creates a larger negative (positive) spike through a larger second capacitor. The data is delayed and inverted again, which finally creates a third tap response to critically damp the pulse at the end of the wire as shown in Section 6.3. The third tap improves the bandwidth of received RZ pulses by generating the latter falling (rising) slope of the pulse equally fast as the initial rising (falling) slope, which is achieved at a small energy overhead since the third capacitor is relatively small. The Matlab analysis and

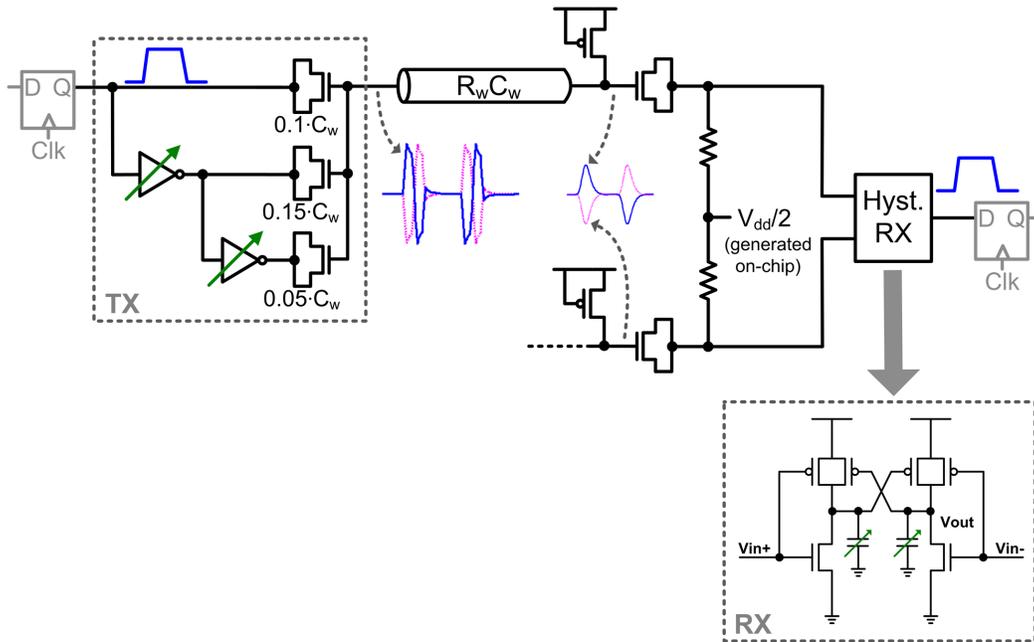


Figure 6.6: Proposed transmitter with receiver circuits with waveforms when ‘001100’ patterns is sent over on-chip links. Note that only 01 (rising) or 10 (falling) patterns generate pulses on the wire. Transceiver remains idle with consecutive 0s or 1s.

SPICE simulations using extracted wire models gave optimal tap values of 1, -1.5 , and 0.5 . Note that the taps sum to zero for this RZ signal, unlike the taps for NRZ signaling [85].

Although using three series capacitors in the transmitter creates overhead in area and capacitive load, the area overhead can be minimized by using NMOS transistors for capacitors instead of creating capacitors using a pitchfork structure [85]. We connected the source-drain of the NMOS transistors to the transmitter output, and the gate of the NMOS transistors to the wire.

6.4.2 Hysteresis Receiver

Previous low swing NRZ signaling schemes [84–86] used clocked comparators followed by SR latches. Because these clocked receivers present a high clock load, and precharge/evaluate every cycle even without data transitions, they tend to consume a significant fraction of the energy of the total communication system. In this design, we adopted a simple clockless hysteresis receiver, often used in off-chip communication [92, 95], to recover NRZ signals from RZ pulses (Figure 6.6).

To detect sharp RZ pulses, hysteresis receivers are simpler than clocked sense-amplifiers [85] or decision feedback equalization (DFE) receivers [86, 88], because they do not need a clock edge carefully positioned on the pulse, a requirement made difficult by link and process variations. Hysteresis receivers consume no energy with idle inputs, unlike clocked receivers with precharge and evaluate [85, 86], and also reduce clock load and simplify timing verification. In exchange, they are less efficient in evaluating switching inputs and need careful noise margin checks. The hysteresis circuit was designed to support bandwidths up to 6Gbps, well above the target for the rest of the link.

The hysteresis receiver evaluates a new output value only when the differential receiver inputs split by more than a certain threshold; otherwise, in case of no data transition, the receiver holds the previous value. Since the hysteresis receiver output directly drives the subsequent flip-flop, drop-in replacement of transceiver and wire blocks between flip-flops could also be more easily done.

The differential NMOS pairs and the cross-coupled PMOS pairs needs to be sized carefully because (1) hysteresis uses contention between NMOS inputs and PMOS pull-ups and (2) it determines the speed of the receiver. However, these transistors cannot be sized too large because they directly add capacitance to the output nodes, leading to excessive hysteresis. Without oversizing the transistors, the amount of hysteresis can be controlled by varying the capacitance of the output nodes as shown in Figure 6.6. Therefore, improving the speed of the hysteresis receiver should be accomplished by having larger differential amplitude in the inputs, rather than sizing the transistors in the receiver itself. Biasing the inputs of the hysteresis receiver will be further discussed in the following section.

6.4.3 Biasing of wire and receiver

A second series capacitor [89] at the end of the wire separates the bias of the long wire and receiver for several reasons. First of all, the inputs of the hysteresis receiver have to be biased around $V_{dd}/2$ since hysteresis is built upon inherent fights between pull-down of NMOS pairs and pull-up of cross-coupled PMOS pairs, as mentioned before. The hysteresis does not exist when the inputs are biased around V_{dd} or GND . We decided to generate a bias voltage of $V_{dd}/2$ using a

on-chip capacitive divider, instead of using an inverter with negative feedback through a resistance [89, 92], because the continuous negative feedback through a resistor implies a fight between the rising (falling) signal and the feedback strength to bring the signal back around $V_{dd}/2$. Isolating the receiver inputs from the long wire with large capacitance, the second series capacitor also minimizes the current requirements of the reference bias, thereby allowing the reference bias to charge only a small capacitance. The on-chip generated bias voltage can be shared among multiple hysteresis receivers.

In contrast to the receiver inputs, the wire requires a higher bias voltage because the series capacitors are all implemented with compact NMOS native devices that exhibit their largest capacitance at high gate voltages. Because the wires send RZ pulses, both differential lines remain at the same DC voltage when there is no data transition; this allows a simple biasing of the differential wires through leaky PMOS transistors. Bias circuits for NRZ signaling, however, are more complicated or impose DC-balanced data restrictions [85, 86].

6.5 Adaptive Pre-Emphasis in DDR Scheme for Further Bandwidth Improvement

In the 3-tap transmitter, a rising transition generates a positive pulse followed by a negative pulse on the wire. If another data transition (in this case, a falling edge) immediately occurs, the trailing negative pulse of the current rising data bit will be adjacent to the leading negative pulse of the successive falling data bit (see Figure 6.7). This suggests that if the two bits were partially overlapped, such that the two negative pulses coincided, the pre-emphasis would effectively double, causing the wire voltage slew rate to also double. Therefore, this circuit supports sending data on the wire at double the bandwidth, as in DDR links, with a new bit each clock phase: the increased pre-emphasis allows the wire to keep up with the circuits by generating suitably sharper pulses. This doubling of the transmitter pre-emphasis, and hence wire performance, happens only when two transitions occur back-to-back. A data transition followed by constant data would not double the pre-emphasis, as the hysteresis receiver allows the wire response to be slower. In other words, the transmitter adaptively employs higher pre-emphasis only when needed, without any

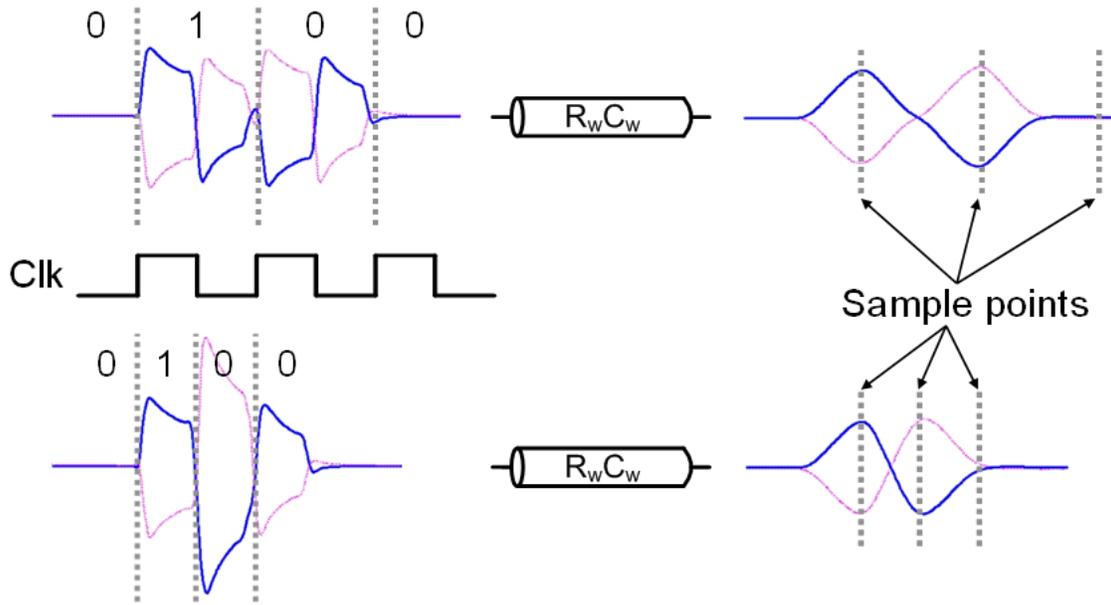


Figure 6.7: Further bandwidth improvement using double-data rate(DDR) scheme.

special encoding. This DDR scheme uses the same circuits as the SDR scheme, but adds a differential amplifier in front of the receiver to improve its performance.

The overall system to enable this technique for further bandwidth improvement is shown in Figure 6.8. Comparing to the SDR scheme in Figure 6.6, the flip-flops are changed to dual-edge flip-flops to send and receive data at both positive and negative edges of the clock. Low V_t transistors are used in these dual-edge flip-flops to achieve better latency. After the second series capacitor at the end of the wire, we added a simple differential amplifier to amplify the pulse swing at the receiver inputs. As we were trying to achieve higher data rates ($\sim 5\text{Gb/s}$), the main limiter was the speed of the hysteresis receiver because we cannot oversize the the cross-coupled PMOS devices as described in Section 6.4.2. Amplifying the receiver input signals is the effective way to improve the latency of the receiver, and this is possible because the second series capacitor isolated the receiver inputs from the long wire.

Figure 6.9 shows simulated waveforms of intermediate nodes in the new communication system, where data is sent on both positive and negative edges of a 2.5GHz clock. The pulse edges are controlled adaptively for different data patterns, and 5Gb/s signaling is demonstrated. Bandwidth density of $\sim 4.5\text{Gb/s}/\mu\text{m}$ is achieved at the expense of additional clocking energy and differential amplifier energy.

