

AN EFFICIENT I/O AND CLOCK RECOVERY DESIGN
FOR TERABIT INTEGRATED CIRCUITS

A DISSERTATION

SUBMITTED TO THE DEPARTMENT OF ELECTRICAL ENGINEERING

AND THE COMMITTEE ON GRADUATE STUDIES

OF STANFORD UNIVERSITY

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

Ming-Ju Edward Lee

August 2001

© Copyright by Ming-Ju Edward Lee 2001

All rights reserved

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

William J. Dally (Principal Advisor)

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Mark A. Horowitz

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Bruce A. Wooley

Approved for University Committee on Graduate Studies:

Abstract

Today in many applications such as network switches, routers, multi-computers, and processor-memory interfaces, the ability to integrate hundreds of multi-gigabit I/Os is desired to make better use of the rapidly advancing IC technology. With many high speed I/Os integrated on a chip, the wire count, component count, and power budget of a system can be significantly reduced, allowing for both reduced costs and expanded capability. Although previously published designs have achieved multi-gigabit bandwidth per channel, the area and power consumption are too large to make terabit integrated circuits feasible.

In this thesis, an efficient I/O and clock recovery design is presented. In a 0.25- μm CMOS technology, the circuits operate at 4-Gb/s, occupy 0.3- mm^2 , and dissipate 180-mW on a 2.5-V supply. Keys to achieving these numbers are a set of circuit techniques applied to the transmitter, the receiver, and the timing circuits. In addition to power and area, resistance to digital noise sources is also critical to enable integration in a VLSI environment. A low-swing input-multiplexed transmitter is used to serialize low-speed data without the speed limitation of traditional input-multiplexing or the area and power penalty of output multiplexing. Since this I/O is intended to be part of a large digital system, pre-emphasis filter is used to drive a backplane with 40-in of PCB trace and two connectors (or other media with a similar loss). A mathematical analysis of the channel and the filter is presented, showing that a 2-tap FIR filter is adequate in such a case. A

capacitively trimmed sense amplifier is used to cancel the receiver offset without sacrificing the speed. This technique increases both the voltage and timing margins, allows small receivers to be built, decreases the power consumption, and increases the input bandwidth. A supply-regulated inverter delay line is used to implement the multi-phase delay-locked loop. Compared to source-coupled delay lines, it dissipates less power and is more portable and easier to design. By regulating the delay line supply with a voltage regulator, the jitter is also significantly reduced. Finally, the Sidiropoulos dual-loop architecture is adopted for the clock recovery. A current-mirror circuit topology is used for both the phase multiplexer and the phase interpolator to achieve a high bandwidth and a good phase linearity. This circuit topology helps the overall timing budget by reducing the receiver clock jitter and dithering. The above circuit techniques were incorporated into two test prototypes, whose experimental data will be described in detail.

Acknowledgments

The name of Professor Bill Dally first came into my life when I was browsing through the MIT faculty web page in my last year as a Berkeley undergraduate student. Being a resident of the San Francisco Bay Area for almost 7 years, I was eager to move to the unfamiliar east coast to attend the graduate school at MIT. I sent an E-mail off to Bill expressing my interest in joining his research group. He replied to inform me that he was moving his whole research group to Stanford. Although I found it a little odd, I didn't give it much thought and went on to look for other faculty advisors at MIT.

But little did I know then that I would eventually give in to the lovely weather of the Bay Area and let go of my venturesome spirit, almost out of a whim, the night before I was supposed to make a decision on which school to attend. And little did I know that I would eventually end up at Bill's research group at a different school. I went around a full circle and finally arrived at the best decision I've ever made in my life. I often reflect upon the past 4 years of study under Bill's guidance and realize again and again how blessed I have been that good fortune always manages to find me even when I am looking elsewhere. No one could have asked for a better advisor. I am loving and enjoying every minute of my work because Bill showed me how wonderful research can be. He motivated me by working harder and yet exuberating more energy and optimism than anybody I've ever seen (rumor has it that he does not sleep). The example and the standard he sets will be with me always.

I would also like to thank Professor Mark Horowitz for his encouragement and for sharing his vast knowledge with me. Mark is a pioneer in this research field, as is evidenced by the many references authored by him, and just knowing that he's here is enough to convince me that Stanford is indeed the best school to be.

I am grateful to Professor Bruce Wooley for being a reader of this dissertation and a member of my oral examination committee. Bruce was my brother-in-law's Ph.D. advisor, and I've heard so much about him since I was a high school student. His involvement really gives this dissertation and my Ph.D. degree a special meaning.

Whenever I am looking for new ideas or advice on my research, I always go to Professor John Poulton. John always forces me to look at things from a totally different perspective. He is a truly creative circuit designer, and I hope I inherited some of that creativity from him over the years.

I feel very proud and fortunate to have been part of a such talented research group: Andrew Chang, Patrick Chiang, Mattan Erez, Steve Greenwood, Sarah Harris, Ujval Kapasi, Steve Keckler, Bruce Khailany, Steve Lacy, Whay Sing Lee, Perter Mattson, John Owens, Li-Shiuan Peh, Scott Rixner, Kelly Shaw, and Brian Towles. Shelley Russell and Pamela Elliot deserve special thanks for their kind assistance over the years and for keeping this research group away from complete chaos.

I've learned so much through interactions with Professor Mark Horowitz's research group. In particular, I would like to thank Ken Chang, Dean Liu, Jaeha Kim, Stefanos Sidosopoulos, Bill Ellersick, Ken Yang, Gu-Yeon Wei, Evelina Yeung, and Ron Ho for their help and intellectual stimulation.

My teammates at Velio Communications, John Edmonson, Ramin Farjad, Dan Hartman, Trey Greer, David Huang, Hiok-Tiaq Ng, Rohit Rathi, Ramesh Senthinathan, and Teva Stone, are among the best circuit designers I have ever met. I joined Velio because of the lure of the start-ups and the opportunity to turn research into products, but stayed because it's just purely fun and fulfilling to work there.

I want to thank my fiancée Lisa for her love and support over the years. I have become a much more content and happier person since I met her. Her presence gives me the will to go through the difficulties in my graduate years.

And of course, I would not have been able to carry myself all the way through the Ph.D. degree without my parents, sister, and brothers. Growing up, I never had to worry about anything but my education, and that's a luxury not many people get to have. Even though we are all scattered around different parts of the world, I am sure that deep, unique bond will exist in our family forever.

Table of Contents

Abstract	v
Acknowledgments	vii
Table of Contents	xi
List of Figures	xiii
List of Tables	xvii
Chapter 1 Introduction	1
1.1 Contributions	2
1.2 Organization.....	5
Chapter 2 Background	9
2.1 Transmitter	10
2.2 Receiver	12
2.3 Timing Circuits	16
2.4 Equalization Filter.....	19
2.5 Summary.....	22
Chapter 3 System Overview	23
3.1 System Architecture.....	24

3.2	Channel Analysis and Equalization	25
3.3	Timing Convention	34
3.4	Signaling Convention	38
3.4.1	Differential vs. Single-Ended Signaling	38
3.4.2	Binary vs. Multi-Level Encoding	40
3.4.3	Uni-directional vs. Simultaneous Bi-directional	43
3.5	Summary	47
Chapter 4	Transmitter	49
4.1	Architecture	49
4.2	Circuit Implementation	53
4.3	Summary	56
Chapter 5	Receiver	59
5.1	Architecture	60
5.2	Circuit Implementation	61
5.3	Summary	68
Chapter 6	Timing Circuits	71
6.1	Multi-Phase Delay-Locked Loop.....	71
6.2	Clock Recovery.....	81
6.2.1	Architecture	81
6.2.2	Circuit Implementation	83
6.3	Summary	88
Chapter 7	Experimental Results	91
7.1	Prototypes	91
7.2	Measurement Results	95
7.3	Summary	105
Chapter 8	Conclusion	107
8.1	Future Work	109
Bibliography	109

List of Figures

Figure 1.1: A basic I/O design.....	2
Figure 2.1: Timing diagram for using multiple clock phases to perform multiplexing. .	11
Figure 2.2: A differential current mode driver.	11
Figure 2.3: A voltage mode driver.	12
Figure 2.4: A demultiplexing receiver architecture.....	13
Figure 2.5: A gate-isolated sense amplifier.....	14
Figure 2.6: A pass-gate sampler.	14
Figure 2.7: A current integrating receiver.	15
Figure 2.8: DLL (top diagram) and PLL (bottom diagram) based multi-phase generation. 16	16
Figure 2.9: DLL (top diagram) and PLL (bottom diagram) based clock recovery.	17
Figure 2.10: A dual-loop clock recovery scheme. The left side shows the architecture and the right side shows a phase interpolator implementation.	18
Figure 2.11: Input and output eye diagram before and after a 1-m, 7-mil, 0.5-oz. PCB trace. 19	19
Figure 2.12: Output eye diagram after a 1-m, 7-mil, 0.5-oz. PCB trace with a two tap equalization filter.	20
Figure 2.13: Dally & Poulton's transmitter architecture.....	21
Figure 2.14: Analog current summing transmitter FIR filter.	21
Figure 3.15: System Architecture of the 4-Gb/s transceiver.	24
Figure 3.16: A typical application of high speed serial links.....	25
Figure 3.17: Frequency response of a 1-m PCB trace.....	26
Figure 3.18: Circuit model of a typical backplane channel utilizing HSPICE's W-element. 28	28
Figure 3.19: Simulated S12 response of Figure 3.18.	28
Figure 3.20: A model of the pre-emphasis filter and the channel.	29
Figure 3.21: Simulated pulse response of Figure 3.18. The bottom plot is a zoomed-in version of the top plot.	30

Figure 3.22: Effect of the ISI on the bit error rate.....	32
Figure 3.23: Effect of the ISI on the bit error rate. The different curves correspond to different levels of Gaussian noise.....	33
Figure 3.24: Effect of the ISI on the bit error rate for a long backplane channel.....	35
Figure 3.25: An abstract eye diagram showing the timing budget.....	35
Figure 3.26: A bundled closed-loop timing system.	37
Figure 3.27: A per-line closed-loop timing system.....	37
Figure 3.28: BER versus the number of filter taps for 4-PAM and 2-PAM signal encoding. The symbol rate of 4-PAM is half that of 2-PAM.	41
Figure 3.29: Simultaneous bi-directional signaling.	43
Figure 3.30: Sum of the magnitude of ISI versus the number of filter taps for unidirectional (PAM2) and bi-directional (BI).	44
Figure 3.31: Near-end and far-end signal of the sample channel.....	45
Figure 3.32: Simultaneous bi-directional signaling waveform without channel loss.	46
Figure 3.33: Simultaneous bi-directional signaling waveform with channel loss.	46
Figure 4.1: Output-multiplexed transmitter architecture.....	50
Figure 4.2: CMOS gate based input-multiplexed transmitter architecture.	52
Figure 4.3: Minimum achievable bit time the configuration in Figure 4.2.	52
Figure 4.4: Transmitter circuit implementation.	54
Figure 4.5: Effect of bit-time on pulse amplitude closure for the 4:1 pseudo-NMOS multiplexer of Figure 4.4.	55
Figure 4.6: Transmitter resynchronization circuit.....	56
Figure 4.7: Two tap transmitter equalization filter is implemented with analog current summing.....	57
Figure 5.1: Receiver architecture.	61
Figure 5.2: Receive sense amplifier with static offset trimming.....	62
Figure 5.3: Trimmed offset versus the calibration value for the capacitively trimmed sense amplifier.	64
Figure 5.4: Trimmed offset versus the calibration value for the capacitively trimmed sense amplifier with different input common-mode levels.	64
Figure 5.5: Aperture time of the capacitively trimmed sense amplifier.....	65
Figure 5.6: Receiver hysteresis versus clock period.	66
Figure 5.7: A second stage of StrongArm latch is inserted to reduce the hysteresis of the input.	67
Figure 5.8: Receiver sensitivity versus clock period.....	68
Figure 5.9: Receiver resynchronization circuit.	69
Figure 6.1: Supply-regulated inverter delay-locked loop architecture.....	72
Figure 6.2: Source-coupled differential delay element.	72
Figure 6.3: Source-coupled delay element with two-element PMOS load and replica-bias. 73	
Figure 6.4: Phase-only comparator employed in the multi-phase DLL.....	75
Figure 6.5: Timing diagram showing incorrect use of a PFD in a DLL.	75
Figure 6.6: Charge pump employed in the multi-phase DLL.	76
Figure 6.7: Linear voltage regulator employed in the multi-phase DLL.	78
Figure 6.8: Simplified model of the linear voltage regulator.	78
Figure 6.9: Level shifter employed in the multi-phase DLL.....	80

Figure 6.10: Simulated jitter due to a 10% supply pulse with 100-ps rise time.....	81
Figure 6.11: Clock recovery architecture.....	82
Figure 6.12: Phase controller architecture.....	83
Figure 6.13: Peripheral loop interpolator.....	84
Figure 6.14: Phase position vs. the interpolation step for current-mirror interpolator.....	84
Figure 6.15: Tri-state inverter based interpolator.....	85
Figure 6.16: Phase position vs. the interpolation step for tri-state inverter based interpolator.....	85
Figure 6.17: Phase multiplexer employed in the peripheral loop.....	87
Figure 6.18: Tri-state inverter based phase multiplexer.....	87
Figure 6.19: PAC for current-mirror phase multiplexer (CM) and tri-state inverter phase multiplexer (TRI).....	88
Figure 7.1: Die photomicrograph of the first prototype chip.....	92
Figure 7.2: Die photomicrograph of the second prototype chip.....	92
Figure 7.3: On-chip noise generator and noise monitor.....	93
Figure 7.4: Experiment setup of the serial link prototype.....	93
Figure 7.5: Differential eye diagram at the transmitter output.....	95
Figure 7.6: Overlap of the bit pattern in Figure 7.5 to show the effective margin. The rectangle shown in the middle is 100-mV by 170-ps.....	95
Figure 7.7: Differential eye diagram after 1m of PCB trace without equalization.....	96
Figure 7.8: Differential eye diagram at the transmitter output with equalization.....	96
Figure 7.9: Differential eye diagram after 1m of PCB trace with equalization.....	96
Figure 7.10: Overlap of the bit pattern in Figure 7.9 to show the effective margin. The margin rectangle shown in the middle is 100-mV by 120-ps.....	97
Figure 7.11: Jitter histogram of the differential transmitter output.....	97
Figure 7.12: Jitter histogram of the differential transmitter output with 1-MHz 200-mV p- p pulses superimposed on the supply.....	98
Figure 7.13: Jitter histogram of the receiver sampling clock with automatic phase control turned on (for input data of Figure 7.5).....	98
Figure 7.14: PMOS open-drain driver for the on-chip clock signals under observation.....	99
Figure 7.15: Jitter histogram of the receiver sampling clock with automatic phase control turned off (for input data of Figure 7.5).....	99
Figure 7.16: Jitter histogram of the receiver sampling clock with automatic phase control turned on (for input data of Figure 7.9).....	100
Figure 7.17: Jitter histogram of the receiver sampling clock with automatic phase control turned on and with 1-MHz 200-mV p-p pulses superimposed on the supply.....	100
Figure 7.18: Receiver single-ended swing versus clock position window. The PASS region has a BER 10^{-12}.....	101
Figure 7.19: Phase position versus the phase step of the clock recovery phase adjustment over a full clock cycle.....	102
Figure 7.20: Phase step size. The numbers at each interval indicate the core DLL phase interval. For example, 0-1 indicates the phase interval between $0\times$ and $45\times$... 102	102
Figure 7.21: Layout of complement phases of interpolators.....	103
Figure 7.22: Power versus bit rate for the transceiver with minimum operating supply.....	

List of Tables

Table 1:	RLGC parameters for the PCB trace in HSPICE	27
Table 2:	Equalization tap weight calculated by the least square method.....	34
Table 3:	Worst case timing budget for our I/O system.	38
Table 4:	Receiver offset breakdown	63
Table 5:	Test chip performance summary.	105

Chapter 1

Introduction

The performance of many digital systems today is limited by the interconnection bandwidth between chips, boards, and cabinets. Although the processing performance of a single chip has increased dramatically since the inception of the integrated circuit technology, the communication bandwidth between chips has not enjoyed as much benefit. Most CMOS chips, when communicating off-chip, drive unterminated lines with full-swing CMOS drivers and use CMOS gates as receivers. Such full-swing CMOS interconnect must ring-up the line, and hence has a bandwidth that is limited by the length of the line rather than the performance of the semiconductor technology. Thus, as VLSI technology scales, the pin bandwidth does not improve with the technology, but rather remains limited by board and cable geometry, making off-chip bandwidth an even more critical bottleneck.

Recently described I/O circuits have increased the absolute I/O bandwidth by an order of magnitude to the Gb/s range [6] [7] [8]. More importantly, they have put this bandwidth on the semiconductor technology-scaling curve by signaling with the incident




Figure 1.1: A basic I/O design.

wave from the transmitter rather than ringing up the line. Figure 1.1 shows an example I/O system. To achieve incident-wave signaling, these circuits use point-to-point interconnect over terminated transmission lines. Low-swing drivers, as opposed to full swing CMOS drivers, are used to minimize power and reduce self-induced noise in the system. On the receive side, inverters are replaced by sensitive receive amplifiers (often clocked regenerative amplifiers) to reduce the required signal swing and achieve a higher bit rate. Precision timing circuits based on delay-locked loops (DLLs) or phase-locked loops (PLLs) are employed in these systems since a critical limitation of the achievable speed is timing accuracy. In cases where significant channel distortion occurs, signaling rate is still limited by the media. Equalization is incorporated in such cases to correct for the distortion and remove this restriction.

1.1 Contributions

A key remaining problem with high-speed I/Os is reducing the area and power of these circuits to enable very high levels of integration. To relieve the pin-bandwidth bottleneck of modern VLSI chips used for network switch fabrics, routers, and CPU-memory interfaces, hundreds of these high-speed I/Os must be integrated on a single chip. A substantial number of the pins on such chips need to use high-speed signaling, not just a few special pins. Besides power and area, an additional requirement for large scale

integration of high-speed I/Os is noise immunity, particularly immunity to power supply noises. In this thesis, we look into these design requirements and describe circuit techniques to improve them [1] [2] [3].