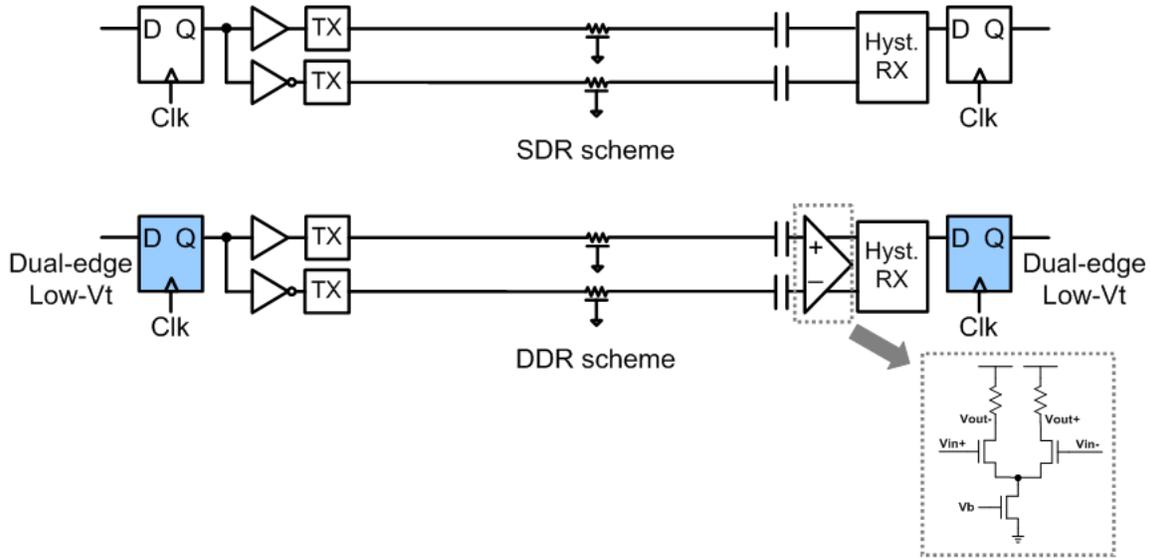


Figure 6.8: Communication system comparison employing single-data rate (SDR) and double-data rate(DDR).

6.6 Measurement Results

We fabricated a 90nm CMOS testchip (Figure 6.15) that included 3-bit 5mm links using conventional repeaters, a previous capacitive driver [85,86], and the proposed SDR and DDR schemes, as shown in Figure 6.10. The proposed and previous schemes both employ optimally twisted differential M5 wires with $0.28\mu\text{m}$ width and $0.28\mu\text{m}$ spacing (2X minimum pitch), while the conventional scheme is optimally repeated with 4 repeaters on single-ended M5 wires with $0.56\mu\text{m}$ width and $0.56\mu\text{m}$ spacing for the same footprint. Note that the design point in the conventional scheme represents an 11% delay increase with 21% lower energy over the optimal delay point in the width/spacing design space. In all schemes, M4 and M6 layers are filled with densely packed orthogonal interconnects.

As briefly mentioned in Section 6.4, the wires were chosen to be 5mm to match “long” wire lengths most commonly found in high-performance systems running at 2.5 GHz. Longer wires would require a larger wire pitch to overcome series losses and are typically flopped in the system architecture. All capacitive driver schemes employed a 100mV differential swing at the end of the wire.

Pseudo-random binary sequence (PRBS) data is generated off-chip and directly sent to the on-chip test structures. The measured BER for both SDR and DDR schemes are less than 10^{-10} .

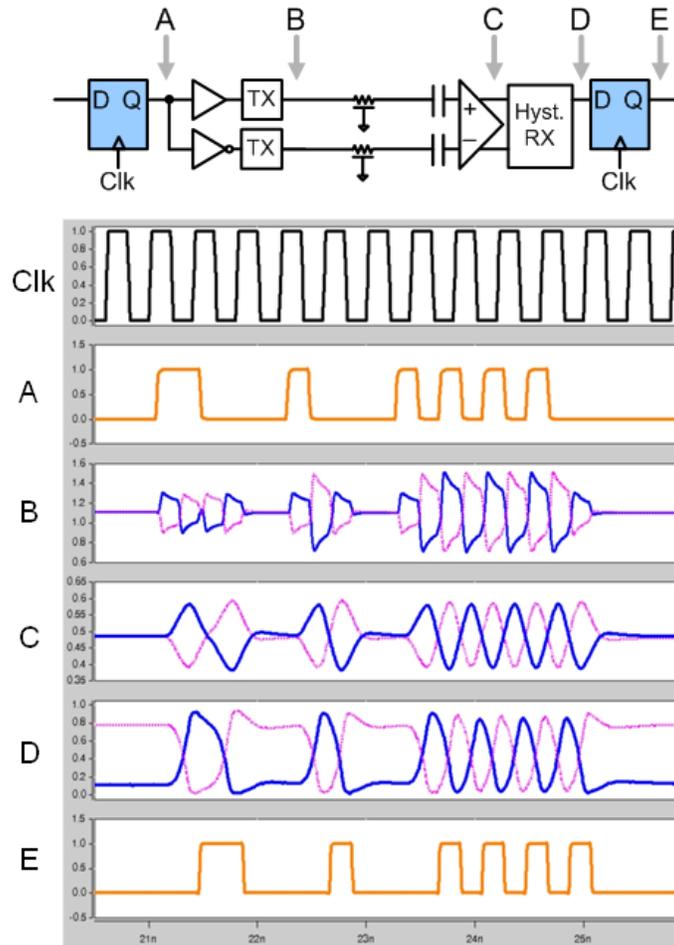


Figure 6.9: Simulated waveforms of intermediate nodes in the DDR communication system.

Energy versus performance characteristics are shown in Figure 6.11. It seems that the performance improvement of the capacitively driven wire over the conventional repeater design shown in [85] does not hold for narrow wires in scaled technologies. If the SR latch delay is included in the critical path of [85], its performance would be worse. The proposed schemes improved the performance over prior approaches to 2.5 Gb/s (SDR) or 4.9 Gb/s (DDR), while achieving energy consumption of 0.24 pJ/b (SDR) or 0.34 pJ/b (DDR). Comparing the receiver input amplitudes from chip measurements to those from extracted simulations, it was observed that the performance was limited by smaller-than-expected capacitance of the NMOS native devices, which resulted in reduced signal swing.

As shown in Figure 6.11, adaptive transmitter pre-emphasis in the DDR scheme provides $\sim 2X$ bandwidth density improvement with 38% increase in energy consumption, due to additional clock

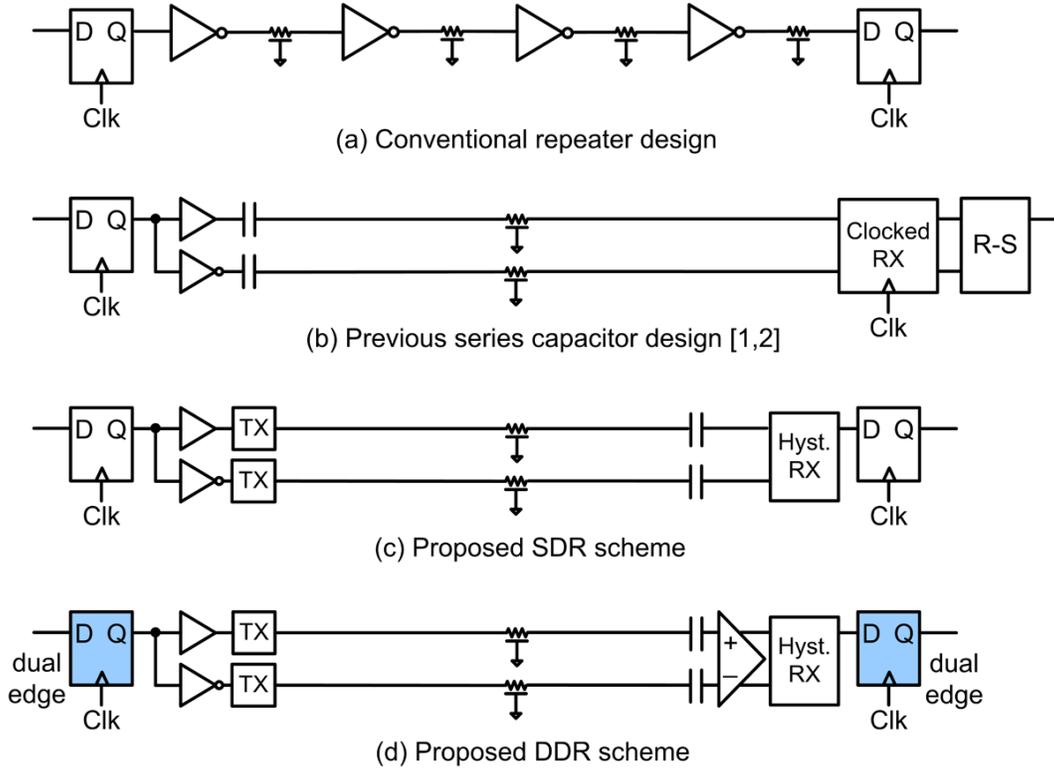


Figure 6.10: Overall block diagrams of four communication schemes: (a) conventional full-swing repeater scheme (b) single series capacitor scheme [85,86] (c) proposed SDR scheme (d) proposed DDR scheme.

energy and the differential amplifier. Measured energy per bit is plotted as a function of data transition activity in Figure 6.12, and it is shown that the energy consumption scales well with low data activity. The total energy consumption at zero data activity approaches 50fJ/b, and according to the extracted simulation, 40% of this energy is consumed by clock energy at the flip-flop, and the remaining 60% is dissipated in the hysteresis receiver.

For observability purposes, multiple on-chip samplers [96] were placed in various places to probe high-speed on-chip signals inside the prototype chip. Figure 6.13 shows the probed waveforms of the transmitter output and receiver input signals through on-chip samplers for a 000010000 data pattern in the DDR scheme, where adaptive pre-emphasis and slew rate control is demonstrated.

Figure 6.14 shows the bandwidth density and energy per bit comparison between the proposed work and representative previous works from literature. Note that each design has its own technology, wire geometry, and wire length, which directly affects the achievable bandwidth and energy

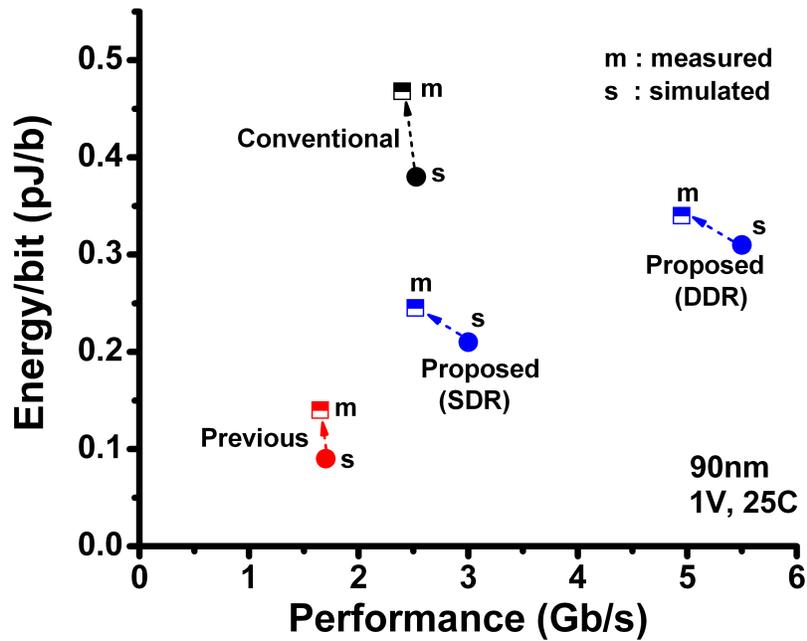


Figure 6.11: Measured energy and performance of conventional, previous [85, 86], and proposed scheme.

consumption of on-chip links. It should be also taken into account that designs with shorter wire lengths would achieve higher bandwidths than those with longer wire lengths (i.e., 10mm), but our proposed design show superior bandwidth density with fairly low energy consumption compared to other designs.

6.7 Summary

The proposed transceiver design for repeaterless on-chip communication demonstrates high bandwidth density, low latency, and low energy consumption. A 3-tap transmitter significantly reduces ISI over long narrow wires, and a simple hysteresis receiver recovers the resulting low-swing RZ pulse. With DDR signaling the transmitter pre-emphasis is adaptively controlled, enabling a data rate of 4.9 Gb/s/ch and bandwidth density of 4.4 Gb/s/m over 5mm on-chip links with 0.34 pJ/b energy consumption.

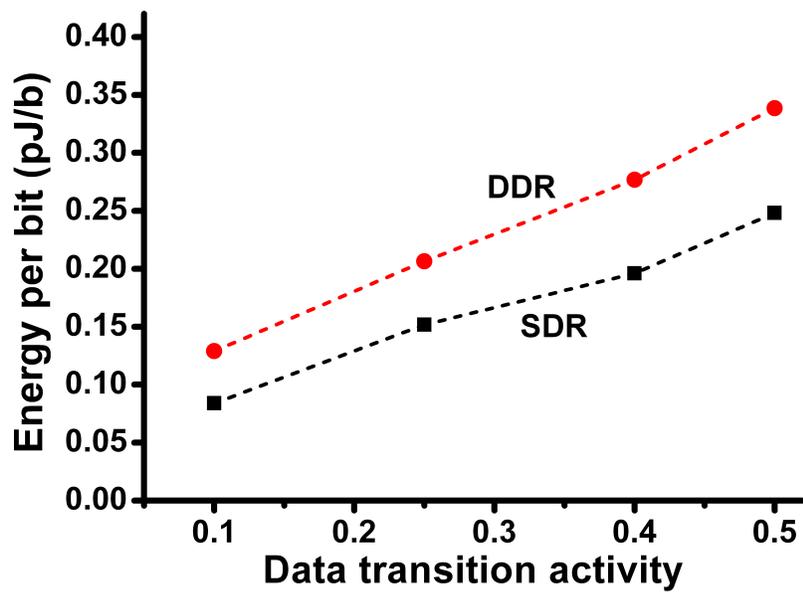


Figure 6.12: Energy versus data activity of the proposed work.

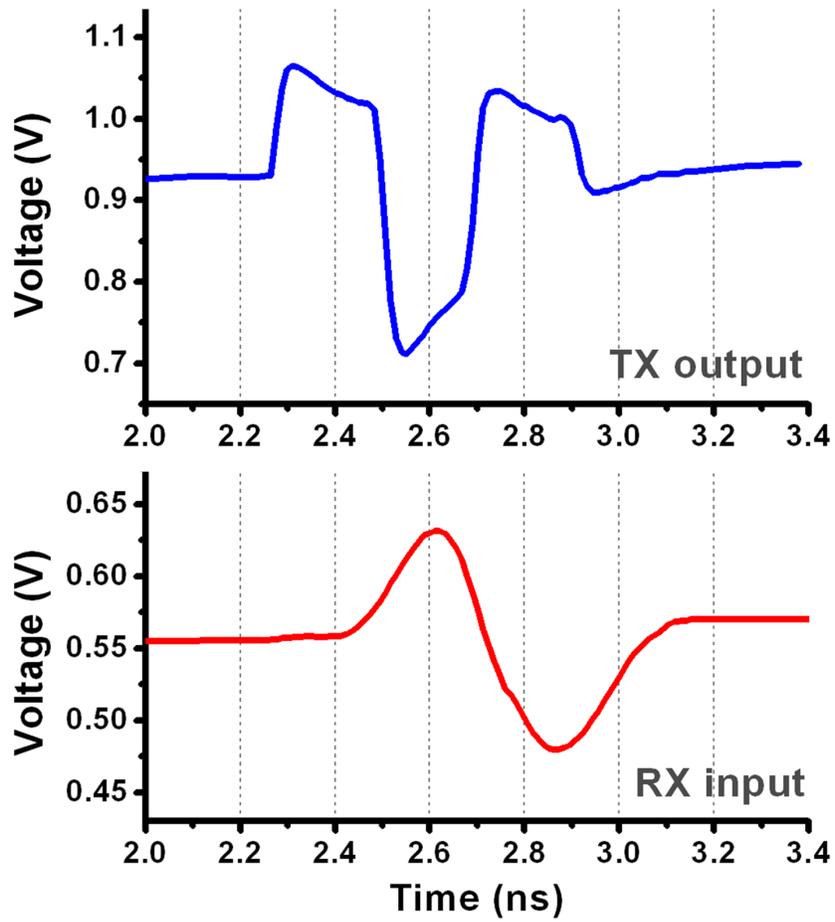


Figure 6.13: Measured waveforms of transmitter output and receiver input signals.

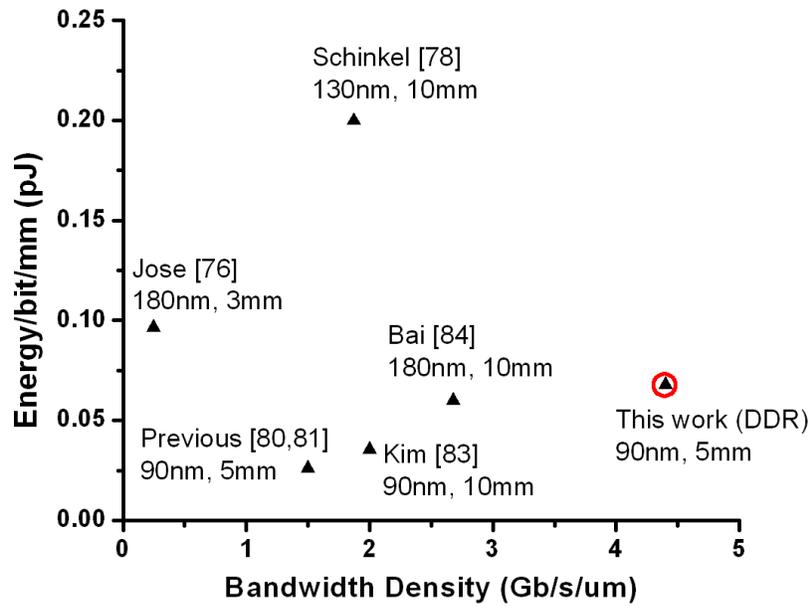


Figure 6.14: Bandwidth density and energy per bit comparison between proposed work and literature.

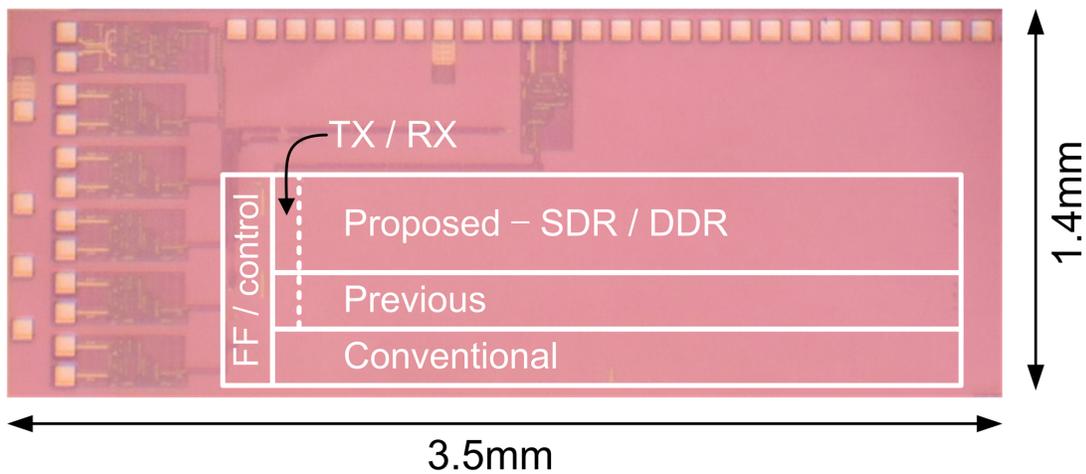


Figure 6.15: Chip micrograph.

CHAPTER 7

Effect of Long Wires on Technology Mapping

Technology scaling reduces gate delays while wire delays may increase. Our work studies the interaction of this phenomenon with technology mapping and its impact on modern EDA flows. In particular, we demonstrate that the use of larger standard cells increases the number of long wires and may undermine circuit delay optimization at 65nm and below. Experiments with 130nm, 90nm, 65nm, and 45nm industrial CMOS technology suggest that limiting the use of larger standard cells in technology mapping becomes more effective at 65nm and 45nm node, resulting in up to 12% improvement in critical path delay on large benchmark circuits.

7.1 Motivation

Over several decades, technology mapping has been extremely useful for reducing the device area of complex logic. Furthermore, recent research in Boolean matching [63, 64] accomplished dramatic efficiency improvements for function matching, facilitating new technology mapping algorithms that can deal with 10-input gates. However, extending these algorithms with proper models of circuit delay and validating them with respect to recent technology nodes remains a major research challenge.

While technology mapping seeks to minimize device count, the bulk of critical path delay has shifted from gates to wires in the last 5 years. In particular, the number of repeaters required is exponentially increasing with each technology step [65, 66], and 10~15% of gates in large microprocessor chips are buffers that break down long interconnects. Extensive literature exists on

optimal buffering [67–69] that employs fairly accurate delay modeling, but does not attempt logic restructuring.

Our work is motivated by the apparent dichotomy between (1) the literature on buffer insertion that improves circuit performance by adding a large number of one-input one-output gates (buffers and inverters) that do not perform any logic operation, and (2) the literature on functional technology mapping, which clusters logic into 5-15 input gates, improving area, but does not evaluate overall circuit performance with respect to current technology nodes.

Previous literature [70–73] suggests that technology mapping must interact with placement of the standard cells and use accurate interconnect models for performance optimization. These works improved the critical delay through either integration of layout information in early logic synthesis stage [70, 71] or iterative re-synthesis with placement information [72, 73], but they have not considered the impact of technology mapping on global buffer counts and the overall circuit performance after place-and-route optimizations.

This chapter proposes to, ironically, undo technology mapping for high-speed designs through reducing the wire delay components in the critical path of large circuits. Foregoing aggressive technology mapping and using a large number of standard cells (but of smaller size) will eliminate the need for excessive buffers during post-placement timing optimization. The discussions and experiments in this work also consider coupling capacitance between adjacent wires which dominates the wire capacitance in most advanced technologies, and we attempt to reduce the parallel run length of neighboring wires.