On the transmitter side, a fast multiplexer is used to serialize on-chip low-speed data into a higher speed bit stream. A low-swing input-multiplexed architecture is used to achieve 4:1 multiplexing with $< 2\tau_4$ bit time, where τ_4 is the fanout of 4 inverter delay. Previous implementations use an output-multiplexed architecture where multiple copies of the output drivers, each sized large enough to drive signals off chip, are placed and multiplexed directly at the transmitter output, where it is connected to a 50- Ω transmission line, to achieve this level of performance [6] [8]. We move the data multiplexing to the input of the output driver and rely on swing reduction to attain the required performance while requiring less area and power. It also improves signal integrity by producing less capacitive load at the output and improving the efficiency of transmitter termination.

For channels which have significant frequency-dependent attenuation, data need to be filtered, usually with a finite-impulse-response (FIR) filter, to be received reliably. Previous designs rely on eye-diagram simulations to provide insights into the filter requirement for a particular channel. Although this method allows designers to get an idea of what the channel output looks like for a given filter configuration, it lacks the ability to quickly quantize the trade-offs between different configurations. In this thesis, we show a mathematical analysis which quantizes the bit error rate improvement with the number of FIR filter taps. This analysis is used to show that a two-tap pre-emphasis filter is sufficient for a backplane channel with 40-in of PCB trace and two connectors.

One of the major drawbacks of previous receiver designs is that they operate with uncanceled input voltage offset [7] [8] [10] [14] [16]. This receiver input offset significantly degrades the timing and voltage margin of the system. In this thesis we introduce a capacitive offset trimming method which reduces the offset to ~ 8 -mV while only degrading the aperture time requirement of the receiver by 6% of the bit time. Besides improving the system margin, cancelling offset also saves power and area by requiring lower signal swing and smaller receivers.

The maximum achievable bit rate in high-performance I/Os is often limited by the timing uncertainty, which is mostly caused by the timing circuits. Previous multi-gigabit I/O designs use source-coupled delay elements to implement either the voltage-controlled delay line (VCDL) in a delay-locked loop (DLL) or the voltage-controlled oscillator (VCO) in a phase-locked loop (PLL) [6] [8] [11]. This type of delay element is mainly used for its low sensitivity to the power supply noise. However, compared with a simple CMOS inverter delay element, it draws significantly more power due to static current consumption. The drawback of a CMOS inverter delay element is that it is much more sensitive to power supply noises. In this design, we use a supply-regulated CMOS inverter based delay-locked loop for the multi-phase generation and clock recovery. The delay line is regulated with a linear voltage regulator to simultaneously achieve supply noise rejection and delay variation. The power saving compared to a source-coupled delay element based delay line is estimated to be 30% for 4 phases and 60% for 8 phases.

The dual-loop clock recovery architecture first described by [12] is adopted in this design. We implement the phase multiplexer and phase interpolator with a current mirror circuit topology to obtain a high bandwidth and a good interpolation linearity. This topology helps the overall timing budget by reducing the receiver clock jitter and dithering.

With these techniques, we were able to construct a compact and low-power I/O at 4-Gb/s with 0.3-mm² of area and 180-mW of power on a 2.5-V supply and in a 0.25- μ m CMOS technology. For reference, the smallest power and area for a similar speed (3.5-Gb/s) and in the same generation of CMOS technology are \sim 300-mW and 0.6-mm² [13]. Besides power and area efficiency, it also exhibits good immunity to power supply noises. With a 200-mV supply noise generated on-chip, the transceiver operates at speed for at least a day (BER < 10⁻¹⁴) with only 50-mV of differential swing. It is now not only feasible but economical to construct a terabit integrated circuit with these I/O circuit techniques. For example, to achieve an aggregate 1-Tb/s I/O bandwidth requires 125 copies of our I/O, 22-W of power, and 37-mm² of area.

1.2 Organization

To provide a background on the status of I/O research at the onset of this research, Chapter 2 briefly describes the previous work. Chapter 3 gives a brief description of the overall I/O architecture and discusses some of the system level design choices to provide a high level view of this design. Much emphasis is given to the analysis of transmitter equalization according to a backplane channel model in HSPICE. Some of the important signaling conventions are discussed. In particular, the pros and cons of single-ended versus differential signaling, binary versus multi-level signal encoding, and uni-directional versus simultaneous bi-directional signaling are reviewed. A brief overview of timing budget and timing convention of the system is also given in this chapter.

Chapter 4 covers the design of the transmitter. A 4:1 multiplexing scheme is used to serialize low-speed parallel data. We first briefly describe existing multiplexing schemes and their limitations. This is followed by a detailed description of our low-swing input-multiplexed design. Transmitter pre-emphasis is realized by replicating the output driver for each tap and summing the current directly at the output.

Chapter 5 presents the design of the receiver. Input data are demultiplexed directly at the input using multiple sense amplifiers. We introduce a capacitive trimming method which reduces the offset of these sense amplifiers to below 8-mV without significant degradation to the aperture time. It increases both the timing margin and the voltage margin and saves power and area by requiring a smaller swing and a smaller receiver.

Chapter 6 describes the timing circuits. A supply-regulated inverter delay line is used to achieve low power consumption and good supply noise rejection simultaneously. We use a dual-loop clock recovery architecture described in [12]. Both the phase multiplexer and the phase interpolator are implemented using a current mirror circuit topology to achieve a high bandwidth and a good phase linearity.

Chapter 7 presents the measurement results from the test prototype fabricated in a 0.25- μm CMOS technology. The experimental setup is first described. We then present measurements of eye diagrams with and without equalization, transmit and receive clock

jitter, receiver timing and voltage window with and without offset cancellation, clock recovery interpolator linearity, plesiochronous frequency tolerance, minimum swing required for BER $< 10^{-14}$, and finally the power consumption breakdown of the I/Os. Chapter 8 concludes this thesis.

Chapter 2

Background

In Chapter 1 we briefly introduced the important features of a state-of-the-art I/O design. In contrast to traditional unterminated CMOS signaling, modern high-performance I/Os use terminated incident wave signaling instead of ringing up the line. Instead of a bus with long stubs, point-to-point interconnects or bus with short stubs have been adopted. These changes largely remove the bit rate limitation due to transmission line reflections and put the signaling speed back on the semiconductor scaling curve like the rest of the integrated circuits.

Merely making the I/O speed scale with the process technology is not enough to satisfy the growing I/O bandwidth demand of ASICs. To push the I/O speed to the maximum, researchers in the past years have introduced innovations both on the circuit level and on the architectural level to all three major blocks of an I/O system, namely the transmitter, the receiver, and the timing circuits. As a result, many designs with signaling rate in the multi-gigabit range have been demonstrated [5] [6] [7] [8].

The first published gigabit serial link design in CMOS is the BULL Serial Link [5]. Many design concepts, such as multi-phase serialization and deserialization, were introduced in this work. Later, [6] [7] [8] pushed the bit rate by a factor of 4 to 4-Gb/s in

the same generation of CMOS technology. Techniques such as transmitter pre-emphasis and output-multiplexing transmitter were introduced in these works to increase the bit rate and improve the signal integrity of the link through a lossy channel. Many serial link designs have been published since these original works. However, there has been very little effort in pushing the bit rate and reducing the power consumption simultaneously to achieve a high level I/O integration. The end result is that, at the onset of this research, although multi-Gb/s designs in CMOS have been demonstrated, only a small number can be integrated on a single chip before the power and area budget explodes.

In this chapter, we briefly describe some of the relevant prior work. Section 2.1 starts with the transmitter. This is followed by the receiver in Section 2.2 and the timing circuits in Section 2.3. For extremely lossy channels, signaling rate is still limited by the transmission media. Section 2.4 describes equalization filter designs which overcome this limit. Finally, Section 2.5 summarizes this chapter.

2.1 Transmitter

Due to limitations of the timing circuits and clock distribution, the on-chip clock period often cannot be made below $6 - 8\tau_4$ (1 – 1.3-GHz in 0.25- μm CMOS technology) [24] [25]. When a smaller bit time is desired, a multiplexer that takes low-speed parallel data and serializes them into a high-speed serial data is required at the transmitter [5] [6] [8]. Because the clock frequency is limited, this is usually done with multiple clock phases, which can be generated easily from a ring oscillator or a delay line. Figure 2.1 shows the timing diagram of how this is done. Phases ϕ_i and $\bar{\phi}_{i+1}$ are AND'd to produce a reference pulse for a bit time. In the simplest and most common system, the transmitter performs 2:1 multiplexing on both edges of the clock [10] [14]. This 2:1 multiplexing improves the bit time to $3 - 4\tau_4$. To increase the bit rate further, the degree of multiplexing must be increased [6] [8]. Because of the high fan-in, the multiplexer is usually the speed bottleneck. Chapter 4 describes different strategies for achieving a targeted $2\tau_4$ bit time (4-Gb/s in 0.25- μm CMOS technology) with 4:1 multiplexing and discusses the pros and cons of each approach.




Figure 2.1: Timing diagram for using multiple clock phases to perform multiplexing.




Figure 2.2: A differential current mode driver.

Another critical design aspect of the transmitter is the output driver. Figure 2.2 and Figure 2.3 show a differential current-mode driver and a differential voltage-mode driver, respectively. A current-mode driver acts as a high-impedance current source. The signal swing is adjusted by varying the amount of current it sinks from the channel. By contrast, a voltage-mode driver acts as low-impedance voltage source. The signal swing is adjusted by varying the supply voltage of the output driver. In order for a voltage-mode driver to act as a voltage, its output impedance must be low. This generally requires large output




Figure 2.3: A voltage mode driver.

transistors and puts significant loading at the output. A current-mode driver, on the other hand, does not suffer from this shortcoming since it only requires its output transistors to completely switch the current from one side to another with a given input swing while remaining saturated during operation. In order to vary the output swing of a voltage-mode driver with a fixed output impedance, a voltage regulator is required to vary its supply. This is expensive and difficult in contrast to simply varying the bias current of a current-mode driver with a servomechanism.

2.2 Receiver

Like the transmitter, the clock period at the receiver is limited to about $6 - 8\tau_4$. The serial data going into the receiver must be demultiplexed first before they can be processed. Figure 2.4 shows a typical implementation where the serial data are demultiplexed directly at the input [7] [8]. The front-end sense amplifiers sample the input on evenly spaced clock phases. Since most sense amplifier designs require a much larger cycle time than aperture time¹, the bit rate can be significantly increased with this architecture.




Figure 2.4: A demultiplexing receiver architecture.

Figure 2.5 shows a gate-isolated regenerative sense amplifier. It was originally used as a flip-flop in the StrongArm microprocessor [7] [15]. The output nodes are pre-charged high when the clock is low. Positive feedback produces a differential CMOS value at the output on the rising edge of the clock. The NMOS connected between node a and b reduces the aperture time by shorting a and b once they fall below $V_{dd} - V_t$, ending the influence of the input on the cross-coupled inverters. The aperture time with this topology is on the order of $0.2\tau_4$.

Alternatively, one could use a pass-gate to sample the input, as shown in Figure 2.6 [8] [16]. The sampler is followed by a regenerative sense amplifier operating on an opposite clock phase to produce a CMOS value. The required aperture time with this topology is on the order of $0.3\tau_4$ for NMOS-based samplers and $0.6\tau_4$ for PMOS-based samplers and is generally larger than a gate-isolated sense amplifier. Except for cases where an analog value is required (for example when a receiver filter is implemented), a

1. Aperture time is defined as the minimum timing window in which the signal must be larger than the receiver sensitivity (including offset) for correct operation.




Figure 2.5: A gate-isolated sense amplifier.




Figure 2.6: A pass-gate sampler.

gate-isolated amplifier, which combines the tasks of sampling and detection, should be used for better aperture time.

The above two receiver designs are based on *sampling* where the input is only sampled inside a very narrow timing window. This approach gives good rejection of timing noise and low-frequency voltage noise if the sampling clock is placed optimally at the center of the eye. In the presence of high frequency noise, however, it is advantageous

to integrate the signal over the bit cell. Figure 2.7 shows a current integrating receiver originally described in [14]. The front-end stage integrate the input when clk is low, and the second stage sense amplifier samples the integrator output on the rising edge. It has the advantage of rejecting high frequency noise that tends to average out over a bit time. One example where an integrating receiver works well is simultaneous bi-directional signaling [17]. This scheme allows sending signals in both directions on a single channel at the same time by subtracting the transmitted signal from the channel before the receiver makes a decision. Because the timing of the subtraction circuit often does not match that of the actual transmitter, using a sampling receiver is unreliable as the sampling instant might coincide with a transient event. An integrating receiver is more robust in such situation as it looks at the whole bit time instead of a particular instant.

Integrating the signal over a bit time has the disadvantage of reducing immunity to low-frequency noise and timing noise. In the presence of phase offset and timing jitter, for example, an integrating receiver would partially integrate over the adjacent bits. The optimal integrating function, $\psi(t)$, is the one which only integrates outside the timing uncertainty and gives more weight when the signal is large, as expressed by

$$\psi(t) \otimes \phi(t) = s(t) \quad (2.1)$$

where $\phi(t)$ is the probability density function of the timing jitter and $s(t)$ is the pulse response of the channel [21]. However, the exact $\psi(t)$ is difficult to implement in reality.




Figure 2.7: A current integrating receiver.

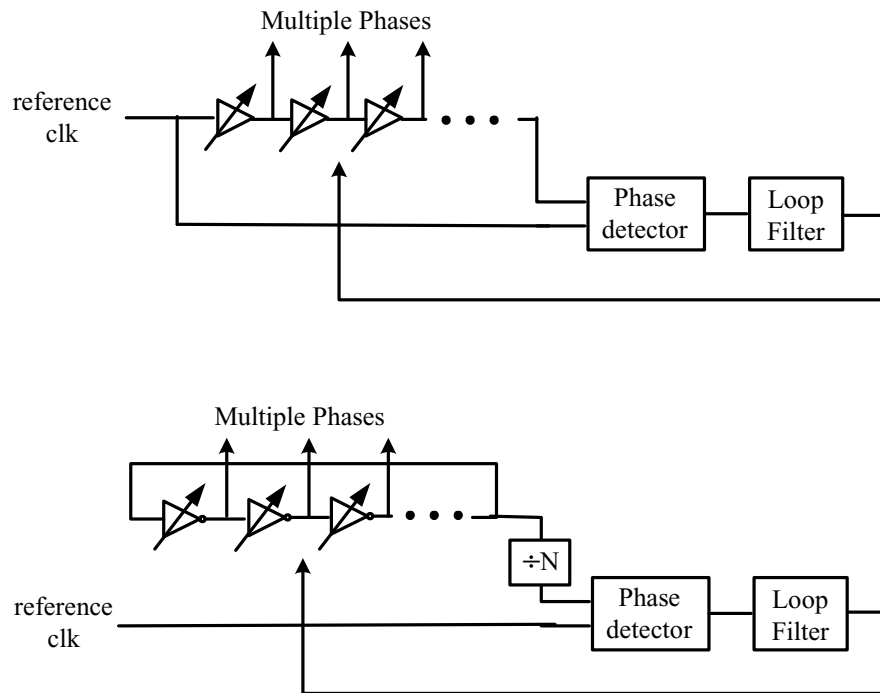


Figure 2.8: DLL (top diagram) and PLL (bottom diagram) based multi-phase generation.

The distinction between an integrating receiver and a sampling receiver becomes vague as bit rate increases since a sampling receiver is really an integrating receiver with a integrating period equal to its aperture time.

2.3 Timing Circuits

As mentioned above, to achieve a signaling rate beyond the frequency limitation of the on-chip clock signal, multiple clock phases are required to perform multiplexing at the transmitter and demultiplexing at the receiver. The multi-phase generation is most often done with either a delay-locked loop (DLL), or a phase-locked loop (PLL) if frequency synthesis is required. Figure 2.8 shows a high level diagram of a DLL and a PLL. For the DLL, two ends of a variable delay line are compared and locked to the same phase using a phase detector and an averaging loop filter. The intermediate nodes can then be tapped off to generate multiple phases. For the PLL, the clock signal generated by a variable oscillator is locked to a multiple of the reference clock frequency. Again the intermediate

nodes of the ring-oscillator can be tapped off to generate multiple phases. Chapter 6 will go into the details of the previous circuit implementations and compare them with our approach.

Another important functionality of the timing circuits is clock recovery, a process in which the receiver clock is aligned to the center of the data signal for maximum timing margin. Figure 2.9 shows two traditional clock recovery schemes based on a DLL and a PLL. The architecture is similar to multi-phase generation. A phase detector measures the instantaneous phase error between the output clock and the reference signal. The reference signal can be either the incoming data¹ or a clock sent by the source. A loop filter averages these measurements, and a clock adjustment is made using either a variable delay line (DLL) or a variable oscillator (PLL).

Assuming a clean clock source is available, a DLL typically produces less clock jitter compared to a PLL since it does not recirculate phase error. In a PLL, any jitter

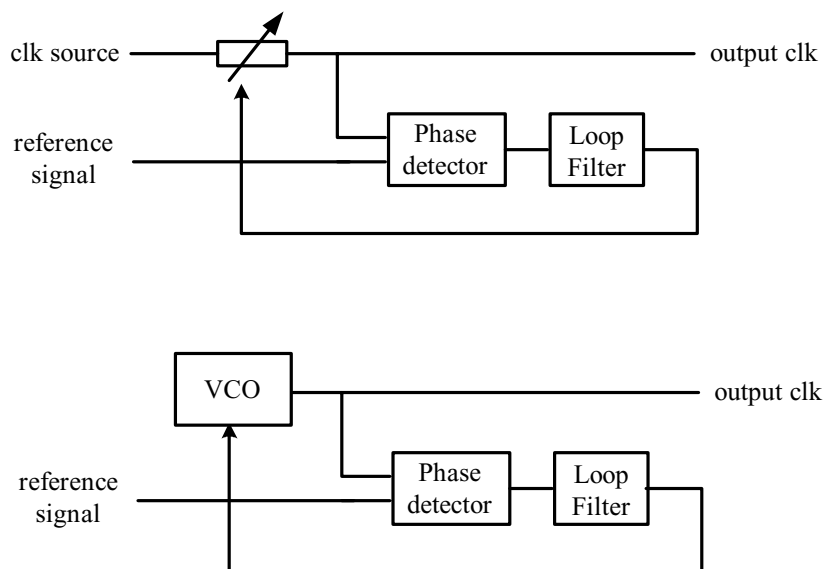


Figure 2.9: DLL (top diagram) and PLL (bottom diagram) based clock recovery.

1. When the incoming data are used directly as the reference, edge detection circuits are required to extract the timing information.