Recent work [77] points out that conglomerating small cells into a large cell may produce non-

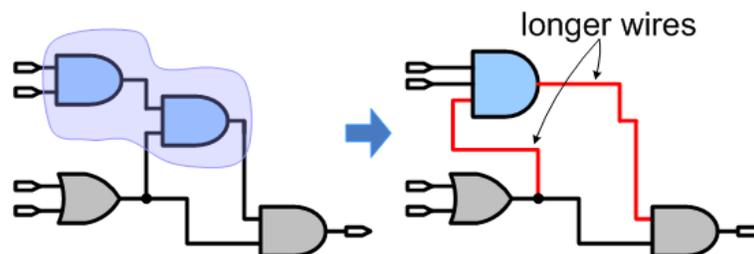


Figure 7.1: Indiscriminate technology mapping may produce longer wires, adversely affecting delay and routing congestion.

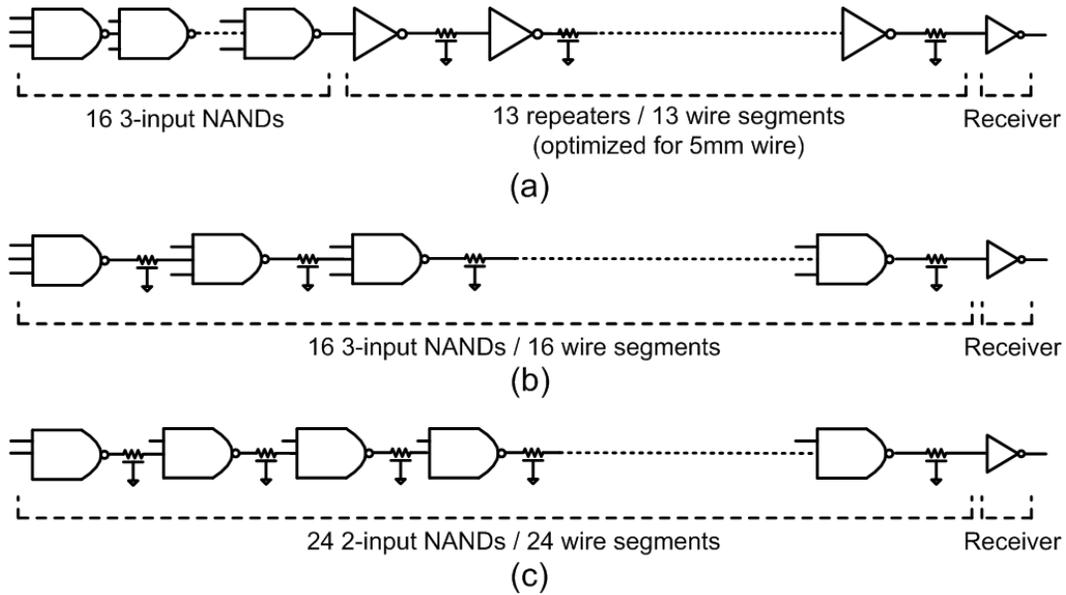


Figure 7.2: Three schemes for comparison of single paths (a) Logic block (16 3-input NANDs) driving an optimally repeated 5mm wire (b) 16 3-input NANDs are placed along the wire (c) 16 3-input NANDs are decomposed into 24 2-input NANDs and placed along the wire.

monotonic interconnects which adversely affect delay and routability as illustrated in Figure 7.1. By limiting the use of large standard cells, our approach inherently blocks the occurrence of this disadvantageous technology mapping, and results in a number of shorter monotonic wires.

Our considerations and conclusions are intended for ASIC/SoC designs rather than FPGA or microprocessor designs. In FPGA designs, programmable interconnect is uniformly buffered and linear wire delays do not significantly depend on whether long nets are broken into shorter segments. On the other hand, technology mapping into LUTs is an important and difficult task, still necessitating technology mapping in FPGAs. In high-end microprocessor designs, clock period is short and the logic between pipelines are often dominated by large fanouts. In this case, the number of inserted buffers cannot be reduced significantly by reducing the length of wires.

7.2 Analysis of Single Paths

As a proof of concept, we conducted a simple experiment as depicted in Figure 7.2. As a baseline for comparison, 16 3-input NAND gates drive a M5 minimum-pitch 5mm wire, which

is optimally repeated in 65nm technology. Traditionally, we would have the combinational logic placed in a denser cluster for minimum area as shown in Figure 2(a). Instead, however, we spread out the 16 NAND gates regularly along the interconnect to implement the logic and also serve as repeaters in Figure 2(b), as previously proposed in [79]. These distributed NAND gates eliminate long wires and the need for repeaters, resulting in actually better performance. This effect could be exploited by decomposing the logic into more gates, i.e., undoing technology mapping. In Figure 2(c), the long wire is divided more finely with more logic gates (24 2-input NANDs) for the same functionality. Note that, in Figure 2(a)-(c), the inputs of the NAND gates which are not in the critical path are tied to Vdd for worst-case rising delays.

HSPICE simulation is done with industrial 65nm CMOS technology, where all three schemes are swept with sizing, and the optimal energy versus delay results are shown in Figure 7.3. Comparing to scheme (a) at iso-energy of 1.1pJ, scheme (b) achieves 13% delay reduction and scheme (c) achieves 18% delay reduction. Overall, (c) improves the energy-delay curve of (a) by a significant amount. At these delay points, Figure 7.4 shows a delay breakdown of the three schemes. By spreading out the NAND gates in (b), logic gate delay is increased since the load capacitance of the NAND gates is increased due to wires, but the repeater delay portion is eliminated and the overall delay is reduced by 13%. Through undoing technology mapping in (c), wire delay is further reduced due to fine chopping of wires and gate delay also slightly decreased due to reduced load capacitance. Interestingly, if long wires are present in the circuit delay, spreading out the gates and using more gates to implement a given logic could actually improve delay since they convert wire delay back to logic delay.

One possible drawback of this approach is that the inputs to the NAND gates in the middle of the wire have to be routed to the intermediate placement locations, but the surrounding logic gates could be restructured and placed nearby the middle of the wire. We scrutinized this in the following section by performing synthesis, placement and routing on large benchmarks.

7.3 Evaluating utility of large cells in technology mapping

To evaluate the utility of technology mapping for general circuits in scaled technologies, we compare pairs of libraries for several benchmarks with each technology. ‘Original’ scheme uses

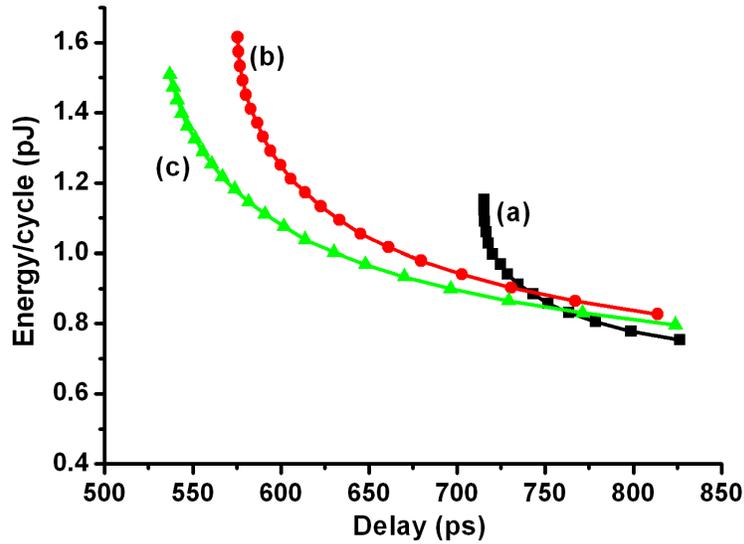


Figure 7.3: Energy versus delay comparison for the three different schemes in Figure 7.2.

original standard cell library without any restriction, while ‘No Large Cells’ scheme is confined to the library where there are only 1-input and 2-input gates available.

7.3.1 Methodology

Figure 7.5 shows the flow chart for both approaches. Starting from the same behavioral netlist, logic synthesis (Synopsys Design Compiler 2007.03-sp2) is applied for each scheme with a restriction on the ‘No Large Cells’ scheme to use only 1-input or 2-input standard cells. After logic synthesis, the structural netlist goes through timing-driven placement, physical synthesis, and timing-driven routing (Cadence SoC Encounter 6.1.2). Post-placement logic restructuring is executed if necessary, but the restriction on the number of inputs of gates still holds in the ‘No Large Cells’ scheme. Finally, timing analysis is performed for both approaches with all back-end parasitics including coupling capacitance. This procedure was done for industrial 130nm, 90nm, 65nm, and 45nm technologies, and benchmark circuits from IWLS 2005 [74] were used (s35932 from ISCAS family and the rest of them from OpenCores family). In the overall flow, the proposed scheme does not add any intermediate steps or iterations to the baseline. In fact, our approach seeks to reduce resource utilization (less standard cells from the library) while also improving delay.

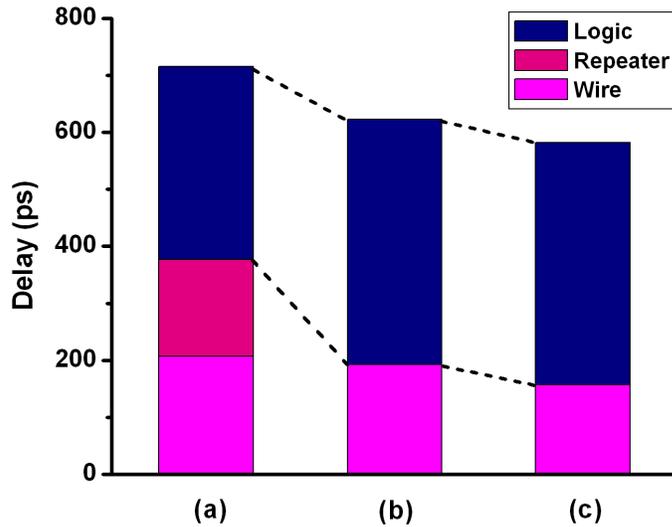


Figure 7.4: Delay breakdown (logic delay, repeater delay, and wire delay) of the three schemes in Figure 7.2 at iso-energy of 1.1pJ.

7.3.2 Experimental Results

One expects the ‘No Large Cells’ approach to increase the gate count due to a more limited standard cell library. However, the critical path could actually benefit from more gates since both the wire capacitance and the number of required buffers are reduced.

Figure 7.6 compares the critical path delay between ‘Original’ and ‘No Large Cells’ configurations for eight benchmarks. Delay of ‘No Large Cells’ scheme is normalized to that of ‘Original’ scheme. The monotonic trend shown in Figure 7.6 illustrates the decreasing utility of large standard cells in technology mapping for more advanced technologies. At 65nm and 45nm technology, discarding large standard cells (3-inputs or more) gave better results (1-12%) in critical path delay than the original technology mapping for all benchmarks. Breaking up the wire into more segments proves to be effective at 65nm and below through reducing the wire delay components. The delay breakdown for benchmark `wb.conmax` is shown in Figure 7.7 across four technology nodes. Gate-dependent delay is defined as the sum of intrinsic gate delay and gate load delay, which is basically the circuit delay when no wire is present. Wire-dependent delay consists of inserted buffer delay and wire load delay, which are the delay elements generated due to routed wires. It can be seen that our approach increases gate-dependent delay by a minimal amount, but the wire-dependent

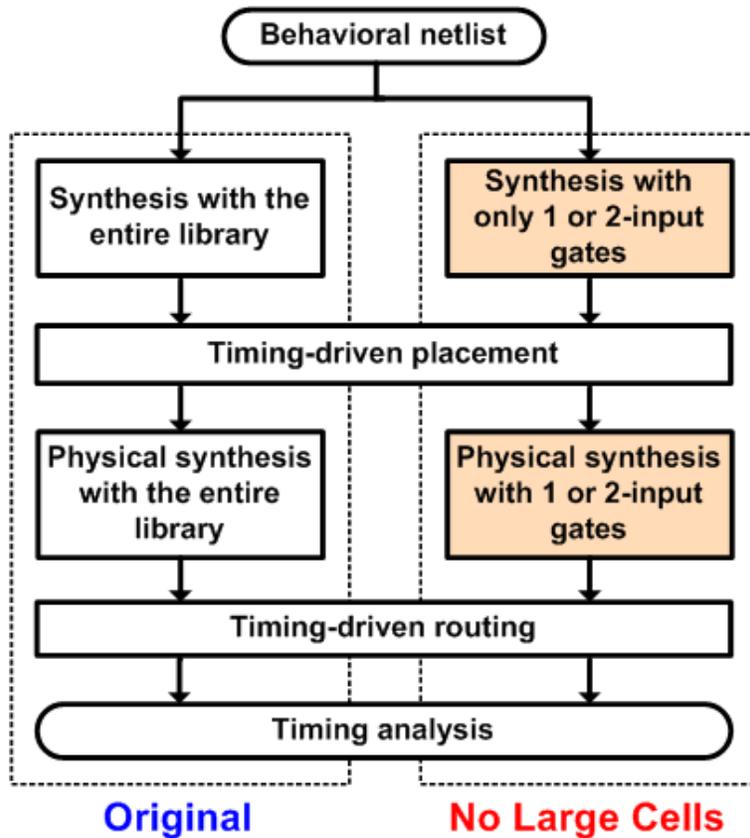


Figure 7.5: Flow chart for the methodology of ‘Original’ and ‘No Large Cells’.

delay component is reduced significantly (35% in 45nm), leading to an overall 12% performance improvement in 45nm node. Note that the relative portion of wire-dependent delay grew considerably at the 45nm node. This is mostly due to the sharp increase of resistance of minimum width wires in 45nm, considering that the capacitance of a unit length wire does not change significantly for each technology step.

Table 1 shows a detailed comparison on several metrics for eight benchmarks for both 65nm and 45nm technology to check whether the ‘No Large Cells’ approach is working as proposed. Typically a large number of buffers are inserted during timing optimization for the given benchmark circuits, and the number of buffers is reduced by 5~54% by breaking the long wires into short wires with more gates. Average wire length and wire capacitance (both total and coupling) show noticeable reduction except for the relatively small s35932 benchmark. The reduction in coupling capacitance is more than that in ground capacitance, which is due to the observed higher routing congestion in intermediate and high metal layers in the ‘Original’ configuration leading to

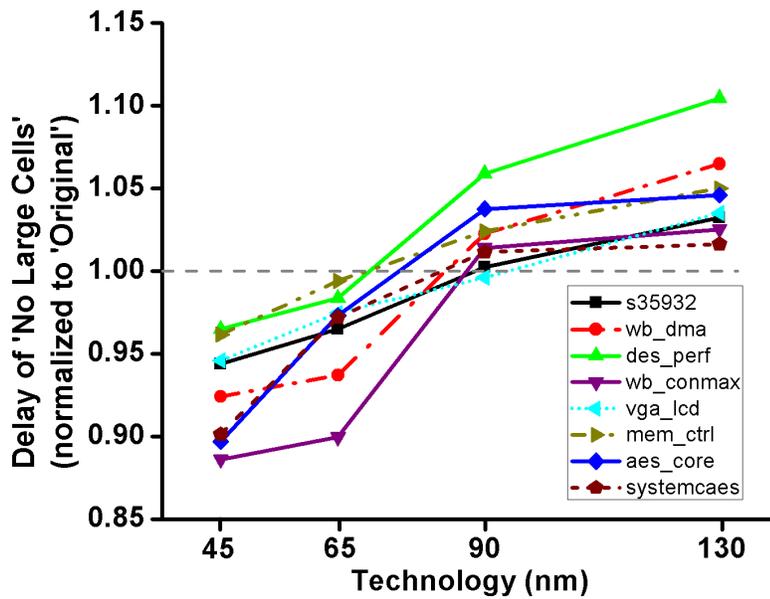


Figure 7.6: Critical path delay comparison of IWLS benchmarks using ‘Original’ and ‘No Large Cells’ approach in 130nm, 90nm, 65nm, and 45nm technology.

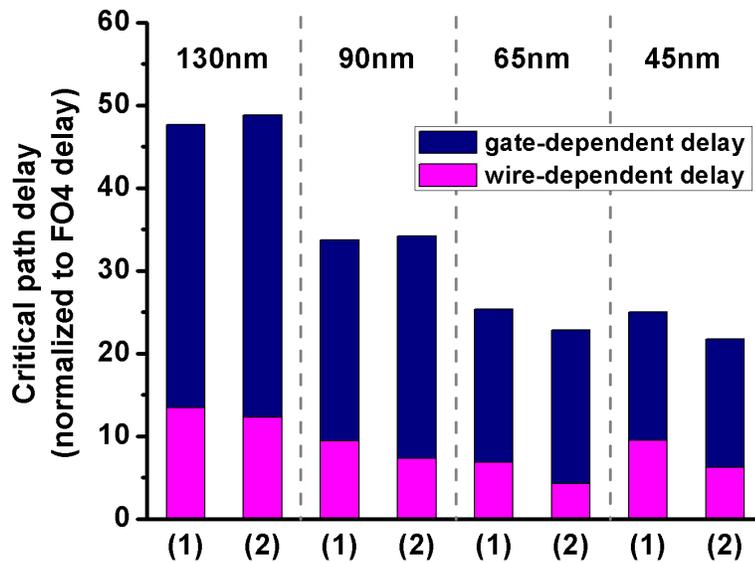


Figure 7.7: Critical path delay breakdown (gate-dependent delay and wire-dependent delay) of benchmark wb_conmax for (1) ‘Original’ and (2) ‘No Large Cells’ approach across four technology nodes.

Table 7.1: Detailed comparison of the benchmarks for ‘Original’ and ‘No Large Cells’ scheme on critical path, average wire length (=total routed wire length/wire count), inserted buffer count, total standard cell count, wire capacitance, and total standard cell area is shown for (a) 65nm and (b) 45nm technology.

(a) 65nm

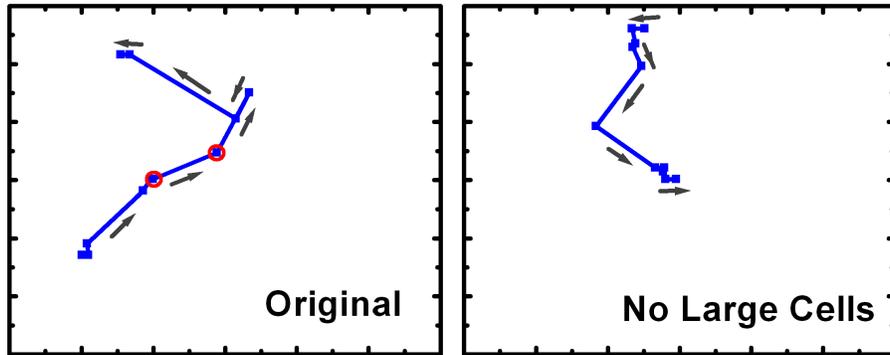
Benchmark	Original						No Large Cells (vs. Original)					
	Critical path delay (FO4)	Avg. wire length (μm)	Buffer count / total cell count	Wire capacitance		Cell area (μm^2)	Critical path delay (FO4)	Avg. wire length (μm)	Buffer count / total cell count	Wire capacitance		Cell area (μm^2)
				Ctot (fF)	Cc (fF)					Ctot (fF)	Cc (fF)	
s35932	10.9	21.5	279/5764	19.3	9.25	35512	-3.5%	+3.4%	-10%/+35%	+14%	+12%	+10.6%
wb_dma	14.9	36.5	180/4968	41.0	28.5	24069	-6.3%	-15.8%	-19%/+19%	-5%	-8%	+4.0%
des_perf	20.1	23.7	511/69733	244.8	128.2	320170	-1.7%	-11.3%	-44%/+12%	-2%	-4%	-1.6%
wb_conmax	25.3	69.7	921/24720	405.5	330.8	112590	-10.0%	-45.0%	-15%/+73%	-21%	-26%	+18.3%
vga_lcd	28.2	40.7	3378/29778	193.3	94.6	264839	-2.5%	-23.5%	-5%/+32%	-9%	-18%	+12.9%
mem_ctrl	21.8	28.1	217/5227	22.9	12.2	50224	-0.6%	-23.1%	-27%/+37%	+2%	-1%	+6.5%
aes_core	25.6	26.0	475/18568	84.1	54.1	108502	-2.7%	-22.9%	-11%/+19%	-15%	-22%	-9.8%
systemcaes	30.1	34.0	832/5678	36.9	24.1	53441	-2.8%	-19.0%	-27%/+20%	+7%	+2%	+20.9%
Average							-3.8%	-19.7%	-20%/+31%	-4%	-8%	+7.7%

(b) 45nm

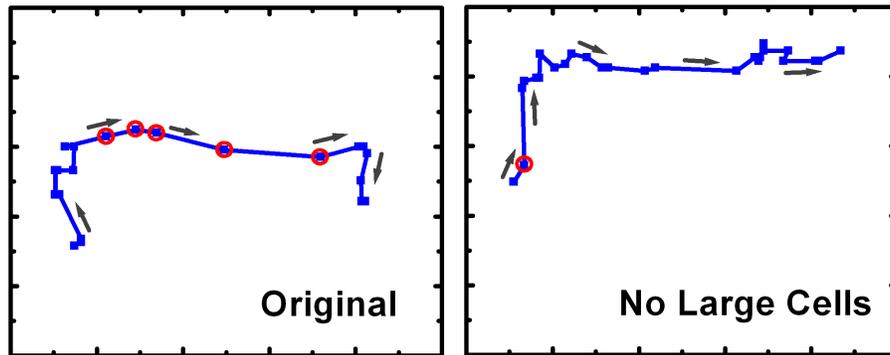
Benchmark	Original						No Large Cells (vs. Original)					
	Critical path delay (FO4)	Avg. wire length (μm)	Buffer count / total cell count	Wire capacitance		Cell area (μm^2)	Critical path delay (FO4)	Avg. wire length (μm)	Buffer count / total cell count	Wire capacitance		Cell area (μm^2)
				Ctot (fF)	Cc (fF)					Ctot (fF)	Cc (fF)	
s35932	12.6	14.0	283/7155	14.3	7.2	20131	-5.6%	-14.5%	-5.3%/+18%	+8%	+7%	+7.5%
wb_dma	15.7	27.3	268/6109	27.8	17.7	13600	-7.7%	-12.5%	-13%/+6%	-9%	-12%	-1.7%
des_perf	23.3	17.3	1599/85297	213.9	122.3	186916	-3.5%	-16.7%	-27%/+20%	-4%	-10%	+3.1%
wb_conmax	25.0	46.3	1874/27280	256.2	231.0	55171	-12.3%	-37.2%	-25%/+70%	-21%	-29%	+20.1%
vga_lcd	31.7	28.6	3807/42890	183.1	108.2	153413	-5.4%	-20.6%	-5%/+28%	-2%	-4%	+8.8%
mem_ctrl	24.3	20.3	307/6285	20.7	12.7	14733	-3.9%	-23.4%	-11%/+32%	-4%	-10%	+7.2%
aes_core	26.2	18.9	1329/17485	63.4	42.6	46640	-10.3%	-19.1%	-54%/+39%	-7%	-4%	+8.4%
systemcaes	30.4	26.4	1031/6557	32.6	21.1	15274	-9.9%	-22.2%	-36%/+16%	-7%	-2%	+15.7%
Average							-7.3%	-17.2%	-22%/+29%	-6%	-8%	+8.6%

increased coupling capacitance. This fact is encouraging because coupling capacitance increasingly dominates the overall wire capacitance with technology scaling. In the s35932 benchmark, the fact that the critical path delay marginally decreased despite an increase in wire length and capacitance suggests further improvement by introducing the proposed approach only on timing-critical nets. Benchmark des3_perf shows a small improvement in critical path delay in spite of a large reduction in the buffer count, especially in 65nm, because the inserted buffer count is a small portion of the total standard cell count (0.7%).