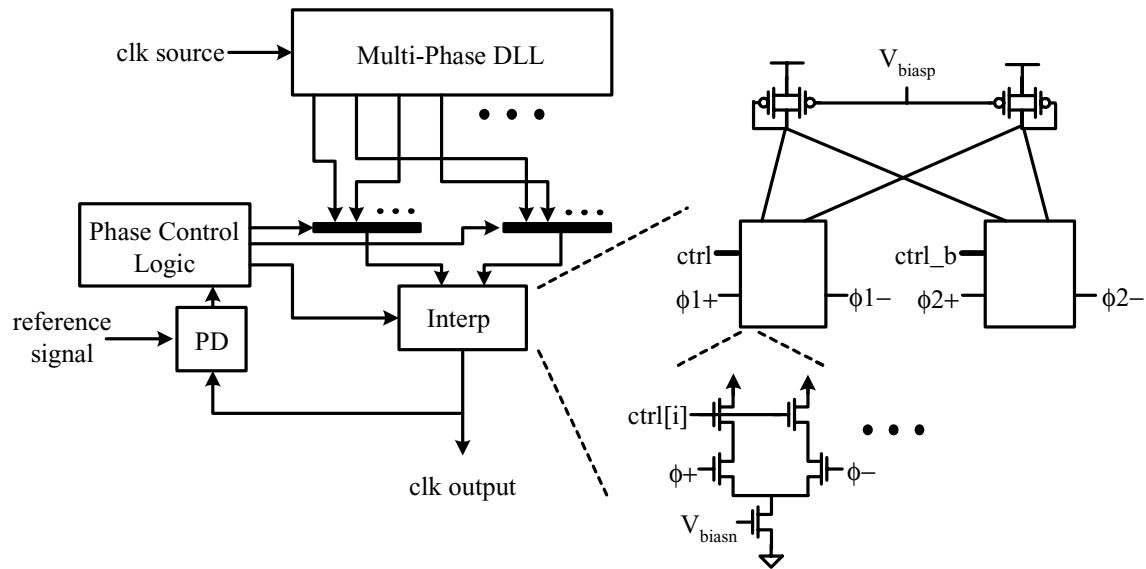


Figure 2.10: A dual-loop clock recovery scheme. The left side shows the architecture and the right side shows a phase interpolator implementation.

introduced during a cycle of operation is fed back through the ring oscillator on the next cycle. [18] shows that delay-element jitter in a PLL is amplified by an *accumulation factor* inversely proportional to the control loop bandwidth. However, the loop bandwidth often needs to be low in order to filter out the noisy reference signal. There is thus a conflicting requirement between reducing the delay-element jitter and reducing the jitter transfer. A DLL, of course, does not have this trade-off since it does not accumulate jitter.

Although a DLL implementation produces less clock jitter, it has several disadvantages. In cases where frequency synthesis in addition to phase alignment is required, however, a PLL is required. Also, by virtue of its phase accumulation, a PLL can accommodate an infinite phase adjustment range. The DLL as described above has only a limited delay adjustment range. This not only prevents its use in plesiochronous clocking but also limits the frequency range over which the clock recovery can operate. [12] introduced a dual-loop DLL architecture, shown in Figure 2.10, to overcome this problem. A core multi-phase DLL similar to the one described above generates evenly spaced clock phases (coarse phases). Two phase multiplexers select two adjacent phases and a phase interpolator takes these two phases and generates finer phases in between. By selecting

different adjacent phases in sequence, the phase can be varied across a full clock cycle. Because the phase range is infinite, it is now possible to support plesiochronous clocking with this architecture. Figure 2.10 also shows a phase interpolator implementation originally described in [12]. By varying the fraction of the total current assigned to the ϕ_1 and ϕ_2 branches with *ctrl*, finer phases can be generated. In Chapter 6, we will analyze the phase interpolator in more detail.

2.4 Equalization Filter

For extremely lossy channels, severe inter-symbol interference (ISI) makes signal detection unreliable. Figure 2.11 shows the input and output eye diagram of a 1-m 7-mil 0.5-oz. PCB trace. The channel ISI causes the eye to close both vertically and horizontally. In the frequency domain, the channel acts as a low pass filter which attenuates the high

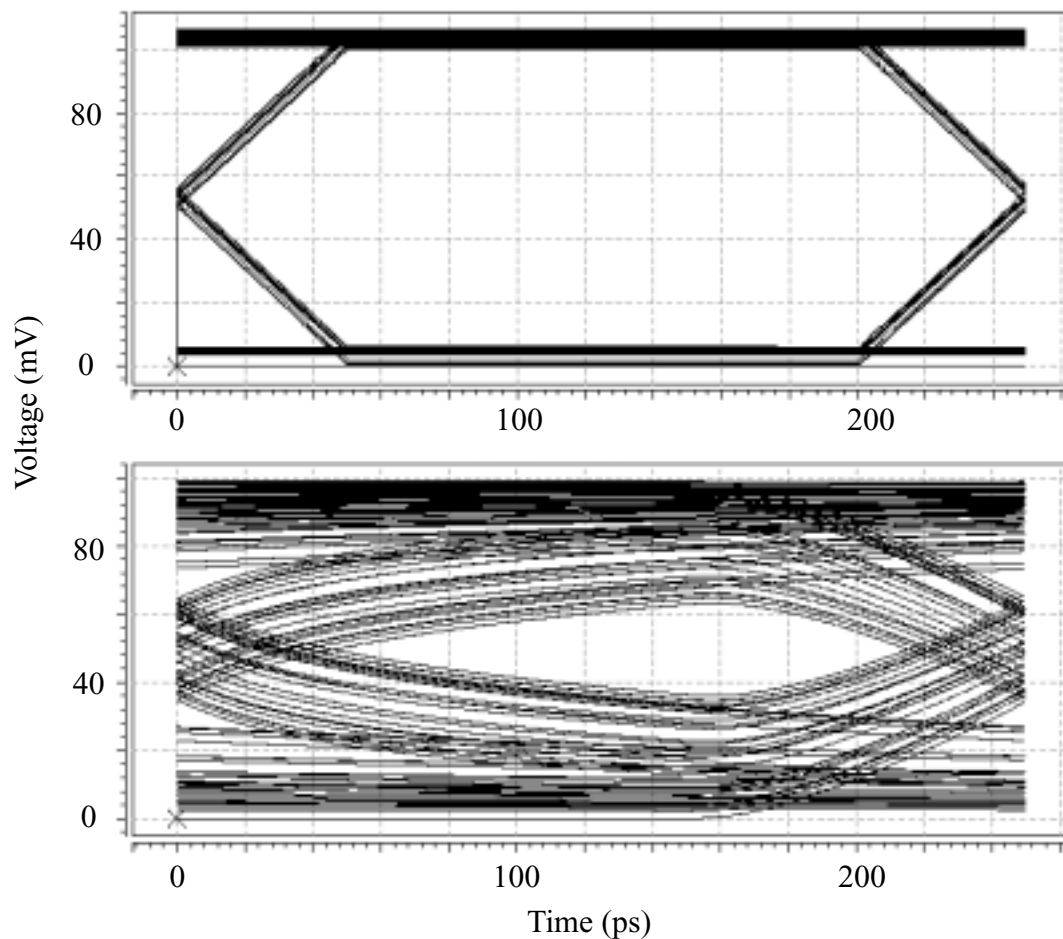


Figure 2.11: Input and output eye diagram before and after a 1-m, 7-mil, 0.5-oz. PCB trace.




Figure 2.12: Output eye diagram after a 1-m, 7-mil, 0.5-oz. PCB trace with a two tap equalization filter.

frequency component. To overcome this problem, an equalization filter can be placed at either the transmitter, the receiver, or both to undo the low-pass filtering [6] [9] [16] [19] [34]. Transmitter equalization, also called signal pre-emphasis, is usually implemented for its simplicity. A simple and common implementation of the filter is a symbol-spaced finite-impulse-response (FIR) filter described by the following equation.

$$V_o(n) = V_i(n) + a_1 \cdot V_i(n-1) + a_2 \cdot V_i(n-2) + \dots \quad (2.2)$$

where a_1, a_2, \dots are the filter tap coefficients. Figure 2.12 shows the eye diagram with a two tap filter. In this case, a high pass filter with $a_1 = -0.24$ is implemented. Compared to the bottom graph of Figure 2.11, both the vertical and horizontal openings are increased. We will analyze FIR equalization filters in more detail in Chapter 3.

[6] first incorporated transmitter equalization into I/O circuits. Figure 2.13 shows the transmitter architecture of this design. A 5-tap FIR filter is implemented in the digital domain (multiply-and-accumulate). A digital-to-analog converter (DAC) converts the resulting 4-bit data into an analog value and drives it off chip. The 4-bit data are encoded as 3 bits of positive drive and 3 bits of negative drive. These six bits directly select which

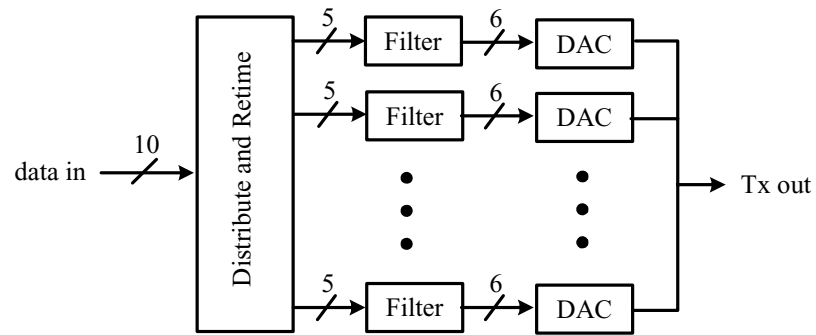


Figure 2.13: Dally & Poulton's transmitter architecture.




Figure 2.14: Analog current summing transmitter FIR filter.

of six pulse generators in the DAC connected that filter are enabled. To achieve a bit rate that is 10 times the on-chip clock frequency, a 10:1 multiplexer is implemented directly at the output using 10 clock phases (we will describe this output-multiplexing scheme in more detail in Chapter 6).

It turns out that the circuitry can be greatly simplified if one combines the multiply-and-accumulate function with the digital-to-analog conversion. This is the approach taken in [9], [19] and this work. Figure 2.14 shows a typical implementation.

The tap coefficients are adjusted via the tail current sources of the output drivers and the addition of the three taps is done directly at the output using analog current summing.

2.5 Summary

In this chapter we briefly described the previous work that is relevant to this thesis. Parallel transmitter and receiver architectures which push the bit rate beyond the on-chip clock frequency limit were described. Current-mode transmitter drivers improve the signal integrity by shielding the transmitted signal from the noisy supply and make it easier to vary the output swing.

On the receiver side, we described a gate-isolated receive amplifier, a pass-gate sampler, and a current integrating receiver. The first two approaches are based on *sampling* in which the input signal is sampled within a very small timing window. The last approach integrates the input signal over a bit time and is more robust against high-frequency noise.

Multi-phase generation and clock recovery based on a DLL and a PLL were presented. The pros and cons of each approach were discussed, leading to a dual-loop DLL architecture which provides an infinite phase range and is compatible with plesiochronous clocking.

Finally, we introduced the concept of channel equalization and presented previous equalization filter designs based on transmitter pre-emphasis. The rest of this thesis is devoted to each topic in more detail and describes how this work improves upon the previous designs to achieve our goals of power- and area-efficiency and noise immunity.

Chapter 3

System Overview

This chapter begins the description of the architecture and the circuit design of a 4-Gb/s inter-chip point-to-point serial link intended for large scale monolithic integration. As mentioned in Chapter 1, although previous designs have achieved bit rate in the multiples of Gb/s range [6] [7] [9] [13] [16], they require large amounts of power and area, making them unsuitable for large scale integration. The large amounts of power and area are a result of both architecture and circuit implementation. This chapter begins by giving a high-level overview of the serial link design. Section 3.1 describes the overall architecture. A target channel environment is introduced and analyzed in Section 3.2, with particular emphasis on determining the number of equalization taps required. Section 3.3 and Section 3.4 describe the timing convention and signaling convention of the transceiver and compare them with the alternatives. A summary is given at the end of the chapter.

3.1 System Architecture

Figure 3.15 shows the overall system architecture of the 4-Gb/s transceiver. On the transmit side, an on-chip multi-phase DLL generates four evenly spaced 1-GHz clock phases to sequence 1-GHz 4-bit-wide data on-chip into a 4-Gb/s bit stream off-chip. This is achieved by first re-synchronizing the 1-GHz data from a single clock domain to per-phase clock domains. Then a fast multiplexer, driven by the four phases, serializes the data. The serialization of parallel data on multiple phases of the clock ensures that the signaling speed is not constrained by the on-chip clock frequency, which is usually limited to about $8\tau_4$ (1-ns for 0.25 μ m technology) for reliable operation. Before going into the channel, the data stream is filtered by a finite-impulse-response (FIR) pre-emphasis filter (represented by multiple drivers implementing a digital-to-analog converter) to cancel out the frequency-dependent attenuation of the channel. Both the transmitter and receiver terminate the line into a digitally-trimmed matched impedance to make system less sensitive to reflections and cross-talks.

The receiver is a mirror image of the transmitter. The serial bit stream is sampled and de-serialized by a 4-phase 1-GHz receive clock generated by a receiver multi-phase DLL. The data are then re-synchronized from the per-phase clock domains to a single

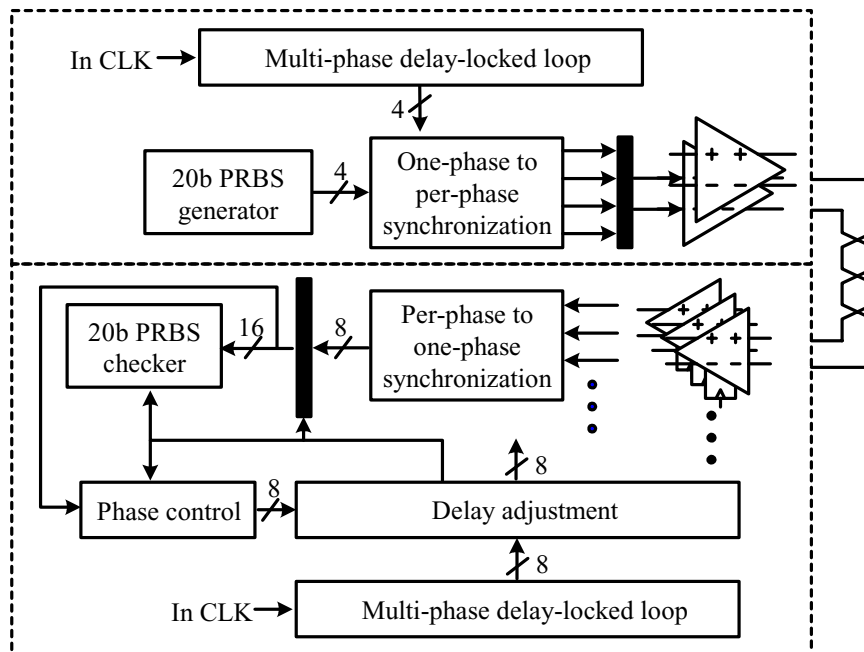


Figure 3.15: System Architecture of the 4-Gb/s transceiver.

clock domain to be processed further by other digital logic. The position of the receive clock is adjusted by a tracking clock recovery unit (delay adjustment and phase control blocks) to the center of the incoming data eye for maximum voltage and timing margin. To facilitate testing, a 20-bit pseudo-random bit sequence (PRBS) generator is integrated with the transmitter and a PRBS checker with the receiver.

3.2 Channel Analysis and Equalization

A realistic inter-chip serial link design must take into account the channel loss and distortion to have reliable communications. Figure 3.16 shows an example application of high speed serial links [20]. It is typical of high bandwidth communication switch systems such as Internet routers and SONET cross-connects. Input data come in, typically through a fiber channel, to a set of line cards. The data streams of the line cards then get routed to a switch card electrically via a backplane. In order to keep up with the exponential growth of bandwidth demands, these intra-system electrical interconnections must have as much bandwidth as possible to provide the maximum switching capability. High speed serial link techniques are typically employed in such interfaces. The printed-circuit-board (PCB)




Figure 3.16: A typical application of high speed serial links.




Figure 3.17: Frequency response of a 1-m PCB trace.

line trace, going from one end on the line card to the other end on the switch card, are typically 30-40 inches. This setup is typical of many digital systems. In this work, we will use this environment to study the equalization requirement. Many other types of transmission media have similar loss magnitude. This design uses transmitter pre-emphasis equalization to cancel the channel attenuation mainly because of its simplicity. Active receiver equalization requires analog signal values at the receiver to work (either through analog sampling and analog signal processing, or analog-to-digital conversion and digital signal processing). This complicates the system where simplicity is important for area, power and robustness considerations. Transmitter pre-emphasis, on the other hand, simply requires repetition of the output driver.

Previous designs have employed transmitter pre-emphasis. Various designs, however, implement different numbers of filter taps for similar channel requirements and symbol rates. In [6], a 5-tap FIR filter was implemented; in [9], a 4-tap FIR filter was implemented; in [19], a 2-tap FIR filter was implemented. The decisions were largely based on simulations and qualitative observations. This work takes a different approach and gives a mathematical foundation to the required number of filter taps and quantizes the trade-offs. In order to do this, a model of the above backplane environment needs to be

constructed. The major components of the backplane channel include the PCB trace and the connectors which connect the line cards/switch cards to the backplane.

Figure 3.17 shows the frequency response of a typical PCB stripline [21]. The cross section of the line is 8×0.7 mils and the actual measured length is 0.988m. The DC resistance is 5.48Ω . GETEK was used as the dielectric material. A set of RLGC parameters were calculated according this frequency response and the material properties to obtain a W-element PCB line model in Hspice. Hspice models the skin effect resistance of the wire and the dielectric conductance of the insulating material using the following equations [22].

$$R = R_o + R_s\sqrt{f} \quad (3.3)$$

$$G = G_o + G_d f \quad (3.4)$$

where R_o is the DC resistance in Ω/m , R_s is the skin effect parameter in $\Omega/m/\sqrt{Hz}$, G_o is the DC conductance of the dielectric material in S/m , and G_d is the dielectric loss parameter in $S/m/Hz$. Table 1 shows the RLGC parameters used for our nominal 50- Ω PCB trace.

Table 1: RLGC parameters for the PCB trace in HSPICE

L	327.5 nH/m
C	130.7 pF/m
R_o	5.48 Ω/m
R_s	$1.313 \times 10^{-3} \frac{\Omega}{m\sqrt{Hz}}$
G_o	0
G_d	$8.28 \times 10^{-12} \frac{\Omega}{m \cdot Hz}$

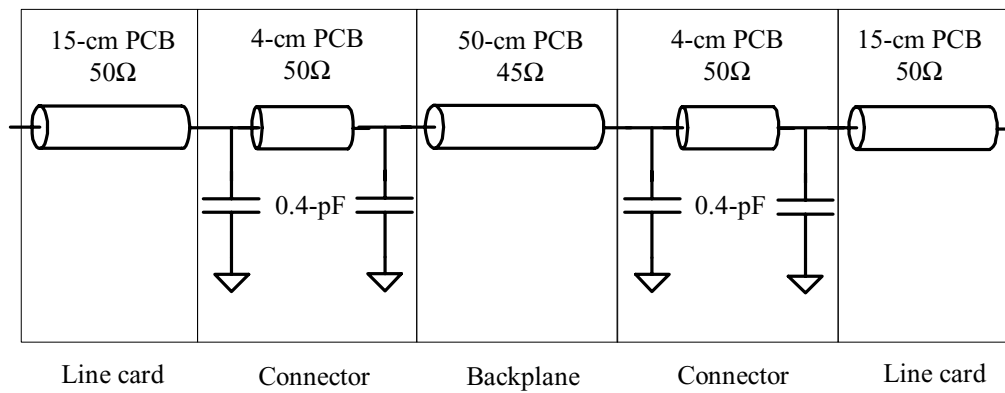


Figure 3.18: Circuit model of a typical backplane channel utilizing HSPICE's W-element.




Figure 3.19: Simulated S12 response of Figure 3.18.

Figure 3.18 shows the complete model of the backplane channel in HSPICE. The PCB lines on the line card and the backplane are 15-cm and 50-cm respectively. The line card to switch card connector is based on a Teradyne part intended for high speed

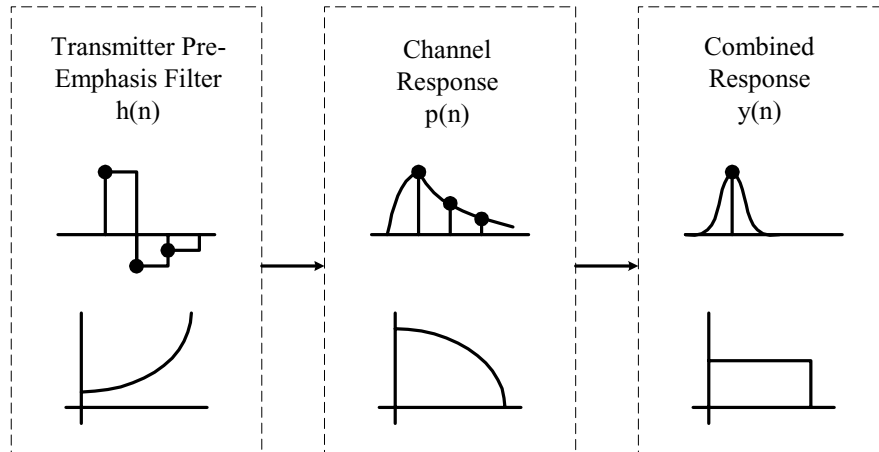


Figure 3.20: A model of the pre-emphasis filter and the channel.

signaling and is represented by two 0.4-pF capacitors for the vias and a short 4-cm PCB line for the internal trace of the connector. The line impedance of the backplane trace is adjusted to be 10% off from that of the line card trace to simulate impedance mismatch. Figure 3.19 shows the S12 response of the channel. The plot shows a larger resonance repeating approximately every 2-GHz and a smaller resonance repeating approximately every 0.13-GHz. When the forward traveling signal encounters the parasitic capacitance, a signal of opposite polarity gets reflected back. If the round trip delay of the signal is close to the signal period, significant attenuation results. The larger resonance exhibited by the S12 response is due to the 4-cm connector trace sandwiched between two 0.4-pF capacitors, and the smaller resonance is due the 50-cm backplane trace sandwiched between two 0.4-pF capacitors and impedance mismatch. The effect of these discontinuities is more channel loss beyond what is already present in the PCB lines and significant reflections at the near-end.

Figure 3.20 shows a representation of the communication link containing the pre-emphasis filter and the channel. The purpose of the pre-emphasis filter is to undo the channel distortion. In the frequency domain, this can be interpreted as having a high-pass pre-emphasis filter to cancel the low-pass filtering effect of the channel, resulting in a flat frequency response. In the time domain, this can be interpreted as transmitting additional

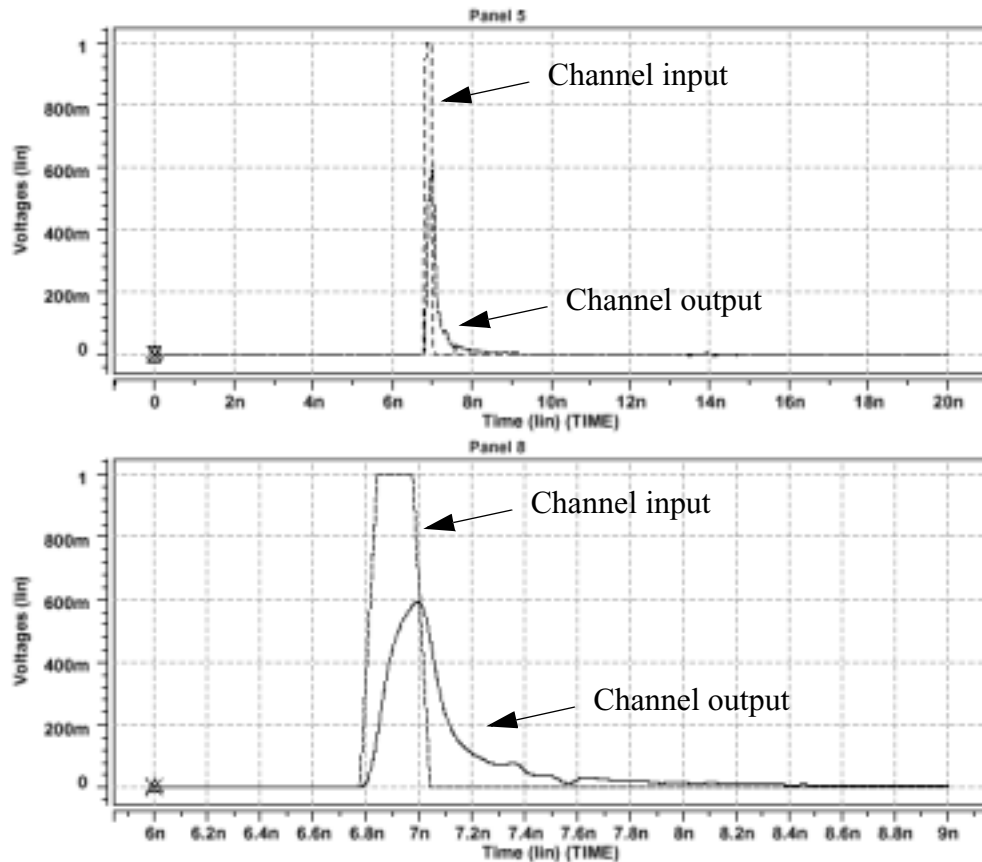


Figure 3.21: Simulated pulse response of Figure 3.18. The bottom plot is a zoomed-in version of the top plot.

bits to cancel the inter-symbol interference, resulting in a delta function. In this work, a method for determining the number of filter taps is presented. The method presented in this work is based on discrete-time signal processing calculations; therefore, a discrete-time model of the channel is required. We use a bit rate of 5-Gb/s for the investigation. Figure 3.21 shows the response of the channel to a 200-ps 1-V pulse with 66-ps rise-time¹. The discrete-time model of the channel $p(n)$ is then obtained by sampling this pulse response every 200-ps.

For a given number of transmitter filter taps, the combined response of the filter and the channel can be written as:

1. The rise time is usually controlled to be around 1/3 to 1/2 of the bit time.

$$\begin{bmatrix} y(0) \\ y(1) \\ \dots \\ y(l+k-2) \end{bmatrix} = \begin{bmatrix} p(0) & 0 & 0 & \dots & 0 & 0 \\ p(1) & p(0) & 0 & \dots & 0 & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & \dots & p(k-1) & p(k-2) \\ 0 & 0 & 0 & \dots & 0 & p(k-1) \end{bmatrix} \begin{bmatrix} h(0) \\ h(1) \\ \dots \\ h(l-1) \end{bmatrix} \quad (3.5)$$

where l is the number of filter taps and k is the channel pulse response length. Ideally, the following relations should be satisfied for the channel output to be free of inter-symbol interference and signal attenuation:

$$\begin{cases} y_{des}(n) = d, n = \Delta \\ y_{des}(n) = 0, n \neq \Delta \end{cases} \quad (3.6)$$

where $y_{des}(n)$ is the desired channel output, d is the received signal amplitude desired, and Δ represents equalizer-channel delay. A non-zero Δ implements a non-causal filter which cancels precursor ISI as well. For the bit rate of interest, no significant precursor ISI is present. Therefore, Δ is 0. In most cases Equation (3.5) has no exact solution for $h(n)$ since it is an over-determined set of linear equations. However, we can find a closed-form solution which minimizes the square of the error. The square of the error can be expressed by

$$\begin{aligned} \mathbf{E} &= \mathbf{Y} - \mathbf{Y}_{des} = \mathbf{P}\mathbf{H} - \mathbf{Y}_{des} \\ \|\mathbf{E}\|^2 &= \mathbf{H}^T \mathbf{P}^T \mathbf{P} \mathbf{H} - 2\mathbf{Y}_{des}^T \mathbf{P} \mathbf{H} + \mathbf{Y}_{des}^T \mathbf{Y}_{des} \end{aligned} \quad (3.7)$$

Next we take the derivative of Equation (3.7) with respect to h and set it to 0.

$$\begin{aligned} \frac{d}{d\mathbf{H}} \|\mathbf{E}\|^2 &= 2\mathbf{H}^T \mathbf{P}^T \mathbf{P} - 2\mathbf{Y}_{des}^T \mathbf{P} = 0 \\ \mathbf{H}^T \mathbf{P}^T \mathbf{P} &= \mathbf{Y}_{des}^T \mathbf{P} \end{aligned} \quad (3.8)$$

If $\mathbf{P}^T \mathbf{P}$ is invertible (i.e. \mathbf{P} is full-rank), then the optimal tap coefficients in the least square sense, \mathbf{H}_{ls} , is given by

$$\mathbf{H}_{ls} = (\mathbf{P}^T \mathbf{P})^{-1} \mathbf{P}^T \mathbf{Y}_{des} \quad (3.9)$$

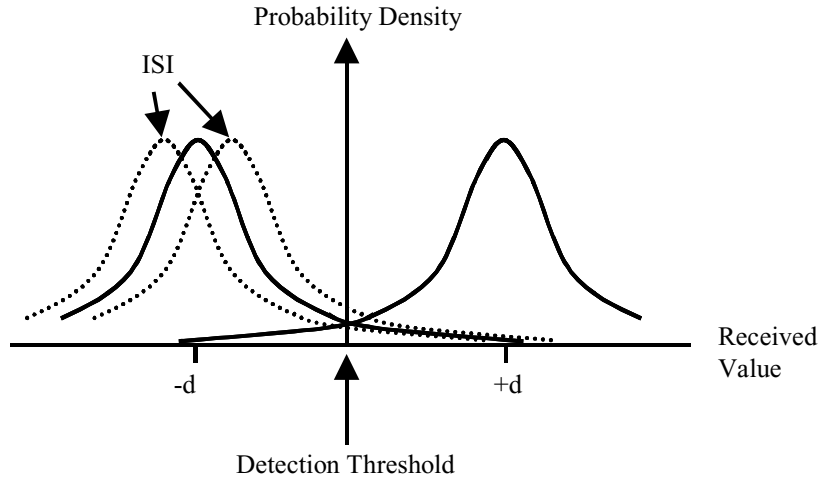


Figure 3.22: Effect of the ISI on the bit error rate.

The residual ISI is

$$\mathbf{E} = \mathbf{P}\mathbf{H}_{ls} - \mathbf{Y}_{des} \quad (3.10)$$

Figure 3.22 shows the effect of the ISI on the bit error rate. Assuming the random noise has a Gaussian distribution with zero mean, the received signal would have a Gaussian distribution centered around $\pm d$ without ISI. With this Gaussian distribution, there is a finite probability of detection error, given by the area under the probability density exceeding the detection threshold. ISI shifts the Gaussian distribution, resulting in either increased (shift to the right in Figure 3.22) or decreased (shift to the left in Figure 3.22) probability of error depending on the exact ISI pattern. Ideally, the received signal without any ISI or noise should be

$$m = \begin{cases} +d, & \text{when 1 is sent} \\ -d, & \text{when 0 is sent} \end{cases} \quad (3.11)$$

With ISI, the received signal for a given ISI pattern, K , is given by

$$s(K) = m + \mathbf{K}^T \mathbf{E} \quad (3.12)$$

where \mathbf{K} has the same length as \mathbf{E} . In our present example, the channel response has 5 taps (see Figure 3.21). For a two-tap filter, the length of \mathbf{K} and \mathbf{E} is 6. \mathbf{K} represents the bit




Figure 3.23: Effect of the ISI on the bit error rate. The different curves correspond to different levels of Gaussian noise.

pattern of the ISI. In this example, if the bit pattern $\{0, 1, 1, 1, 0, 1\}$ is sent, then the ISI pattern, \mathbf{K} , at the last bit is $[-1 \ +1 \ +1 \ +1 \ -1 \ +1]^T$. For binary signal encoding, the probability of error for a given \mathbf{K} is

$$P_e(\mathbf{K}) = \frac{1}{\sqrt{2\pi}\sigma} \int_{s(\mathbf{K}) \times \text{sign}(m)}^{\infty} \exp\left(-\frac{u^2}{2\sigma^2}\right) du \quad (3.13)$$

where σ^2 is the Gaussian noise variance. Finally, the overall probability of error is

$$P_e = \sum_{\mathbf{K}} \text{Pr}\{\mathbf{K}\} P_e(\mathbf{K}) = \frac{1}{2^{\text{length}(\mathbf{K})}} \sum_{\mathbf{K}} P_e(\mathbf{K}) \quad (3.14)$$

where the second equality is valid for random data.

Figure 3.23 shows the bit error rate (BER) versus the number of pre-emphasis filter taps for sending 5-Gb/s binary bit stream down the backplane channel. The three

curves correspond to different levels of random Gaussian noise. The noise variance are chosen such that, if the received signal is a delta function (i.e. no ISI) with amplitude d , the bit error rate would be 10^{-10} , 10^{-20} , and 10^{-30} , respectively. That is, the BER numbers specified are the minimum achievable by each curve. We can then determine how many taps are required to get the BER close to the minimum level. As indicated by the plot, a simple two-tap filter brings the BER very close to the minimum level. For reference, Table 2 shows the tap weights calculated by the least square method. The tap weight after the

Table 2: Equalization tap weight calculated by the least square method

	Tap 1	Tap 2	Tap 3	Tap 4
2-tap equalizer	1	-0.309	0	0
3-tap equalizer	1	-0.298	-0.034	0
4-tap equalizer	1	-0.298	-0.032	-0.026

second tap is vanishingly small, indicating the need for higher precision if additional benefit beyond two taps is to be realized. Since this analysis indicates a good BER performance over a typical backplane with a simple two-tap filter, the serial link design in this work only includes a two-tap filter to minimize power and area.

In order to see what would happen when the channel becomes more lossy, the backplane PCB trace in Figure 3.18 is increased from 50-cm to 200-cm. Figure 3.24 shows BER versus the number of filter taps for this modified channel. The plot shows that there is a significant gain up to 5 or 6 taps and that the BER does not quite reduce to its minimum even with 10 taps.

3.3 Timing Convention

Figure 3.25 shows three major components of a timing budget on an abstract eye diagram [21]. The rise time, t_r , is the time required for the waveform to switch states. The aperture time, t_a , is the duration over which the signal must be above the voltage margin




Figure 3.24: Effect of the ISI on the bit error rate for a long backplane channel.




Figure 3.25: An abstract eye diagram showing the timing budget.

for the receiver to operate reliably. The timing uncertainty, t_u , is the peak-to-peak timing error from the nominal waveform. The minimum bit time must be greater than the sum of the three components, as expressed by:

$$t_b \geq t_r + t_a + t_u \quad (3.15)$$

t_m is the timing margin of the signal. Depending on the voltage margin, we can draw different *margin rectangles* inside the eye, as shown in Figure 3.25. A smaller allowable voltage swing would result in a larger timing margin, t_m .

For a good transmitter design with 50- Ω on-chip termination, the rise time, t_r , can be as low as $1\tau_4$ delay. The receiver aperture time, t_a , is a function of the receiver topology. For a gate-isolated sense-amplifier, t_a is on the order of $0.2 - 0.3\tau_4$. $t_r + t_a$ relates to the raw speed of the transistor and is about 150-ps in 0.25 μm CMOS technology. The timing uncertainty, t_u , on the other hand, is determined by noise, which results in timing jitter, and mismatch, which results in static offset, in a system. The matching properties of the system has much to do with the timing convention, which is described next.

Modern high-speed inter-chip communication systems usually employ either source-synchronous timing (also know as bundled close-loop timing) or per-line closed-loop timing, which is employed in this design. Figure 3.26 shows a typical source-synchronous timing system. The source sends a clock signal along with the data, and the receiving end uses a delay-locked loop to align its sampling clock with the transmitted clock (plus setup time). Figure 3.27 shows a typical per-line closed-loop timing system. The timing information is extracted directly from the data stream by detecting the presence of data transitions. The advantage of source-synchronous timing is that it is simpler to implement and does not require any special encoding to ensure enough transitions are present in the data signal. It also allows one timing circuit to be shared across a group of data signals. Whereas source-synchronous timing can be used in a multi-drop bus environment through *round-trip distribution* [35], per-line closed-loop timing must be point-to-point. However, source-synchronous timing has many more uncancelled skews that make its bit rate much lower. Delay measurements of commercial parts have shown skews of 50-60 ps per meter of printed-circuit board trace, per connector, or per package pin [17]. For a transceiver communicating over a backplane, skews of >250-ps can be expected between clock and data lines. Clearly, per-line closed-loop timing has to be used to operate at 4-Gb/s.