It is not surprising that the standard cell count is increased by 12~54%, but the standard cell area overhead is only 8.6% on average at 45nm technology since 1-input and 2-input gates are typically smaller than complex gates. This area increase would not necessarily result in comparable die area increases in modern microprocessors or SoC designs because embedded memories and



(a) Benchmark wb_dma at 65nm node



(b) Benchmark systemcaes at 45nm node

Figure 7.8: Critical path comparison between ‘Original’ and ‘No Large Cells’ configuration for benchmarks (a) wb_dma at 65nm technology node and (b) systemcaes at 45nm technology node is shown (dots with circles represent inserted buffers).

hard IP blocks consume a large portion of the total chip area, making standard cell area a relatively lesser concern [75, 76]. Also, in designs with hierarchical floorplans, increasing the area of one partition does not affect the area of the entire chip, and designs requiring high I/O bandwidth (such as network processors) are pad-limited. Furthermore, by expanding this work to remove large cells only from timing critical paths, similar delay results with much smaller area increases are expected since standard cells on critical paths are responsible for only a small fraction of overall cell area.

The critical paths and signal directions of benchmarks wb_dma (65nm node) and systemcaes (45nm node) for configurations ‘Original’ and ‘No Large Cells’ are visualized in Figure 7.8. For benchmark wb_dma, the path is noticeably shorter, has fewer long wires and no inserted buffers in the ‘No Large Cells’ configuration, yielding an improvement of 6.3% in critical path delay. Benchmark systemcaes in 45nm node is also a good example of effectively converting wire delay

into gate delay. In the ‘Original’ case, five buffers are inserted to send the signal to the distant location, whereas a number of small standard cells are spread out to serve as a repeater while also performing logic operation in the ‘No Large Cells’ case.

In addition to circuit performance, power consumption is considered in our analysis. Table 2 shows both dynamic and leakage power consumption of the final netlists for the two schemes in 65nm technology. We used randomized switching data with average activity factor of 0.2 for each benchmark, and measured power using Synopsys NanoSim. For a few benchmarks, power consumption of the ‘No Large Cells’ scheme is actually lower than that of the ‘Original’ scheme, due to the interaction of the appreciably lower buffer count and smaller wire capacitance. For the `vga_lcd` benchmark, buffer count is not significantly reduced by the simplified technology mapping, resulting in 7.1% and 4.8% power increase in 65nm and 45nm, respectively. The power overhead for the `wb_dma`, `mem_ctrl`, and `s35932` benchmarks is insignificant. Overall, despite the increased gate count, the capacitance of 1-input and 2-input gates is small, leading to comparable overall power consumption of the ‘No Large Cells’ scheme as that of the ‘Original’ scheme.

The leakage power overhead in Table 2 is largely proportional to the standard cell area increase in Table 1. More precisely, the reason why the leakage overhead is slightly larger than the area overhead is that small standard cells have shorter stacks of transistors leading to less stack effect and more leakage power. However, the additional leakage power is relatively small ($\sim 1/100$ of dynamic power in all benchmarks) and the net effect on total power as seen in Table 2 is very low for these typical high-performance designs.

Our results on full integrated circuits motivate placement-aware technology mapping and post-placement logic restructuring, which can indeed improve timing. However, commercial tools available to us only partially include this feature, and in its absence, we demonstrate that large standard cells are not particularly useful on critical paths. The arguments from Section 7.2 suggest that even with placement-driven technology mapping and post-placement logic restructuring, large cells will be less useful on critical paths. An additional advantage of our approach is that breaking down large cells into smaller ones improves routability by enhancing the ability to reduce routing congestion [77, 78].

Throughout these benchmark experiments for critical path delay optimization, we execute synthesis, placement, and routing for the same circuit. As a result the size of the circuit and the length

Table 7.2: Dynamic and leakage power comparison between ‘Original’ and ‘No Large Cells’ scheme for (a) 65nm and (b) 45nm technology.

(a) 65nm

Benchmark	Dynamic power		Leakage power	
	Original (mW)	No Large Cells	Original (mW)	No Large Cells
s35932	5.3	+0.6%	47.2	+14.8%
wb_dma	2.8	+1.1%	34.1	+5.9%
des_perf	26.4	-4.5%	448.4	0%
wb_conmax	11.7	-7.9%	145.3	+22.7%
vga_lcd	22.6	+7.1%	408.1	+11.5%
mem_ctrl	2.4	+0.5%	35.3	+11.3%
aes_core	9.4	-9.0%	95.6	+2.3%
systemcaes	4.0	+6.5%	37.9	+25.2%
Average		-0.7%		+11.7%

(b) 45nm

Benchmark	Dynamic power		Leakage power	
	Original (mW)	No Large Cells	Original (mW)	No Large Cells
s35932	2.7	+1.2%	67.9	+20.0%
wb_dma	1.3	+5.1%	48.1	+0.3%
des_perf	11.9	+0.4%	706.2	+6.1%
wb_conmax	10.3	-5.4%	242.8	+22.6%
vga_lcd	9.4	+4.8%	460.1	+7.0%
mem_ctrl	1.4	+1.5%	45.9	+14.6%
aes_core	3.6	+2.1%	162.7	+12.7%
systemcaes	1.8	-4.4%	59.4	+23.7%
Average		+0.5%		+13.4%

of long wires will also decrease for each technology step, which is why the wire-dependent delay in the ‘Original’ approach in Figure 7.7 decreases at each technology node from 130nm to 65nm. However, when technology scaling is used to double the number of on-chip transistors, the chip size and longest wires do not shrink. If technology mapping is skipped under this assumption (higher levels of integration for scaled technologies), wire delay will dominate due to inter-module communication and we suggest that the performance improvement using the proposed approach would increase.

7.4 Summary

Our work offers a first-of-a-kind careful analysis of technology mapping across four technology nodes. While this step has been commonly used in logic synthesis flows, we point out that the use of large standard cells in it appears unnecessary and even harmful for high-performance designs at 65nm and below (low power designs could still benefit from technology mapping through reduced leakage). This is a consequence of uneven scaling of wire and gate delay, as well as the fact that technology mapping essentially trades gate counts for an increased number of long wires (as shown in Table 1). Empirical trends observed for large benchmark circuits mapped to 130nm, 90nm, 65nm, and 45nm libraries suggest that the 65nm node is an inflection point for the utility of large cells in technology mapping.

CHAPTER 8

Conclusions

In modern VLSI designs, on-chip communication resources are increasingly being the bottleneck of the overall chip performance and energy consumption. To alleviate such problematic challenges in the perspective of a circuit designer, our work presented various circuit techniques and methodologies for high-performance and energy-efficient on-chip communication. These includes new techniques to improve the overall wire resistance or capacitance through reducing the effective miller coupling factor (MCF) in the presence of optimal repeaters (Chapter 2, 3, and 4), as well as new repeaterless signaling techniques for better energy-delay tradeoff (Chapter 5 and 6).

Edge encoding technique was introduced in Chapter 2 for energy-efficient multi-cycle interconnects. A simple encoder separated the incoming rising and falling transitions, eliminating the worst-case MCF of 2. Pulsed registers with transparency windows showed further performance improvement and energy benefits. Performance and energy analysis in the presence of PVT variation confirmed the robustness benefits over prior work.

In Chapter 3, an alternating repeater technique was presented for better energy-delay tradeoff. Alternating MCF of 0 and MCF of 2 between repeater segments guaranteed average MCF of 1 over every data pattern, and proved useful for multi-drop global busses.

Chapter 4 proposed a crosstalk-aware pulse width modulation (PWM) based circuit technique which compressed two bits of information into transition type and different pulse widths. Crosstalk-aware signaling used a pre-correction scheme to nullify the effects of crosstalk while pulses were propagated over long wires. The encoder and decoder could be self-calibrated against global and local variation.

To facilitate repeaterless full-swing signaling, self-timed regenerators (STRs) have been presented in Chapter 5. Without breaking the global wire, STR circuits were placed regularly along the interconnect and showed superior performance-energy trade-off than conventional full-swing repeaters in 90nm CMOS.

In Chapter 6, a transceiver system with adaptive pre-emphasis was designed to improve the achievable bandwidth in a RC-dominated repeaterless on-chip link while exploiting the energy benefits of low-swing signaling. A 3-tap FIR filter in the transmitter effectively eliminated inter-symbol interference (ISI) in narrow wires and a receiver employed hysteresis to recover the low-swing pulse. As a result, bandwidth density $4.4\text{Gb/s}/\mu\text{m}$ with 0.34pJ/b energy was achieved for 5mm links in 90nm CMOS.

The interaction between technology mapping, wire scaling, and repeater insertion was explored in Chapter 7. In the case of using large standard cells results in excessive buffer insertion in a large IC design, it was pointed out that the buffer insertion as well as the critical path delay could be improved by limiting the use of complex cells. For non-critical paths, however, large standard cells would still prove useful to reduce area.

For general on-chip links in microprocessors or ASIC designs, robustness, signal integrity, and stability are very important concerns, and noise margin cannot easily be sacrificed. Full-swing signaling maximizes the noise margin under nominal supply, and to meet certain performance and slew constraints, the link designers are forced to use repeaters. The techniques presented in Chapter 2 and 3 work well in these environments, and deliver improvement in on-chip link energy due to reduced coupling capacitance by transient and spacial separation of worst-case MCF. Both techniques showed robust operation across PVT variation corners.

In certain on-chip links where the point-to-point interconnect paths are well-defined such as links between crossbar and cores in a multi-core microprocessor, more aggressive approaches should be considered for better performance and power. Among the techniques using full-swing signaling, the one proposed in Chapter 4 showed the highest benefits in terms of energy-efficiency and performance. This was mostly because 2 bits were compressed onto one wire and hence the wire pitch could be doubled, leading to considerably less wire resistance and coupling capacitance. However, these improvements are achieved at the expense of certain calibrations. The work in Chapter 4 is readily applicable on long interconnects only when crosstalk from source to destination is

well-known and pre-characterized for different data patterns.