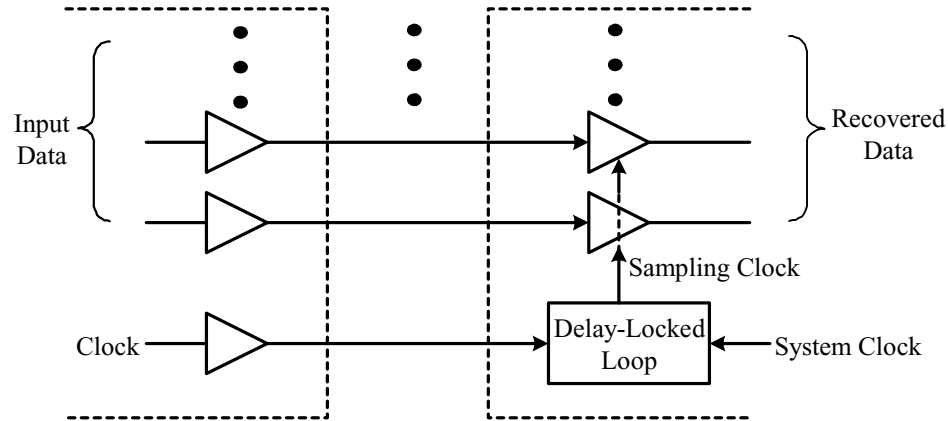


Figure 3.26: A bundled closed-loop timing system.




Figure 3.27: A per-line closed-loop timing system.

Table 3 lists the expected worst case timing budget of our 4-Gb/s system. Assuming that the transmitter drive current is 10-mA and the rise time is approximately a bit time (250-ps) for a very lossy channel, then the portion of the rise time which eats into the timing budget for an offset-cancelled receiver differential sensitivity of 20-mV is

$$t_r = 250ps \times \frac{20mV}{500mV} = 10ps \quad (3.16)$$

The receiver aperture time, t_a , for a gate-isolated sense amplifier is on the order of 0.2-0.3 τ_4 (~30-ps in our technology). The pk-pk transmitter clock jitter is on the order of 20-ps. Since we are using a dual-loop clock recovery architecture in which the receiver clock might dither between 1-2 steps, the expected receiver clock jitter is around 50-ps. From actual lab measurement of the silicon, the transmitter and receiver clock phases have

about 15-ps and 30-ps of offsets. The total of the above timing budget is 155-ps. This leaves about 0.38-UI (95-ps) of margin for the channel ISI.

Table 3: Worst case timing budget for our I/O system.

Rise time (t_r)	10-ps
Receiver aperture (t_a)	30-ps
Tx clock jitter	20-ps
Rx clock jitter (include dithering)	50-ps
Tx phase offset	15-ps
Rx phase offset	30-ps
Total	155-ps

3.4 Signaling Convention

3.4.1 Differential vs. Single-Ended Signaling

Differential signaling requires two wires and pins per channel, whereas single-ended signaling requires only one wire and pin per channel. Due to self-induced power supply noise, however, differential signaling usually requires less than twice as many pins compared to single-ended signaling, as explained below. Although less efficient in terms of pin utilization, differential signaling has many advantages which make it more robust and better suited for a large digital system. These are described in more details below.

Self-induced power supply noise. A differential driver, unlike a single-ended driver, always draws a constant amount of current from the power supplies, resulting in very little AC power supply current. The stable power supply current draw helps reduce power supply noise due to wire inductance (i.e. $L di/dt$ noise). The following example demonstrates how the pin-efficiency of single-ended signaling is not twice as much as that of differential signaling as a result of additional power supply pins required. Let us assume realistically that we want to design a high bandwidth switch fabric chip which requires 100 4-Gb/s high speed serial links for a total bandwidth of 400-Gb/s both into and out of the chip. If we were to use single-ended current-mode drivers, each putting 4-mA of current on the line (100-mV with 25- Ω double termination) with 100-ps rise-time (40% of

the bit time), then the total noise generated on the power supply inductance, L_{sup} , when all 100 drivers are switching simultaneously, is:

$$V_{noise} = L_{sup} \times \frac{di}{dt} = L_{sup} \times \frac{4mA \times 100}{100ps} = 4 \times 10^9 L_{sup} \quad (3.17)$$

If V_{noise} were to be kept below 100-mV, L_{sup} would need to be < 0.025 -nH. A differential driver always sinks a constant amount of current, greatly reducing the di/dt noise. As technology scales and supply voltage decreases, this advantage will only become more important. In the above example, for a typical wire bond inductance of 2-3 nH, one single-ended driver would require one supply pin for the switching current, whereas a good differential driver requires none. In other words, the pin count requirement is equal. In reality, the pin count advantage of differential signaling is not as big due to transient glitches. However, the fact that the pin inefficiency of differential signaling becomes less significant as bit rate increases and supply voltage decreases remains the same.

Return current. With differential signaling, the return current is a constant DC value. In an environment where the return current paths are shared among a group of channels (as is the case for PCB), cross-talk among adjacent channels is significantly reduced compared to single-ended signaling, where the switching currents in the shared current path couple into other channels.

References. A differential signal serves as its own receiver reference. Unlike the transmitter generated reference which is shared among a group of single-ended lines, the differential lines are usually tightly coupled (or even twisted) and easily make many noise sources common mode to the receiver.

Signal swing. The voltage difference between a 1 and a 0 for differential signaling (henceforth called the differential swing) is twice that of the value for single-ended signaling (henceforth called the single-ended swing). For many drivers whose single-ended swings are limited¹, differential signaling can provide more noise margin.

1. For example, a current mode driver needs to keep its output transistor(s) in saturation. As technology scales and power voltage decreases, the swing of the output driver becomes more limited.

In summary, differential signaling creates less noise and has better noise immunity compared to single-ended signaling. Its disadvantage, namely the pin inefficiency, will become less significant as bit rate increases and supply voltage decreases.

3.4.2 Binary vs. Multi-Level Encoding

One method to increase the achievable bit rate is to encode multiple bits in a data symbol using multi-level signaling. Instead of two voltage levels, a digital-to-analog converter (DAC) can be used to encode multiple bits on multiple voltage levels. For example, one can encode 2 bits/symbol on 4 voltage levels, reducing the required bandwidth by half while achieving the same bit rate [9]. The decrease in signal bandwidth can potentially reduce the amount of ISI; however, since ISI is a proportional noise, it can potentially increase as well due to a higher voltage swing requirement for multi-level encoding. We can write a corresponding set of Equation (3.11) – Equation (3.14) for 4-level signaling as follows.

$$m = \begin{cases} +3d, & \text{when 10 is sent} \\ +d, & \text{when 11 is sent} \\ -d, & \text{when 01 is sent} \\ -3d, & \text{when 00 is sent} \end{cases} \quad (3.18)$$

$$s(\mathbf{K}) = m + \mathbf{K}^T \mathbf{E} \quad (3.19)$$

$$P_e(\mathbf{K}) = \begin{cases} \frac{\sqrt{2}}{\sqrt{\pi}\sigma} \int_{s(\mathbf{K}) \times \text{sign}(m) - (|m| - d)}^{\infty} \exp\left(-\frac{u^2}{2\sigma^2}\right) du, & \text{if } m = \pm d \text{ (inner levels)} \\ \frac{1}{\sqrt{2\pi}\sigma} \int_{s(\mathbf{K}) \times \text{sign}(m) - (|m| - d)}^{\infty} \exp\left(-\frac{u^2}{2\sigma^2}\right) du, & \text{if } m = \pm 3d \text{ (outer levels)} \end{cases} \quad (3.20)$$

$$P_e = \sum_K Pr\{K\} P_e(K) = \frac{1}{4^{\text{length}(K)}} \sum_K P_e(K) \quad (3.21)$$

Gray code is used for the multi-level encoding since crossing to the most immediate level(s) only results in one bit error. The format of K now includes ± 3 . For example, if the length of e is 4 and the bit pattern {00, 11, 01, 10} is sent, then $K = [-3 +1 -1, +3]^T$. Figure




Figure 3.28: BER versus the number of filter taps for 4-PAM and 2-PAM signal encoding. The symbol rate of 4-PAM is half that of 2-PAM.

3.28 shows the result of these modifications for 4-level signaling. For comparison, the corresponding curves for binary signaling are also included. The curve for 4-level signaling is at 2.5-GSymbols/s (5-Gb/s) and that for binary signaling is at 5-GSymbols/s (5-Gb/s). As indicated by the plot, binary signaling performs slightly better at 2 taps of equalization. Figure 3.28 also indicates that the BER for 4-level signaling decreases faster than that for binary signaling. This is due to the longer ISI for binary signaling since its symbol rate is higher. The plot shows that neither has a significant advantage as far as channel equalization is concerned.

Multi-level signaling only makes sense when the channel bandwidth is severely constrained or when the circuit speed is limited. The energy per bit required for multi-level encoding compares unfavorably with that required for binary encoding. Assuming the fixed noise source (such as receiver offset and sensitivity) has amplitude V_{NF} and the

uncancelled proportional noise (such as crosstalk, detection threshold variation, and transmitter offset) is some fraction K_N of the signal swing, V_{SW} , then the required V_{SW} for a given signal-to-noise ratio, K_{NM} , can be found as

$$V_N = V_{NF} + K_N V_{SW} \quad (3.22)$$

$$V_{SW} > 2(N-1)K_{NM}V_N$$

$$V_{SW} > \frac{2(N-1)K_{NM}V_{NF}}{1-2(N-1)K_{NM}K_N}$$

The energy per bit required for a fixed power supply is thus

$$E_{bit} = \frac{2(N-1)K_{NM}V_{NF}t_{bit}V_{dd}}{Z \log_2 N (1-2(N-1)K_{NM}K_N)} \quad (3.23)$$

where t_{bit} is the bit time and Z is the transmission line impedance. E_{bit} increases considerably with N . In particular, there is an upper limit on K_{NM} , which can be expressed as

$$K_{NM} < \frac{1}{2(N-1)K_N} \quad (3.24)$$

For $K_N = 15\%$, the signal-to-noise ratio K_{NM} is limited to 3.33 and 1.11 for binary signaling and 4-level signaling respectively. In order to have any margin against noise (i.e. $K_{NM} > 1$), K_N cannot exceed 50% and 16.7% for binary signaling and 4-level signaling respectively. These numbers indicate considerable noise immunity degradation for multi-level signaling. The ease of implementation, favorable energetics, noise immunity,




Figure 3.29: Simultaneous bi-directional signaling.

comparable equalization performance, and adequate circuit speed (which is discussed in Chapter 3 – Chapter 5) make binary signaling the clear choice in this design.

3.4.3 Uni-directional vs. Simultaneous Bi-directional

Another method to increase the effective pin bandwidth is to send bits in both directions simultaneously over the same channel through simultaneous bi-directional signaling [23]. The effective wire density and pin count of the system can be doubled. As shown in Figure 3.29 a replica transmitter with matched delay produces the same waveform as the main transmitter. The receiver subtracts this waveform from the signal on the transmission line to cancel out the component which is due to its own transmitted data.

Since simultaneous bi-directional signaling can operate at half the bit rate and still achieve the same effective bandwidth per pin as uni-directional signaling, its ISI is

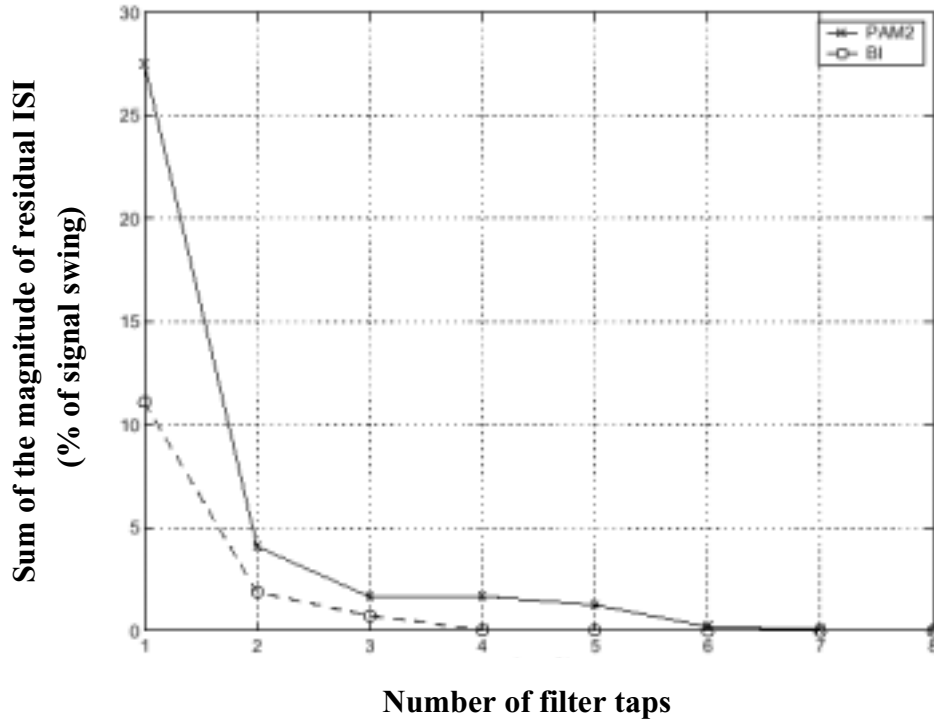


Figure 3.30: Sum of the magnitude of ISI versus the number of filter taps for unidirectional (PAM2) and bi-directional (BI).

smaller. Figure 3.30, shows the sum of residual ISI magnitude versus the number of filter taps for both uni-directional and bi-directional signaling. The tap weight and residual ISI are again calculated from the least square method presented above. Since ISI is one form of proportional noise, it is expressed as a percentage of the signal swing. The maximum sum of residual ISI is 11% and 27% for bi-directional and uni-directional respectively without equalization and 2% and 4% with a two-tap pre-emphasis filter (50% would render the signal undetectable).

Although simultaneous bi-directional signaling reduces ISI due to frequency-dependent channel attenuation, it is much more susceptible to many other forms of proportional noise such as near-end crosstalk, channel reflections, and replica offset.

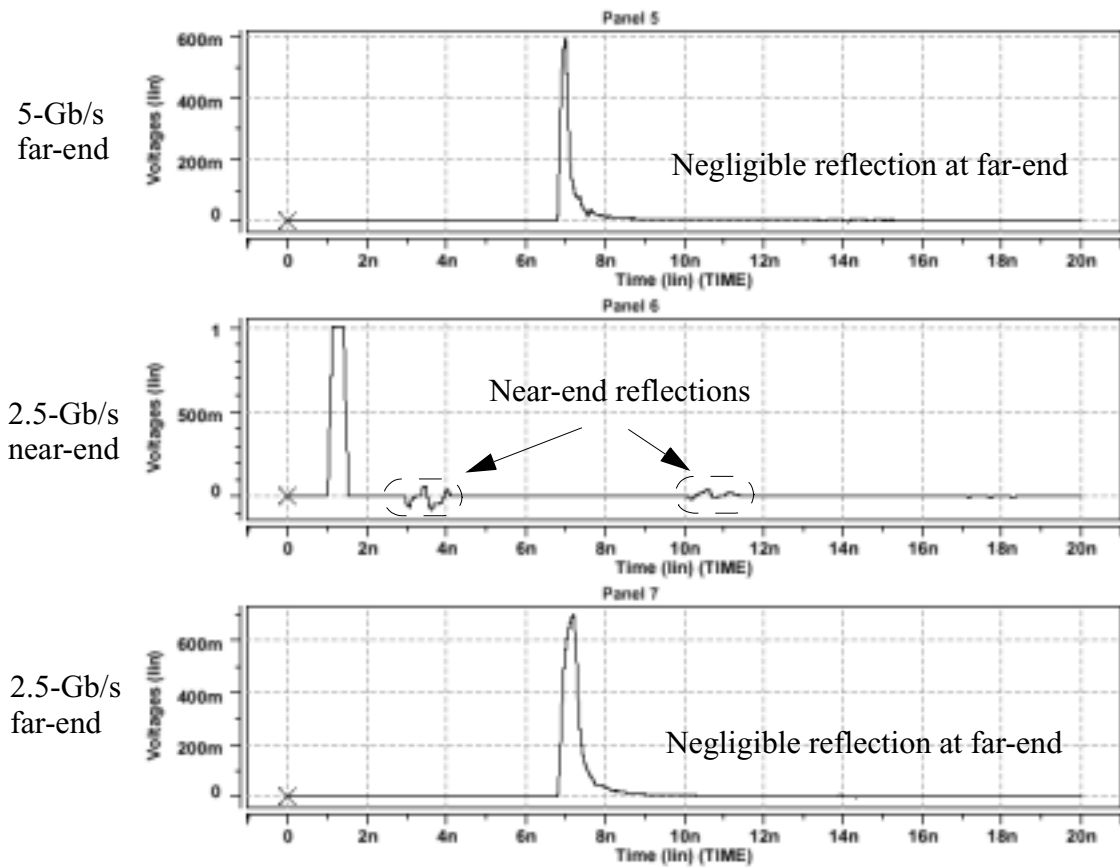


Figure 3.31: Near-end and far-end signal of the sample channel.

Figure 3.31 shows the near-end and far-end voltage for the backplane channel of Figure 3.18 at 5-Gb/s and 2.5-Gb/s. As shown by the plots, there is negligible far-end reflection at both 5-Gb/s and 2.5-Gb/s due to channel attenuation and double termination. However, we can see significant near-end reflections, which total to about 35% of the received signal. The near-end reflections only affect simultaneous bi-directional signaling, resulting in 35% additional proportional noise compared to uni-directional signaling. Although the near-end reflections can be reduced with a filter, the fact that it requires a long filter length and that its arrival time depends critically on the exact length of the channel makes it difficult and expensive to implement in reality.

Replica offset is another significant proportional noise only present in simultaneous bi-directional signaling. In the presence of channel attenuation, the effect of it is more pronounced. Figure 3.32 shows simultaneous bi-directional signaling waveform for no channel attenuation. The transmitter subtracts its own signal from the signal on the channel via a replica driver to obtain the received signal. Because the main transmitter sees additional package parasitics not present at the replica transmitter output and because the replica transmitter is usually a scaled-down version of the main transmitter, a mismatch of both voltage and delay is present between the main and replica transmitter. With a current-integrating receiver¹ [17], we can consider both the voltage and delay




Figure 3.32: Simultaneous bi-directional signaling waveform without channel loss.




Figure 3.33: Simultaneous bi-directional signaling waveform with channel loss.

1. A sampling sense amplifier is unreliable as the sampling instant might happen during the delay mismatch or during the transient of the replica transmitter. A current-integrating receiver attenuates these effects.

mismatch as a voltage proportional noise. The problem is compounded by the channel loss. Figure 3.33 shows the simultaneous bi-directional signaling waveform for a 0.5 channel gain (or a channel attenuation of 2). The transmitter output swing is now larger than the received input swing, making any proportional noise on the transmitter output more pronounced. Assuming the channel gain is A_C and the proportional noise due to replica mismatch on the transmitter output is V_{NP} then the effective proportional noise on the received input is

$$K_{RP} = \frac{V_{NP}}{A_C} \quad (3.25)$$

For example, with a 10% replica mismatch and a channel gain of 0.69¹, K_{RP} is 15%.

The above analysis indicates that, although simultaneous bi-directional signaling suffers from smaller channel attenuation due to its lower bit rate requirement, other proportional noise is much worse. In particular, for our backplane channel it has an additional 35% of proportional noise due to near-end reflections and 15% of proportional noise due to an estimated 10% replica transmitter mismatch that are not present in uni-directional signaling. These noise sources more than overwhelm the channel attenuation advantage. We can also see that for this backplane channel, the sum of near-end reflections and replica mismatch (totaling 50%) renders the signal completely undetectable without even considering other sources of noises.

3.5 Summary

This chapter presented the high-level architecture of the 4-Gb/s transceiver and explains its timing conventions and some of its signaling conventions. Per-line closed-loop timing is used to eliminate many of the mismatch-dependent static timing uncertainties and achieve a much higher bit rate compared to source synchronous systems. Since this transceiver is intended to be used in a large digital system, its equalization is designed to cancel the frequency-dependent attenuation of a backplane. A least-square

1. For the backplane channel of Figure 3.18, the main tap is 0.69 of the transmitted amplitude at 2.5-Gb/s, as shown by Figure 3.31.

analysis of the channel was presented and a two-tap FIR filter was deemed to be adequate in this environment. The signaling convention of this design was also presented. We use uni-directional and differential signaling with binary encoding mainly because of noise generation and noise immunity concerns. In the next three chapters, we discuss the transceiver circuits in more detail.

Chapter 4

Transmitter

This chapter presents the transmitter design. The purpose of the transmitter is to filter the data according to the channel and drive the resulting signal off chip with the least amount of power, area and noise. To alleviate the frequency requirement of the timing circuits and the digital logic, we use a 4:1 multiplexer to serialize low-speed parallel data on 4 evenly-spaced phases of the 1-GHz clock, giving a bit rate of 4-Gb/s. A low-swing input-multiplexed architecture is used to achieve a good compromise between speed, power, area, and transmitter output loading compared with previous designs.

Section 4.1 starts with architectural considerations and compares this approach with previous designs. The circuit implementation and analysis are presented in Section 4.2, followed by a summary at the end. The discussion of the transmitter timing circuit, namely the delay-locked loop for generating multiple phases, is deferred until Chapter 6.

4.1 Architecture

The shortest achievable clock period in a given technology is generally limited to be no less than about $8\tau_4$ (roughly 1-ns in 0.25 μm) for adequate margins [24] [25]. Although it is possible to use a faster clock, it puts significant burden on the timing circuit design, clock distribution, and data synchronization. In this design, we keep the clock

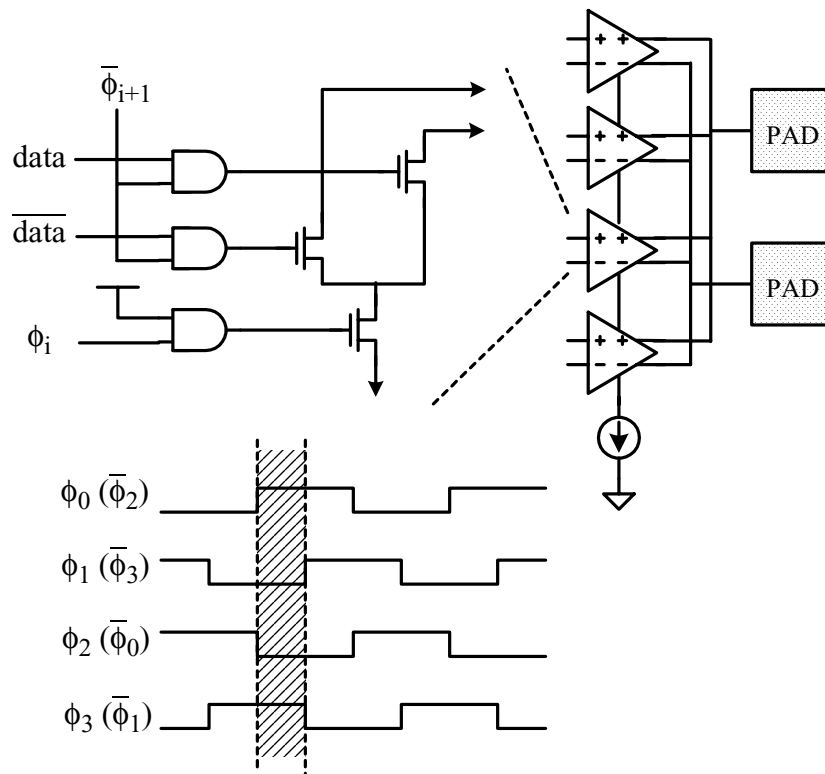


Figure 4.1: Output-multiplexed transmitter architecture.

period above the $8\tau_4$ limit and rely on the front end circuits to multiplex and de-multiplex data. On the transmitter side, this means a fast multiplexer is needed to take a parallel signal with this clock period and multiplex it, using multiple clock phases, into a serial signal with a shorter bit time, $2\tau_4$ in the present case. As shown in Figure 4.1 previously published transmitter designs achieve high bandwidth by multiplexing directly at the output pin where both a low time constant ($25\ \Omega$ double termination impedance and a few pF capacitive load) and small swings are present. Two adjacent clock phases are used to generate a short differential current pulse equal to a bit time. The minimum bit time previously reported has been on the order of τ_4 [6] [8].

Although fan-out delay numbers have been used extensively to report the performance of an output-multiplexed architecture, the minimum bit time achievable with this architecture will cease to scale with the process technology in the near future. In previous and current CMOS technologies, the bandwidth at the output pad is large compared to the bandwidth on-chip due to the output transmission line, which essentially

creates a shunting resistance without any capacitive loading. Furthermore, the ESD protection devices at the output have thus far remained a smaller portion of the total output capacitive load or simply ignored. In the near future, this picture will change dramatically since the doping of the P-N junction, on which most ESD protection devices rely to provide a low-resistance discharge path, will increase with the scaling of the transistor dimensions. While the capacitive load of the transmitter driver and the termination resistor decreases with shrinking transistor dimensions, that of the ESD protection devices increases with increasing silicon doping. Therefore, the bandwidth at the output pad will quickly cease to have any advantage compared to the bandwidth on chip.

To make matters worse, the output-multiplexed architecture requires multiple copies of the output driver, each sized large enough to drive signals off chip. Besides the area penalty, the total capacitive loading at the output significantly reduces the effectiveness of the transmitter termination resistor. For example, assuming a total capacitive loading of 1.5-pF and a 50- Ω termination at the transmitter output, then for 4-Gb/s data speed, the approximate impedance seen by the reflected signal is

$$R_{term} = 50 \parallel \frac{1}{2\pi(2 \times 10^9)(1.5 \times 10^{-12})} = 26\Omega \quad (4.26)$$

In other words, approximately 25% of the reflected signal would be reflected back again toward the receiver.

A more area-efficient method which also produces less capacitive loading at the transmitter output is to perform multiplexing at the input of the transmitter. A popular scheme is to use static CMOS gates to perform multiplexing and buffering to drive the final output driver as shown in Figure 4.2 [10] [17]. However, this multiplexer is unable to achieve bit rate higher than $4\tau_4$ (2-Gb/s in 0.25 μm) because of the bandwidth limit of CMOS gates. Figure 4.3 shows the maximum bit rate vs. the degree of multiplexing. The shaded area denotes the achievable bit time. The speed is initially limited by the achievable clock frequency at 2:1 multiplexing ratio. Above this point, high multiplexer fan-in becomes the bottleneck and the achievable speed gradually decreases.

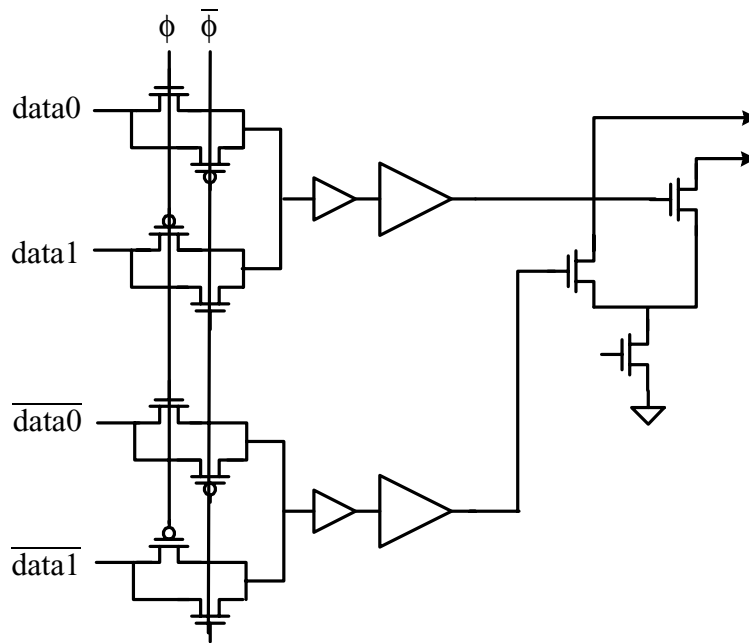


Figure 4.2: CMOS gate based input-multiplexed transmitter architecture.

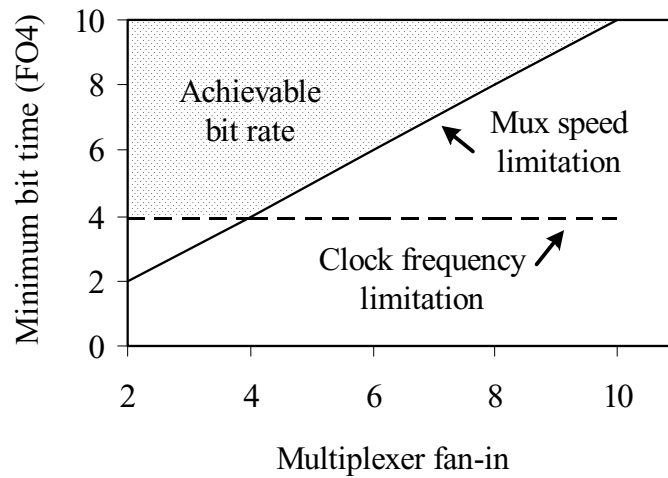


Figure 4.3: Minimum achievable bit time the configuration in Figure 4.2.

This speed limitation is not an inherent property of the process technology but of the circuit topology. Since the transmitter output swing is much smaller than the full CMOS swing, signal attenuation is another degree of freedom in optimizing speed and power. In the output-multiplexed architecture, all of the signal attenuation occurs at the bottleneck point, the output. This signal attenuation trades gain for higher bandwidth.

Figure 4.2, on the other hand, applies no signal attenuation at its bottleneck node, the multiplexer. A good balance between maximizing speed and minimizing power is achieved by an input-multiplexed architecture where the voltage swing and the capacitive fan-out from the multiplexer to the final output are carefully chosen to meet the required speed and final output swing. In order to do this, circuit topologies which allow direct trade-off between signal swing, bandwidth and power are employed. We will describe the circuit topologies in the next section. For a practical implementation, the minimum bit time achievable using this strategy is about $2\tau_4$.

For the current process technology, a general rule of thumb can be derived from the above discussion in selecting a transmitter architecture given the system performance requirement. If the target bit time is above $4\tau_4$ (500-ps in $0.25\ \mu\text{m}$), a CMOS gate based input-multiplexed architecture should be used due to its simplicity. For bit rate between $2\tau_4$ and $4\tau_4$, a swing-optimized input-multiplexed architecture can be used to achieve higher speed without the area penalty of an output-multiplexed architecture. For the fastest bit rate, an output-multiplexed architecture should be employed. As technology scales and the ESD limitation becomes more severe, however, a swing-optimized input-multiplexed architecture might achieve a higher bit rate than an output-multiplexed one.

4.2 Circuit Implementation

Figure 4.4 shows the transmitter circuit diagram. It consists of a 4:1 multiplexer, a pre-amplifier, and an output driver. The transmitter employs dual pseudo-NMOS multiplexers at its input, one for the signal and one for its complement. Each multiplexer is switched by two series NMOS that are gated by two adjacent clock phases. Thus, input d_i is enabled onto the pre-amplifier input during phase ϕ_i . ϕ_0 through ϕ_3 are generated by a multi-phase DLL described in Chapter 6. Figure 4.5 shows the pulse-amplitude-closure (PAC) versus the bit time for the pseudo-NMOS multiplexer implementation driving the pre-amplifier. The speed of this circuit is mainly determined by the resistance of PMOS and the total capacitance at the output node. Increasing the PMOS size relative to the NMOS size would increase the speed while reducing the swing. However, the ratio of the PMOS and NMOS sizes has to be chosen such that the swing at the multiplexer output is

enough to completely switch the output driver in the worst case (fast PMOS and slow NMOS). In this design, the swing at the multiplexer output is approximately 1/3-1/4 of the full swing. If we allow a maximum of 10% PAC, this circuit can operate at $2\text{-}\tau_4$, or 250-ps in 0.25- μm technology, as shown in Figure 4.5. An intermediate stage of pre-amplifier is inserted between the multiplexer and the output driver to reduce the size of the power-hungry multiplexer. Since the source-coupled pre-amplifier stage and the output driver stage are not the speed bottleneck, no swing reduction is applied to them at full output drive and they provide mostly current gain.

The transmitter and receiver both include 50- Ω PMOS termination resistors with 18 bits of thermometer-coded control. The adjustment step is 5%. In order for the PMOS transistor to work well as a resistor, the output swing should be kept well inside its linear

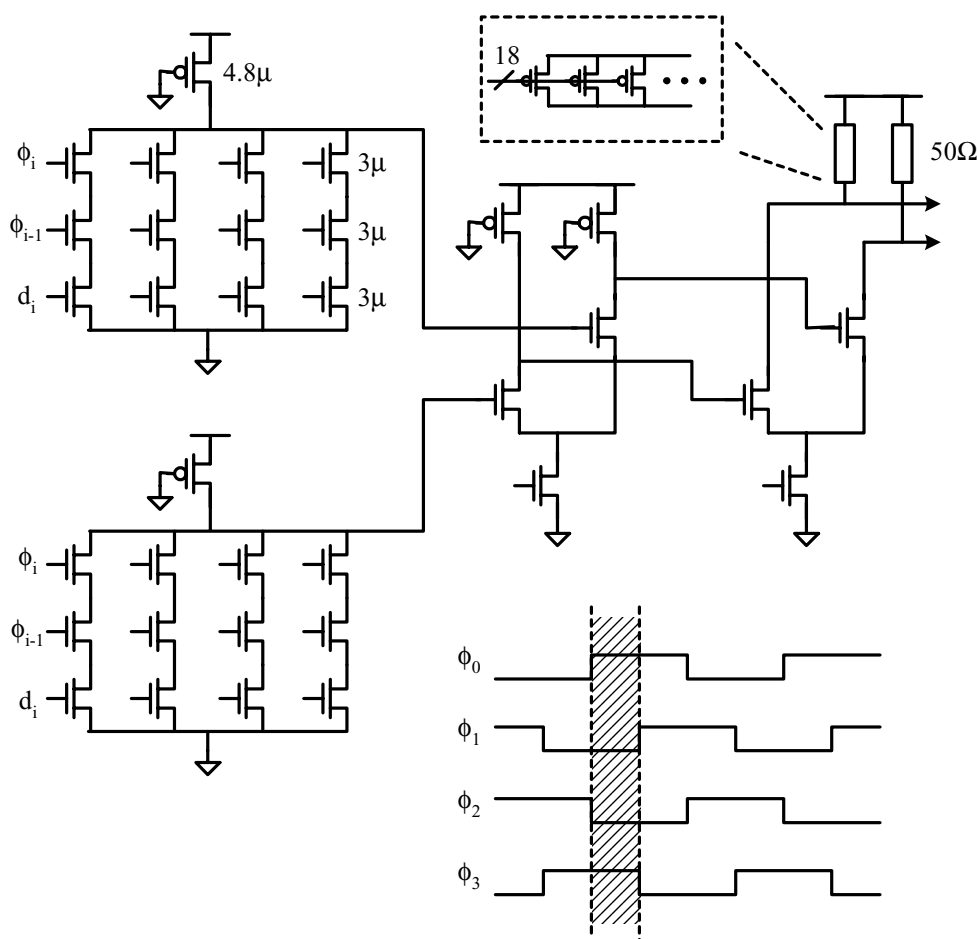


Figure 4.4: Transmitter circuit implementation.

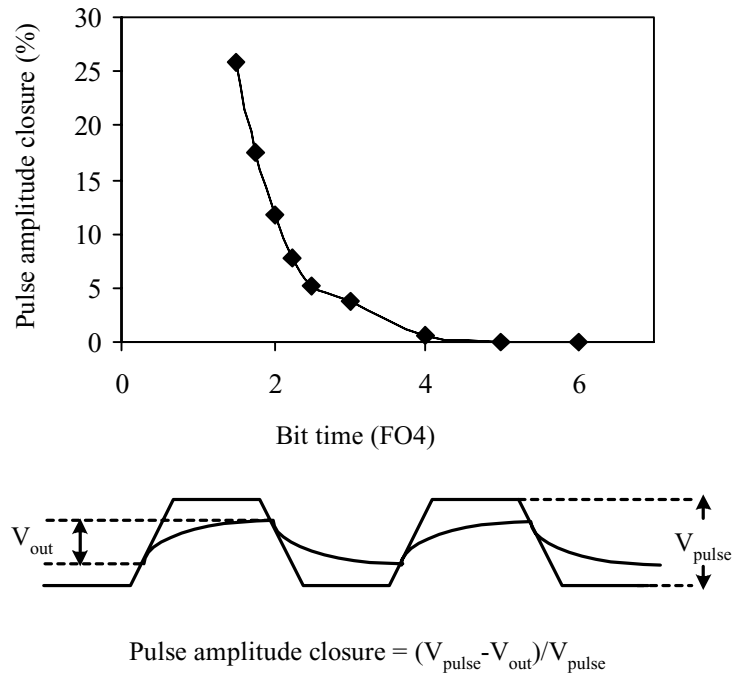


Figure 4.5: Effect of bit-time on pulse amplitude closure for the 4:1 pseudo-NMOS multiplexer of Figure 4.4.

regime. In our implementation, for example, to avoid more than 10% of resistance variation, the single-ended swing needs to be $< 200\text{-mV}$.