On the other hand, if the aggressor behavior is well-known and once the constraint on the noise margin could be relaxed, the designer could adopt low-swing signaling such as the one in Chapter 6 to minimize the on-chip link energy consumption. Compared to previous low-swing works, new signaling techniques featuring adaptive pre-emphasis showed better bandwidth density combined with low energy consumption. However, adaptive pre-emphasis placed certain timing requirements such that the delay between different taps should match half cycle of the system clock.

This work could be continued in several different directions including the ones described below. First, in scaled technologies, the interaction of gate resistance, wire resistance, and supply regulation could be further investigated. It was pointed out that the wire resistance substantially increases as CMOS technology continues to scale ($\sim 2.5X$ with each technology step as shown in Chapter 1), but once the supply voltage reduces down to near-threshold or sub-threshold region, the gate resistance also greatly increases such that it is comparable or actually higher than the long wire resistance. At some inflection point, the conventional repeater insertion or existing on-chip signaling techniques would not be as efficient as they have been with the nominal supply, and to improve the overall processor and on-chip link performance, the existing circuit techniques should be revisited. Also, if the mid-voltage supply could be regulated from a nominal supply with high efficiency, mid-voltage computation and signaling should be emphasized in a different way.

Second, more circuit techniques could be introduced to bridge the gap between on-chip communication circuits and off-chip communication circuits. Comparing to on-chip communication, off-chip communication tends to have higher power budgets to treat and recover a signal, allowing complicated circuitry with more functionality such as to clock data recovery and decision feedback equalization. However, nowadays several on-chip signaling circuit techniques in the literature including the one introduced in Chapter 6 show similarities with previously proposed off-chip circuit techniques. Smaller gate delay in nanometer CMOS technologies could possibly allow higher integration of transistors for on-chip transmitter and receiver circuits with permissible energy consumption. Also, if general communications schemes such as channel coding or compression are selectively adopted into the well-known on-chip communication channel, better performance or energy-efficiency could be achieved.

BIBLIOGRAPHY

- [1] International technology roadmap for Semiconductors 2003, <http://public.itrs.net>.
- [2] S. Rusu, *et al.*, “A dual-core multi-threaded Xeon processor with 16MB L3 cache,” *Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, pp. 118-119, 2006.
- [3] G. Konstadinidis, *et al.*, “Implementation of a third-generation 16-core 32-thread chip-multithreading SPARC processor,” *Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, pp. 84-85, 2008.
- [4] S. Ramprasad, N. R. Shanbhag, and I. N. Hajj, “A coding framework for low-power address and data busses,” *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 7, no. 2, pp. 212-221, June 1998.
- [5] K. Lee, S.-J. Lee, and H.-J. Yoo, “SILENT: serialized low energy transmission coding for on-chip interconnection networks,” *Proc. of Int. Conf. on Computer-Aided Design (ICCAD)*, pp. 448-451, 2004.
- [6] G. Konstadinidis, *et al.*, “Implementation of a third-generation 1.1-GHz 64-bit microprocessor,” *Journal of Solid-State Circuits (JSSC)*, vol. 37, no. 11, pp. 1461-1469, Nov. 2002.
- [7] P. J. Restle, *et al.*, “The clock distribution of the Power4 microprocessor,” *Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, pp. 144-145, 2002.
- [8] N. Bindal, *et al.*, “Scalable sub-10ps skew global clock distribution for a 90nm multi-ghz IA microprocessor,” *Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, pp. 346-347, 2003.
- [9] R. Kumar, V. Zhyban, and D. M. Tullsen, “Interconnections in multi-core architectures: understanding mechanisms, overheads and scaling,” *Proc. of Int. Symp. on Computer Architecture (ISCA)*, pp. 408-419, 2005.
- [10] J. Seo, D. Sylvester, D. Blaauw, H. Kaul, and R. Krishnamurthy, “A robust edge encoding technique for energy-efficient multi-cycle interconnect,” *Proc. of Int. Symp. on Low Power and Electronics and Design (ISLPED)*, pp. 68-73, 2007.
- [11] J. Seo, H. Kaul, R. Krishnamurthy, D. Sylvester, and D. Blaauw, “A robust edge encoding technique for energy-efficient multi-cycle interconnect,” *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 2010, to appear.
- [12] R. Arunachalam, *et al.*, “Optimal shielding/spacing metrics for low power design,” *Proc. of IEEE Computer Society Annual Symposium on VLSI*, pp. 167-172, 2003.
- [13] S. Wong, *et al.*, “An empirical three-dimensional crossover capacitance model for multilevel interconnect VLSI circuits,” *IEEE Transactions on Semiconductor Manufacturing*, Vol. 13, pp. 219-227, 2000.
- [14] <http://www.eas.asu.edu/ptm/interconnect.html>
- [15] K. Hirose and H. Yasuura, “A bus delay reduction technique considering crosstalk,” *Proc. of DATE*, pp.441-445, 2000.

- [16] M. Khellah, *et al.*, “A Skewed repeater bus architecture for on-chip energy reduction in microprocessors,” *Proc. of International Conference on Computer Design*, pp. 253-257, 2005.
- [17] A. B. Kahng, S. Muddu, and E. Sarto, “Interconnect optimization strategies for high-performance VLSI designs,” *Proc. of International Conference on VLSI Design*, pp. 464-469, 1999.
- [18] H. Kaul, J. Seo, M. Anders, D. Sylvester, and Ram Krishnamurthy, “A robust alternate repeater technique for high performance busses in the multi-core era,” *Proc. of Int. Symp. on Circuits and Systems*, pp. 372-375, 2008.
- [19] C. J. Akl and M. A. Bayoumi, “Reducing interconnect delay uncertainty via hybrid polarity repeater insertion,” *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, Vol. 16, pp. 1230-1239, 2008.
- [20] B. Victor and K. Keutzer, “Bus encoding to prevent crosstalk delay,” *Proc. of Int. Conf. on Computer-Aided Design (ICCAD)*, pp. 57-63, 2001.
- [21] P. P. Sotiriadis, A. Wang, and A. Chandrakasan, “Transition pattern coding: an approach to reduce energy in interconnect,” *Proc. of European Solid-State Circuits Conference (ESSCIRC)*, pp. 348-351, 2000.
- [22] M. Khellah, *et al.*, “Static pulsed bus for on-chip interconnects”, *Symposium on VLSI Circuits Dig. Tech. Papers*, pp. 78-79, 2002.
- [23] H. Deogun, *et al.*, “A dual-V_{dd} boosted pulsed bus technique for low power and low leakage operation,” *Proc. of Int. Symp. on Low Power and Electronics and Design (ISLPED)*, pp. 73-78, 2006.
- [24] H. Kaul, *et al.*, “Design and analysis of spatial encoding circuits for peak power reduction in on-chip buses,” *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, Vol. 13, pp. 1225-1238, 2005.
- [25] A. B. Kahng, *et al.*, “On Switch Factor Based Analysis of Coupled RC Interconnects,” *Proc. of DAC*, pp. 79-84, 2000.
- [26] J. Eble, V. De, D. Wills, and J. Meindl, “Minimum repeater count, size, and energy dissipation for gigascale integration (GSI) interconnects,” *Proc. of Int. Interconnect Technology Conference*, pp. 56-58, 1998.
- [27] H. Partovi, R. Burd, U. Salim, F. Weber, L. DiGregorio, and D. Draper, “Flow-through latch and edge-triggered flip-flop hybrid elements,” *Dig. Tech. Papers on IEEE Int. Solid-State Circuits Conference (ISSCC)*, pp. 138-139, 1996.
- [28] J. Tschanz, S. Narendra, Z. Chen, S. Borkar, M. Sachdev, and V. De, “Comparative delay and energy of single edge-triggered and dual edge-triggered pulsed flip-flops for high-performance microprocessors,” *Proc. of Int. Symp. on Low Power Electronics and Design (ISLPED)*, pp. 147-152, 2001.

- [29] K. Bowman, J. Tschanz, M. Khellah, M. Ghoneima, Y. Ismail, and V. De, "Time-borrowing multi-cycle on-chip interconnects for delay variation tolerance," *Proc. of Int. Symp. on Low Power Electronics and Design (ISLPED)*, pp. 79-84, 2006.
- [30] K. Bernstein, *et al.*, "Design and CAD challenges in sub-90 nm CMOS Technologies," *Proc. of Int. Conf. on Computer-Aided Design (ICCAD)*, pp. 129-136, 2003.
- [31] P. Bai, *et al.*, "A 65nm logic technology featuring 35nm gate lengths, enhanced channel strain, 8 Cu interconnect layers, low-k ILD and $0.57\mu\text{m}^2$ SRAM cell," *Intl. Electron Devices Meeting (IEDM) Technical Digest*, pp. 657-660, Dec 2004.
- [32] M. Anders, N. Rai, R. Krishnamurthy, and S. Borkar, "A transition-encoded dynamic bus technique for high-performance interconnects," *IEEE Journal Solid-State Circuits (JSSC)*, Vol. 38, pp. 709-714, May 2003.
- [33] P. P. Sotiriadis, A. Wang, and A. Chandrakasan, "Transition pattern coding: an approach to reduce energy in interconnect," *Proc. European Solid-State Circuits Conf. (ESSCIRC)*, pp. 348-351, 2000.
- [34] S. Komatsu, M. Ikeda, and K. Asada, "Bus data encoding with coupling-driven code-book method for low power data transmission," *Proc. European Solid-State Circuits Conf. (ESSCIRC)*, pp. 297-300, 2001.
- [35] H. Kaul, D. Sylvester, D. Blaauw, "Performance optimization of critical nets through active shielding," *IEEE Trans. Circuits and Systems - I*, Vol. 51, pp. 2417-2435, Dec 2004.
- [36] K. Nose, and T. Sakurai., "Two schemes to reduce interconnect delay in bi-directional and uni-directional buses," *VLSI Circuits Symp. Digest*, pp. 193-194, 2001.
- [37] A. B. Kahng, S. Muddu, E. Sarto, and R. Sharma, "Interconnect tuning strategies for high-performance ICs," *Proc. Design, Automation and Test in Europe (DATE)*, pp. 471-478, 1998.
- [38] C. J. Akl and M. A. Bayoumi, "Reducing delay uncertainty of on-chip interconnects by combining inverting and non-inverting repeaters insertion," *Intl. Symp. Quality Electronic Design (ISQED)*, pp. 219-224, 2007.
- [39] J. Seo, D. Sylvester, and D. Blaauw, "Crosstalk-aware PWM-based on-chip global signaling in 65nm CMOS," *Symp. VLSI Circuits Dig. Tech. Papers*, pp. 88-89, 2004.
- [40] P. Wang, G. Pei, and E. Kan, "Pulsed wave interconnect," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 12, no. 5, pp. 453-463, May 2004.
- [41] D. Boijort and O. Svanell, "Pulse width modulation for on-chip interconnects," Master Thesis, Linkping University.
- [42] A. J. Joshi and J. A. Davis, "Wave-pipelined multiplexed (WPM) routing gigascale integration (GSI)," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 13, no. 8, pp. 899-910, Aug. 2005.