Before the data can be fed into the multiplexers from a single clock domain (for example, from the PRBS generator), they need to be resynchronized to per-phase clock domains to ensure proper timing margins. Figure 4.6 shows the resynchronization circuit. The digital clock, $dclk$, which is generated from the transmitter multi-phase DLL and has the same phase position as ϕ_0 nominally, is also used to clock the preceding digital circuits such as the PRBS generator. The second tap data are delayed by one bit time from the first tap data.

The two-tap FIR filter is implemented by summing two transmitter drivers directly at the output pin (effectively a 2-bit digital-to-analog converter), as shown in Figure 4.7 [9] [19]. The tap coefficients are adjusted by varying the output current of the two output

drivers. They can also be made programmable for different channels with simple current mirrors.

4.3 Summary

The architecture and circuit implementation of the transmitter were discussed and compared with previous approaches. A clock frequency of 1-GHz is chosen to ease the timing circuit design, clock distribution, and data synchronization. To achieve a bit rate of




Figure 4.6: Transmitter resynchronization circuit.

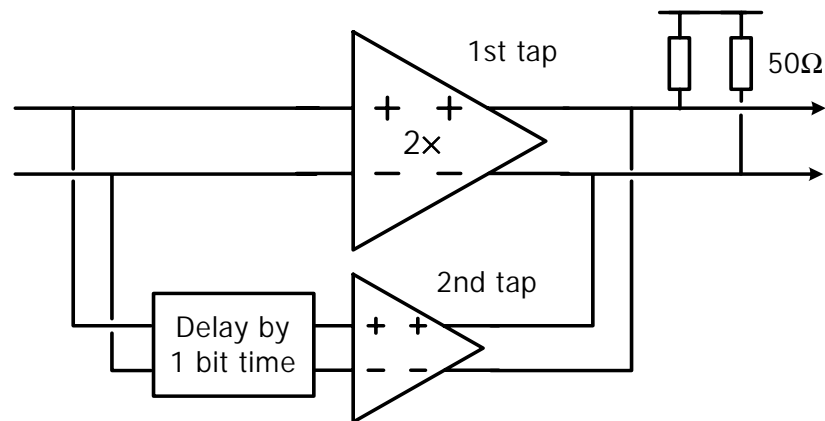


Figure 4.7: Two tap transmitter equalization filter is implemented with analog current summing.

4-Gb/s, we use a swing-optimized input-multiplexed transmitter to serialize parallel data. This implementation allows for an area-efficient design while achieving a 2× speed increase compared to a CMOS implementation. As ESD protection design becomes increasingly difficult due to higher silicon doping in deep-submicron technologies, the advantage of an output-multiplexed transmitter becomes less significant, making an input-multiplexed design a more attractive choice in the future. The transmitter is source-terminated with a digitally controlled PMOS resistor and achieves two-tap equalization through direct analog current summing at the output.

Chapter 5

Receiver

A major part of a serial I/O's voltage budget is for overcoming the inaccuracy of the receiver. Many designs today use insensitive receivers with uncanceled offset, resulting in a large amount of wasted power just to overcome these imperfections. The voltage margin of an I/O system can be expressed as

$$V_{NM} = \frac{1}{2}V_{SW} - K_N V_{SW} - V_{NF} \quad (5.1)$$

where V_{NM} is the noise margin, V_{SW} is the signal swing, K_N is the proportional noise factor, and V_{NF} is the fixed noise. The required swing is given by

$$V_{SW} = \frac{V_{NM} + V_{NF}}{0.5 - K_N} \quad (5.2)$$

$$\frac{dV_{SW}}{dV_{NF}} = \frac{1}{0.5 - K_N} \quad (5.3)$$

For $K_N = 20\%$, any reduction in the fixed noise (such as receiver offset) results in 3 times as much reduction in the required swing for a given noise margin. For example, if the

receiver offset is reduced from 40-mV to 10-mV, the required swing would decrease by 90-mV.

Besides increasing the voltage margin, a smaller voltage swing requirement also increases the timing margin, as pointed out in Chapter 2. This chapter introduces an offset trimming method which reduces the input sense amplifier offset to < 8 -mV. Besides increasing the voltage and timing margin, this scheme also decreases power and increases the input bandwidth since small sense amplifiers can be built without creating excessively large offset. To relieve the frequency requirement of the receiver, we place multiple sense amplifiers at the receiver input to perform 1:4 demultiplexing on multiple phases of the 1-GHz clock, resulting in a receiver bit rate of 4-Gb/s.

Section 5.1 describes the architecture of the receiver. The circuit implementation is presented in Section 5.2, followed by a summary at the end. The discussion of the multi-phase generation and clock recovery circuits is deferred until Chapter 6.

5.1 Architecture

Like the transmitter, the receiver achieves a bit rate higher than the clock frequency by using multiple phases of the clock to sequence the data stream. Multiple copies of the receive sense amplifiers are connected to the receiver input to directly sample the incoming data on different phases of the clock. Figure 5.1 shows the receiver architecture. Because this design uses an edge-sampling tracking clock recovery scheme (described in Chapter 6), a total of 8 sense amplifiers are attached to the input to achieve 1:4 demultiplexing. With this architecture, the bit rate is limited by the aperture time, not the cycle time, of the sense amplifier.

The switches at the front end are used to short the differential inputs during offset calibration, which is described in Section 5.2. The sampled data are processed by a resynchronization circuit to bring them from per-phase clock domains to a single clock domain. To ease the speed requirement of the subsequent digital circuits such as the PRBS

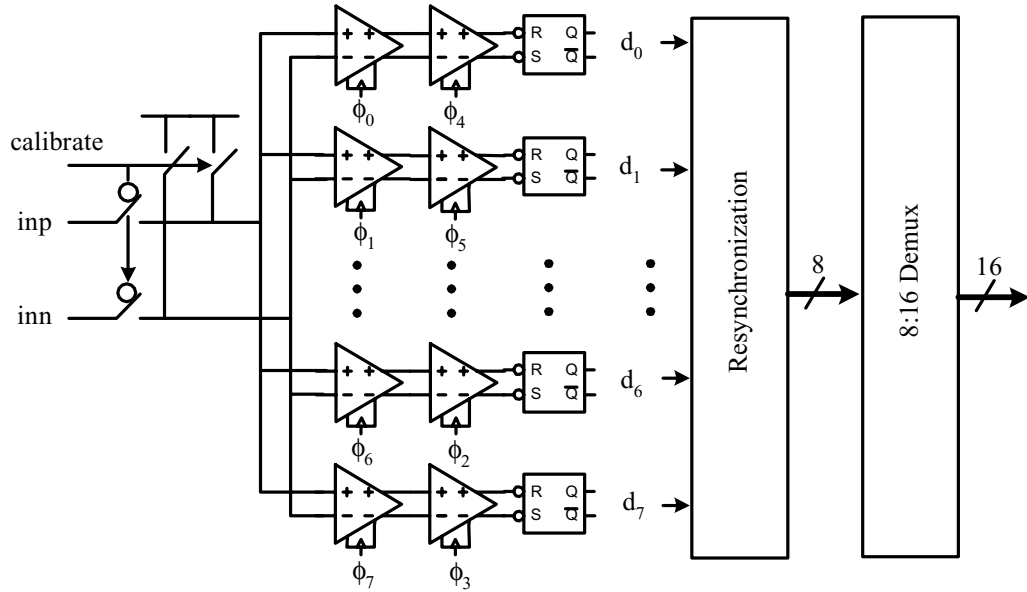


Figure 5.1: Receiver architecture.

checker and the clock recovery control, the data are further demultiplexed by an 8:16 demultiplexer, bringing the final clock frequency to 500-MHz with a 4-Gb/s bit rate.

5.2 Circuit Implementation

The receive sense amplifier, shown in Figure 5.2, is a modified version of the StrongArm sense amplifier with capacitively trimmed offset voltage [15]. We trim the sense amplifier by placing 4-bit binary-weighted PMOS capacitors on the two integrating nodes directly above the input transistors (node a and b). Digitally adjusting the capacitance while shorting the inputs unbalances the amplifier to cancel the offset voltage. It is important to pre-charge node a and b to V_{dd} with transistors $M7$ and $M8$ to increase the effectiveness of trimming capacitors. Otherwise, node a and b would only get charged to approximately $V_{dd} - V_t$ and the effect of the trimming would not be as pronounced. Although this implementation uses trimmable capacitance to introduce imbalance, other means, such as trimmable current shown in dashed-line in Figure 5.2, are possible [26].

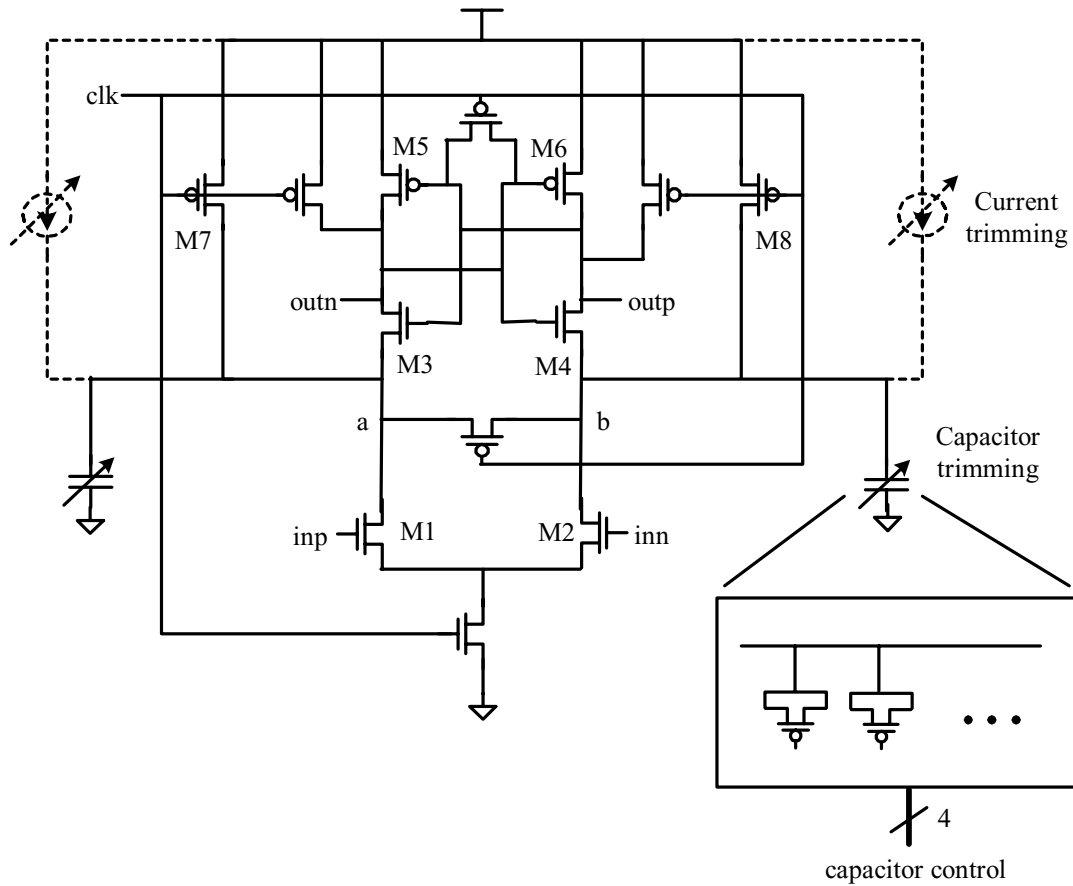


Figure 5.2: Receive sense amplifier with static offset trimming.

Furthermore, this technique can be applied to other receiver topologies, such as the current integrating receiver and the Yukawa sense amplifier, to cancel offset [8] [14].

It remains to be seen how much offset is to be expected from this receiver design in order to determine how much offset cancellation is required. [27] develops and verifies by experiment a model for the dependence of the variations in threshold voltage and transconductance on transistor size. The standard deviations can be expressed by

$$\sigma_{V_t} = \frac{A_{V_t}}{\sqrt{WL}} + S_{V_t}D \quad (5.1)$$

$$\frac{\sigma_{\beta}}{\beta} = \frac{A_{\beta}}{\sqrt{WL}} + S_{\beta}D \quad (5.2)$$

where $\beta = \mu_n C_{ox}$, D is the distance between devices to be matched, and A_{V_t} , A_β , S_{V_t} , S_β are experimentally determined parameters. [27] shows that the distance parameters S_{V_t} and S_β are quite small. Also, in this study we are mainly interested in matching properties of transistors that are close together. We therefore ignore the second term. From the A_{V_t} and A_β parameters provided by the fab, the receiver offset is simulated and shown in Table 4. The simulation is done in the worst corner (slow transistors, 2.25-V, 100°C). Table 4 shows that the offset mainly comes from the matching between the input transistors M1 and M2. The overall offset can be calculated by doing a statistical sum as follows.

$$V_{offset} = \sqrt{15.4^2 + 8^2 + 1.2^2} = 17.4 \text{ mV} \quad (5.3)$$

A 6σ coverage ($1 - 2 \times 10^{-9}$) requires 105-mV of offset cancellation capability.

Table 4: Receiver offset breakdown

Offset due to matching between	Offset
M1 and M2	15.4-mV
M3 and M4	8-mV
M4 and M5	1.2-mV

Offset calibration can be done either in software via on-chip scan interface or in hardware. The trimming capacitors in this design introduce up to ± 120 -mV of offset out of the 16 steps in 8-mV increments. With simple bang-bang control in which one of the dithering states are arbitrarily chosen, the worst case offset after cancellation is 8-mV. A more sophisticated averaging algorithm can be used to halve this number. The adjustment step can be reduced by decreasing the trimming capacitor size or increasing the receiver

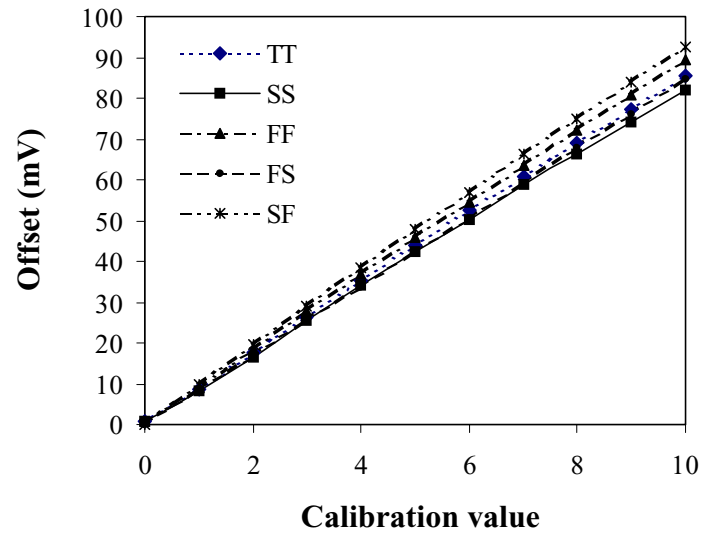


Figure 5.3: Trimmed offset versus the calibration value for the capacitively trimmed sense amplifier.

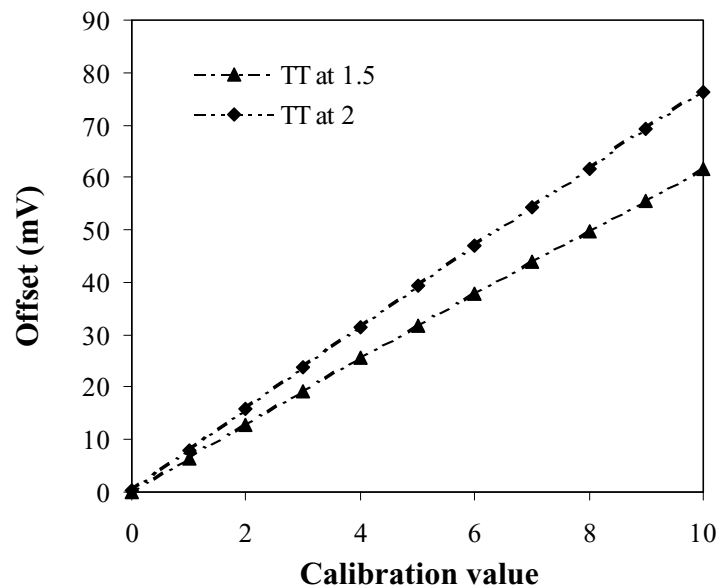


Figure 5.4: Trimmed offset versus the calibration value for the capacitively trimmed sense amplifier with different input common-mode levels.

regeneration current. Figure 5.3 shows the introduced offset versus the calibration value across the temperature (0°C - 100°C) and the supply (2.25V - 2.75V) corners in simulation. The step size variation is about 10% ($\pm 5\%$). Accounting for this variation, the worst case offset after calibration is approximately given by:

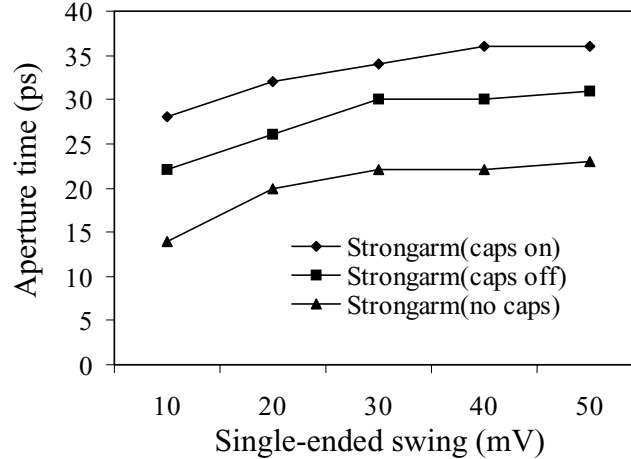


Figure 5.5: Aperture time of the capacitively trimmed sense amplifier.

$$V_{\text{off,after}} = 4mV + 0.1 \times V_{\text{off,before}} \quad (5.4)$$

where $V_{\text{off,after}}$ is the offset voltage after calibration and $V_{\text{off,before}}$ before. Figure 5.4 shows the variation of offset with two different input common-mode levels. The variation is around 20%. If the common-mode level is expected to change (for example when pre-emphasis setting is adjusted), the offset must be recalibrated to avoid this variation.