- [43] V. Venkatraman, and W. Burleson, "An energy-efficient multi-bit quaternary current-mode signaling for on-chip interconnects," *Proc. of Custom Integrated Circuits Conference (CICC)*, pp. 301-304, 2007.
- [44] M. Ghoneima, Y. Ismail, M. Khellah, J. Tschanz, and V. De, "Serial-link bus: a low-power on-chip bus architecture," *Proc. of Int. Conf. on Computer-Aided Design (ICCAD)*, pp. 541-546, 2005.
- [45] K. Lee, S.-J. Lee, and H.-J. Yoo, "SILENT: serialized low energy transmission coding for on-chip interconnection networks," *Proc. of Int. Conf. on Computer-Aided Design (ICCAD)*, pp. 448-451, 2004.
- [46] J. Park, J. Kang, S. Park, and M. P. Flynn, "A 9Gbit/s serial transceiver for on-chip global signaling over lossy transmission lines," *Proc. of Custom Integrated Circuits Conference (CICC)*, pp. 347-350, 2008.
- [47] N. Binkert, R. Dreslinski, *et al.*, "The M5 simulator: modeling networked systems," *IEEE Micro*, vol. 26, no. 4, pp. 52-60, July/Aug. 2006.
- [48] S. Pant and D. Blaauw, "Circuit techniques for suppression and measurement of on-chip inductive supply noise," *European Solid-State Circuits Conference (ESSCIRC)*, pp. 134-137, 2008.
- [49] D. Sylvester and K. Keutzer, "Getting to the bottom of deep submicron II: A global wiring paradigm," *Proc. ISPD*, pp. 193-200, 1999.
- [50] P. Saxena, *et al.*, "The scaling challenge: can correct-by-construction design help?" *Proc. ISPD*, pp. 51-57, 2003.
- [51] T. Takayanagi, *et al.*, "A dual-core 64-bit ultrasparc microprocessor for dense server applications," *IEEE J. Solid-state Circuits*, Jan. 2005.
- [52] A. Nalamalpu, *et al.*, "Boosters for driving long onchip interconnects—Design issues, interconnect synthesis, and comparison with repeaters," *IEEE Trans. on CAD*, Jan. 2002.
- [53] H. Kaul and D. Sylvester, "Transition aware global signaling (TAGS)," *Proc. of Int. Symp. on Quality Electronic Design (ISQED)*, pp. 10-14, 2002.
- [54] H. Huang and S. Chen, "Interconnect accelerating techniques for sub-100-nm gigascale systems," *IEEE Trans. on Very Large Scale Integration (VLSI) Systems*, Nov. 2004.
- [55] R. Ho, *et al.*, "Efficient on-chip global interconnects," *Symp. VLSI Circuits Dig. Tech. Papers.*, pp. 271-274, 2003.
- [56] L. Zhang, *et al.*, "Driver pre-emphasis techniques for on-chip global buses," *Proc. ISLPED*, pp. 186-191, 2005.
- [57] R. T. Chang, *et al.*, "Near speed-of-light signaling over on-chip electrical interconnects," *IEEE J. Solid-State Circuits*, pp. 834-838, May 2003.

- [58] A. P. Jose, *et.al.*, “Near speed-of-light on-chip interconnects using pulsed current-mode signaling,” *Symp. VLSI Circuits Dig. Tech. Papers*, pp. 108-111, 2005.
- [59] J. Seo, P. Singh, D. Sylvester, and D. Blaauw, “Self-timed regenerators for high-speed and low-power interconnects,” *Proc. of Int. Symp. on Quality Electronic Design (ISQED)*, pp. 621-626, 2007.
- [60] P. Singh, J. Seo, D. Blaauw, and D. Sylvester, “Self-timed regenerators for high-speed and low-power global interconnects,” *IEEE Trans. on Very Large Scale Integration (VLSI) Systems*, vol. 16, no. 6, pp. 673-677, Jun. 2008.
- [61] M. Kamon, *et.al.*, “Fasthenry: A multipole-accelerated 3-d inductance extraction program,” *IEEE Trans. on Microwave Theory and Techniques*, Sep. 1994.
- [62] Y. Cao, *et. al.*, “New paradigm of predictive mosfet and interconnect modeling for early circuit simulation,” *Proc. CICC*, pp. 201-204, 2000.
- [63] A. Abdollahi and M. Pedram “A new canonical form for fast boolean matching in logic synthesis and verification,” *Proc. Design Automation Conference*, pp. 379-384, 2005.
- [64] G. Agosta *et al.*, “A unified approach to canonical form-based boolean matching” *Proc. Design Automation Conference*, pp. 841-846, 2007.
- [65] D. Sylvester and K. Keutzer, “Getting to the bottom of deep submicron II: A global wiring paradigm,” *Proc. International Symposium on Physical Design*, pp. 193-200, 1999.
- [66] P. Saxena *et al.*, “Repeater scaling and its impact on CAD,” *IEEE Transactions of Computer-Aided Design of Integrated Circuits and Systems*, Vol. 23, No. 4, pp. 451-463, 2004.
- [67] C. Alpert *et al.*, “Buffer insertion with accurate gate and interconnect delay computation,” *Proc. Design Automation Conference*, pp. 479-484, 1999.
- [68] Y. Ismail and E. Friedman, “Optimum repeater insertion based on a CMOS delay model for on-chip RLC interconnect,” *Proc. International ASIC Conference*, pp. 369-373, 1998.
- [69] A. Nalamalpu and W. Burleson, “Repeater insertion in deep sub-micron CMOS: ramp-based analytical model and placement sensitivity analysis,” *Proc. International Symposium on Circuits and Systems*, pp. 766-799, 2000.
- [70] R. Otten and R. Brayton, “Planning for performance,” *Proc. Design Automation Conference*, pp. 122-127, 1998.
- [71] M. Pedram and N. Bhat, “Layout driven technology mapping,” *Proc. Design Automation Conference*, pp. 99-105, 1991.
- [72] A. Lu *et al.*, “Combining technology mapping with post-placement resynthesis for performance optimization,” *Proc. International Conference on Computer Design*, pp. 616-621, 1998.

- [73] G. Stenz *et al.*, “Performance optimization by interacting netlist transformations and placement,” *IEEE Trans. on Very Large Scale Integration (VLSI) Systems*, Vol. 19, No. 3, pp. 350-358, March 2000.
- [74] IWLS 2005 benchmarks. <http://iwls.org/iwls2005/benchmarks.html>.
- [75] E. Wein, and J. Benkoski, “Hard macros will revolutionize SoC Design,” *EE Design*, August 2004. <http://www.eetimes.com/showArticle.jhtml?articleID=26807055>.
- [76] T. Chen *et al.*, “MP-trees: a packaging-based macro-placement algorithm for mixed-size designs,” *Proc. Design Automation Conference*, pp. 447-452, 2007.
- [77] S. Plaza, I. Markov, and V. Bertacco, “Optimizing non-monotonic interconnect using functional simulation and logic restructuring,” *Proc. International Symposium on Physical Design*, pp. 95-102, 2008.
- [78] R. Shelar, P. Saxena, X. Wang, and S. Sapatnekar, “An efficient technology mapping algorithm targeting routing congestion under delay constraints,” *Proc. International Symposium on Physical Design*, pp. 137-144, 2005.
- [79] M. Moreinis *et al.*, “Logic gates as repeater (LGR) for area-efficient timing optimization,” *IEEE Trans. on Very Large Scale Integration (VLSI) Systems*, Vol. 14, No. 11, pp. 1276-1281, November 2006.
- [80] R. T. Chang, *et al.*, “Near speed-of-light signaling over on-chip electrical interconnects,” *IEEE J. Solid-State Circuits*, vol. 38, no. 5, pp. 834-838, May 2003.
- [81] A. P. Jose and K. L. Shepard, “Near speed-of-light on-chip interconnects using pulsed current-mode signaling,” *Symp. VLSI Circuits Dig. Tech. Papers*, pp. 108-111, 2005.
- [82] L. Zhang, *et al.*, “Driver pre-emphasis techniques for on-chip global buses,” *Proc. Int. Symp. on Low Power Electronics and Design (ISLPED)*, pp. 186-191, 2005.
- [83] D. Schinkel, *et al.*, “A 3-Gb/s/ch transceiver for 10-mm uninterrupted RC-limited global on-chip interconnects,” *IEEE J. Solid-State Circuits*, vol. 41, no. 1, pp. 297-306, Jan. 2006.
- [84] A. P. Jose and K. L. Shepard, “Distributed loss-compensation techniques for energy-efficient low-latency on-chip communication,” *IEEE J. Solid-State Circuits*, vol. 42, no. 6, pp. 1415-1424, June 2007.
- [85] R. Ho, *et al.*, “High speed and low energy capacitively driven on-chip wires,” *IEEE J. Solid-State Circuits*, vol. 43, no. 1, pp. 52-60, Jan. 2008.
- [86] E. Mensink, *et al.*, “A 0.28pJ/b 2Gb/s/ch transceiver in 90nm CMOS for 10mm on-chip interconnects,” *Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, pp. 414-415, 2007.
- [87] N. Tzartzanis and W. Walker, “Differential current-mode sensing for efficient on-chip global signaling,” *IEEE J. Solid-State Circuits*, vol. 40, no. 11, pp. 2141-2147, Nov. 2005.

- [88] B. Kim and V. Stojanovic, "A 4Gb/s/ch 356fJ/b 10mm equalized on-chip interconnect with nonlinear charge-injecting transmit filter and transimpedance receiver in 90nm CMOS," *Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, pp. 66-67, 2009.
- [89] J. Bae, *et al.*, "A 0.6pJ/b 3Gb/s/ch transceiver in 0.18um CMOS for 10mm on-chip interconnects," *Proc. Int. Symp. Circuit and Systems*, pp. 2861-2864, 2005.
- [90] J. Seo, R. Ho, J. Lexau, M. Dayringer, D. Sylvester, and D. Blaauw, "High bandwidth and low energy on-chip signaling with adaptive pre-emphasis in 90nm CMOS," *Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2010, to appear.
- [91] J. Kim, *et al.*, "A 5.6-mW 1-Gb/s/pair pulsed signaling transceiver for a fully AC coupled bus," *IEEE J. Solid-State Circuits*, vol. 40, no. 6, pp. 1331-1340, June 2005.
- [92] L. Luo, *et al.*, "A 3Gb/s AC coupled chip-to-chip communication using a low swing pulse receiver," *IEEE J. Solid-State Circuits*, vol. 41, no. 1, pp. 287-296, Jan. 2006.
- [93] L. Luo, *et al.*, "A 36Gb/s ACCI multi-channel bus using a fully differential pulse receiver," *Proc. Custom Int. Circuits Conf.*, pp. 773-776, 2006.
- [94] M. Hossain, *et al.*, "A 14-Gb/s 32mW AC coupled receiver in 90-nm CMOS," *Symp. VLSI Circuits Dig. Tech. Papers*, pp. 32-33, 2007.
- [95] N. Miura, *et al.*, "A 11Gb/s inductive-coupling link with burst transmission," *Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, pp. 298-299, 2008.
- [96] R. Ho, *et al.*, "Applications for on-chip samplers for test and measurement of integrated circuits," *Symp. VLSI Circuits Dig. Tech. Papers*, pp. 138-139, 1998.