As process technology scales, transistor matching becomes more difficult. This scheme allows small and low-power receive sense amplifiers to be constructed without the penalty of large offset voltage. In serial link applications, this translates to an increase in both voltage and timing margins. A small receiver also reduces the input capacitance and increases the input bandwidth.

Figure 5.5 shows the simulated aperture time of three variations of the capacitively trimmed sense amplifier: with all capacitors *switched on*, with all capacitors *switched off*, and without any trimming capacitor *attached*. The simulation assumes a clock rise/fall time of 100-ps. The worst case aperture time with all trimming capacitors switched on is ~ 35 -ps ($1 \tau_1$), ~ 15 -ps (6% of the 250-ps bit time) worse than the bare sense amplifier. The




Figure 5.6: Receiver hysteresis versus clock period.

increase in aperture time due to offset cancellation does not impose a significant degradation to the overall timing budget.

Figure 5.6 shows the hysteresis of the sense amplifier versus the clock frequency. When the clock period falls below $6\tau_4$, hysteresis degrades considerably. This limit is due to insufficient pre-charge time, which leaves significant residual memory at the internal nodes of the sense amplifier, preventing the next regeneration from making a correct decision. Since this limit is much larger than 8 times the receiver aperture time, the receiver bit rate is limited by the clock frequency, not the receiver aperture time, suggesting a higher bit rate should be achieved by a higher demultiplexing width instead of clock frequency.

An SR latch is used at the sense amplifier output to produce a stable digital signal which does not have the pre-charge cycle of the sense amplifier. It is important, however, that the front-end sense amplifier be followed by another stage of sense amplifier clocked by the opposite phase before the SR latch. This is shown in Figure 5.7. If the front-end sense amplifier is connected to the SR latch directly, hysteresis on the order of a few mV

would be present even at 1-GHz (4-Gb/s) due to data-dependent load of the SR latch (i.e. the input capacitance of the SR latch is different for state 1 and 0).

Simulation indicates that the regeneration time constants of the first stage and second stage sense amplifiers are approximately 60-ps and 70-ps. The overall gain of the two stages is given by

$$A = \exp\left(\frac{t_{cyc}/2}{60p}\right) \cdot \exp\left(\frac{t_{cyc}/2}{70p}\right) \quad (5.5)$$

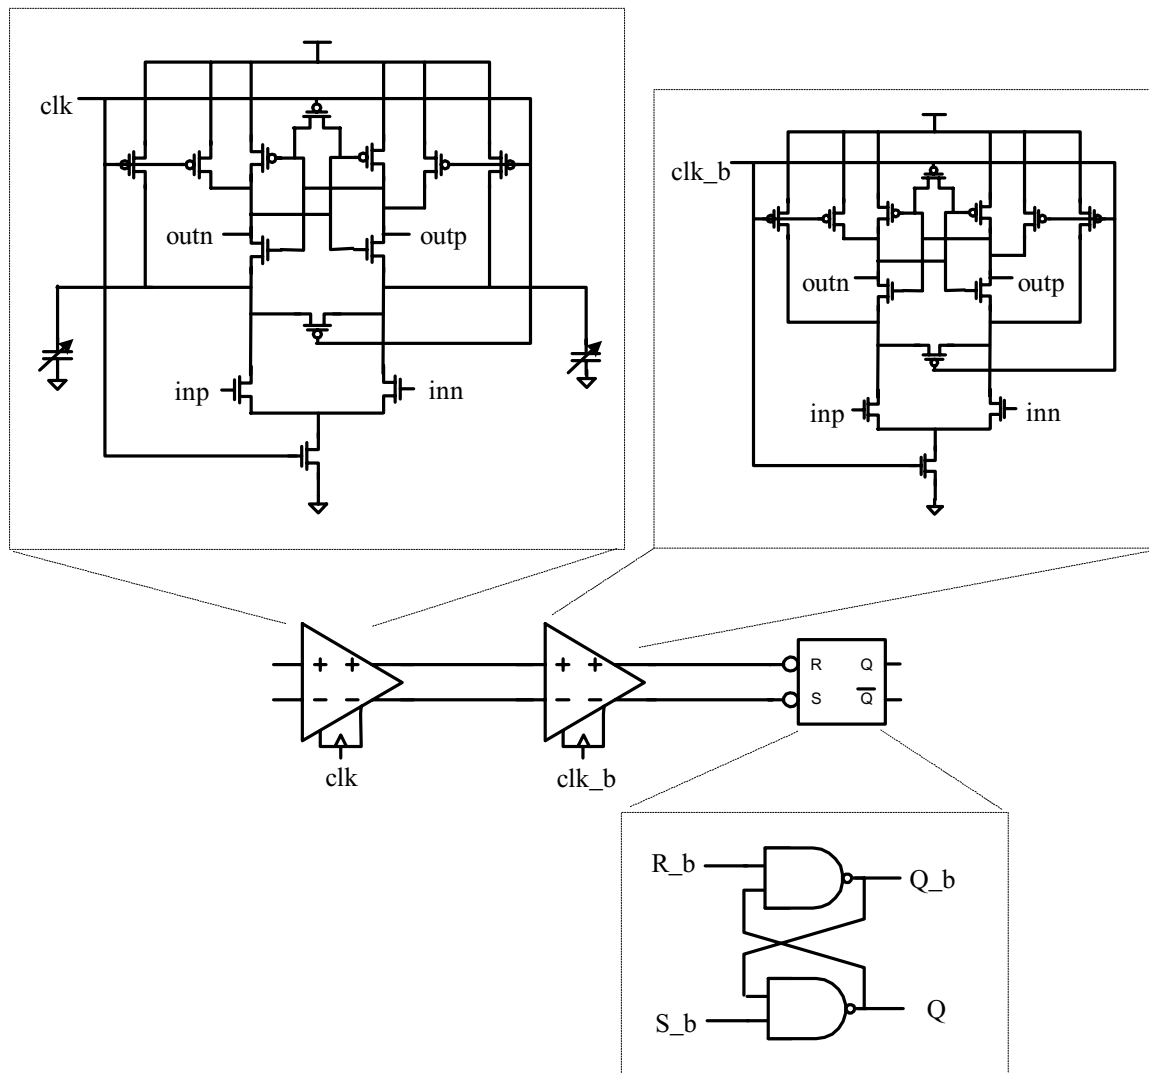


Figure 5.7: A second stage of StrongArm latch is inserted to reduce the hysteresis of the input.

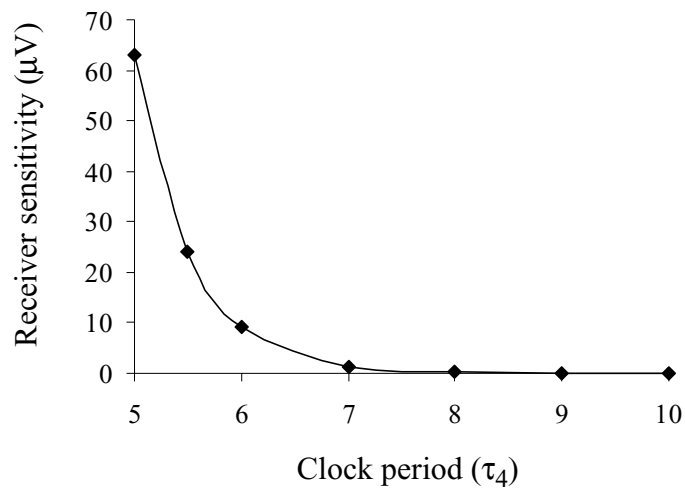


Figure 5.8: Receiver sensitivity versus clock period.

Figure 5.8 shows the sensitivity of the receiver given this total gain, assuming that at least 1-V is required before the SR-latch to make a reliable decision. The plot indicates that the sensitivity of the receiver is much smaller than the hysteresis (Figure 5.6) and is not a factor in the overall receiver voltage budget.

The sampled data must be resynchronized before they can be processed by the subsequent digital circuits in a single clock domain. Figure 5.9 shows the resynchronization circuit. The digital clock, *dclk*, which is generated in the receiver multi-phase DLL and has the same phase position as ϕ_0 nominally, is also used as the clock source for the subsequent digital circuits such as the PRBS checker and the clock recovery control.

5.3 Summary

Eight offset-cancelled sense amplifiers connected to the receiver input perform data and edge (for clock recovery) sampling on 8 phases of the receive clock to achieve a bit rate 4 times the on-chip clock frequency. With this architecture, the bit rate is limited by the aperture time, not the cycle time, of the sense amplifier, given enough

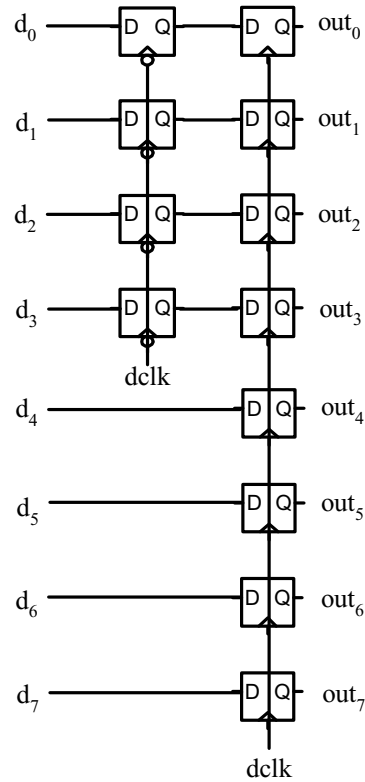


Figure 5.9: Receiver resynchronization circuit.

demultiplexing width. Simulations show, however, that with a demultiplexing width of 8, the receiver bit rate is still limited by the clock cycle time. This suggests that, at 4-Gb/s, a higher bit rate should be achieved by a higher demultiplexing width instead of clock frequency. Offset-cancellation is implemented by placing trimmable capacitance at the internal nodes of the sense amplifiers to create imbalance. This scheme allows small and low-power receiver to be built without compromising the offset voltage of the circuit. Aperture time penalty of the offset cancellation is only 6% of the target bit time and has negligible effect on the achievable speed.

Chapter 6

Timing Circuits

In Chapter 3, we introduced three major components of a timing budget: signal rise time (t_r), receiver aperture time (t_d), and timing uncertainty (t_u). Signal rise time and receiver aperture time relate to the characteristics of the transmitter and the receiver, which are the subject of the last two chapters. In this chapter, we discuss the design of the timing circuits, which affect the amount of timing uncertainty in the system. A multi-phase delay-locked loop is needed to perform serialization on the transmit side and deserialization on the receive side. To accurately sample the incoming data with the maximum timing margin, the receiver also requires a clock recovery unit to position its clock at the center of the data eye. Our design goal of the timing circuits is to perform these functions with low power, small area, and noise immunity for dense integration. Section 6.1 describes the multi-phase delay-locked loop design. The clock recovery unit employing the dual DLL architecture [12] is described in Section 6.2. Finally, a summary is given at the end to conclude this chapter.

6.1 Multi-Phase Delay-Locked Loop

The delay-locked loop generates evenly spaced clock phases at 1-GHz to sequence the multiplexer at the transmitter and the demultiplexer at the receiver. As shown in Figure

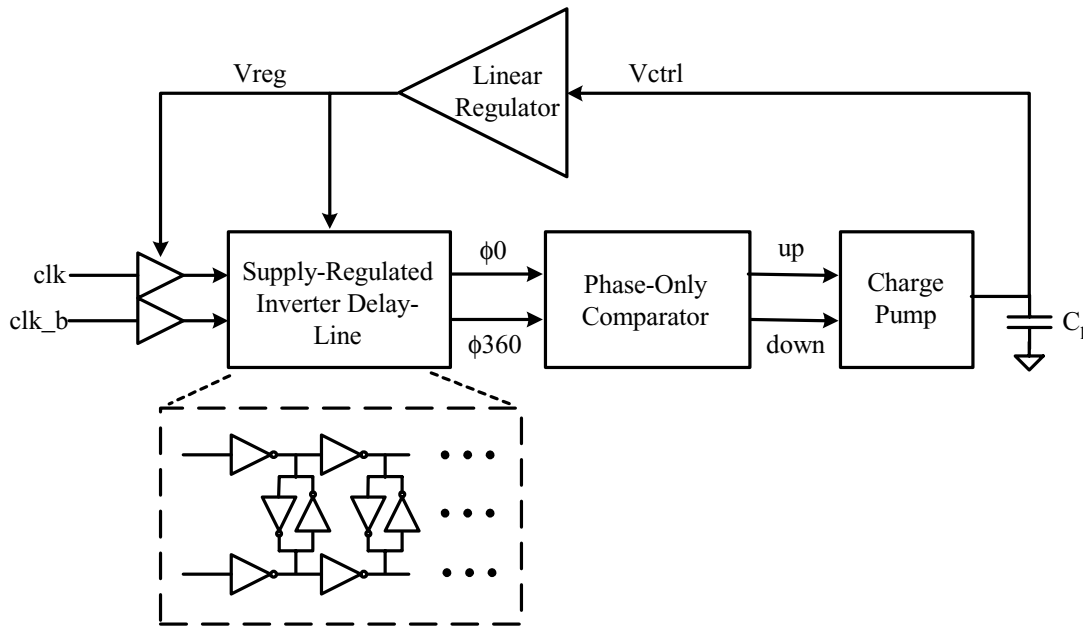


Figure 6.1: Supply-regulated inverter delay-locked loop architecture.

6.1, the DLL uses a differential CMOS inverter delay line with a regulated power supply. Operating differentially generates fine phases and the complementary outputs simplify the level shifter. The delay-line only needs to generate 180° of delay, avoiding jitter due to excessively long delay-lines. The true and complement outputs of each stage are cross-coupled with weak inverters to minimize skew. Delay is adjusted by varying the supply voltage with a linear voltage regulator. In [28], a PLL is built using a similar approach in




Figure 6.2: Source-coupled differential delay element.




Figure 6.3: Source-coupled delay element with two-element PMOS load and replica-bias.

which the delay of a CMOS inverter ring oscillator is adjusted by varying its supply current through a current regulator.

Source-coupled differential structures, shown in Figure 6.2 are widely used as delay elements due to their low supply sensitivity [6] [8] [11] [12] [13] [19]. The delay of such structures is adjusted by varying the resistance of the load, R_l , while the delay-line swing is set by the product of its tail current, I_t , and R_l . Changes in power supply voltage change only the common-mode voltage of the swing and has little effect on the delay, assuming the load resistance is linear. A common implementation of load resistance is a two-element PMOS structure shown in Figure 6.3. The tail current is adjusted via a replica bias circuit so that the output swings between V_{dd} and $V_{dd} - V_{c-sc}$, where the load is approximately resistive. The delay per stage is inversely proportional to the conductance of the load, which is given by:

$$g_{load} = 2\beta_p(V_{c-sc} - V_{tp}) \quad (6.1)$$

Compared to a CMOS inverter delay element, a source-coupled delay element requires more power. Also, as technology scales and V_{dd}/V_t ratio shrinks, a source-coupled delay element runs out of headroom and a CMOS inverter implementation becomes much more portable and simpler to design.

The power consumed in an N-phase differential inverter delay line and differential source-coupled element delay-line are:

$$P_{inv} = N \times \frac{C_t V_{c-inv} V_{dd}}{T_c} = \frac{C_t V_{c-inv} V_{dd}}{T_d}, \text{ inverter element based} \quad (6.2)$$

$$P_{sc} = \frac{N}{2} \times \frac{C_t (V_{dd} - V_{c-sc}) V_{dd}}{T_d}, \text{ source-coupled element based} \quad (6.3)$$

where T_d is the stage delay, T_c is the clock cycle time, V_{c-inv} is the inverter delay-line control voltage, V_{c-sc} is the source-coupled delay-line control voltage, and C_T is the total capacitance charged at each stage during a cycle. C_T is roughly the same for both with the same accuracy requirement. The inverter delay-line requires twice as many delay element compared to the source-coupled element delay-line since the availability of differential signals is assumed. If $V_{c-inv} = k \times (V_{dd} - V_{c-sc})$, this analysis indicates that an inverter delay-line consumes $2k/N$ of the power consumed by a source-coupled delay-line, where k is usually between 1 and 1.5. Intuitively, the factor N comes from the fact that an inverter delay-line does not consume any static current. For our case, N is 4 and the power reduction is $0.75 \sim 0.5$. As N increases (in applications such as clock recovery or when the width of multiplexing is increased), power saving is more significant.

Since the two branches of the inverter delay-line are not truly differential, cross-coupled inverters are inserted to minimize skew between the two lines. These cross-coupled inverters are weak compared to the delay inverters, but are very effective in eliminating skew and hence reduce phase spacing imbalances. In this design, the size of the cross-coupled inverters are 1/4 the size of the delay inverters. Inserting input clock skews of up to 200-ps results in phase spacing difference by less than 40-ps. Without the cross-coupled inverters, this number would increase to 400-ps.




Figure 6.4: Phase-only comparator employed in the multi-phase DLL.

Figure 6.4 shows the schematic of the phase-only comparator [29]. The job of the phase-only comparator is to generate *up* and *down* pulses with durations proportional to the phase difference between ϕ_0 and ϕ_{360} of the delay-line. It is different in functionality from a phase frequency detector (PFD) widely used in a phase-locked loop. A PFD cannot be used in a delay-locked loop since a delay-line, unlike an oscillator, cannot wrap its phase to force the PFD to the correct state when incorrectly initialized or disturbed. Figure 6.5 shows how an incorrect initial state pegs a delay-locked loop at one end of its delay range when a PFD is used. In this example, we assume that ϕ_{360} is derived from *clk* and is




Figure 6.5: Timing diagram showing incorrect use of a PFD in a DLL.

compared against clk_b (ϕ_0) so that the maximum delay in each of the differential branches is only 180 degrees to avoid excessively long delay-lines. An *up* pulse decreases the delay and a *down* pulse increases the delay. As shown in the timing diagram, because the positive edge of ϕ_0 appears first at startup, the loop tries to shorten the delay of the delay-line even though the delay is too short. Eventually the delay-line would be pegged at its shortest-delay end without any means of correction. A phase-only comparator resets itself on the falling edge of either ϕ_0 or ϕ_{360} so that this situation would not occur. Short *up* and *down* pulses with equal duration at lock avoid deadband which results in static phase offset.

Figure 6.6 shows the schematic of the charge pump. It takes the *up* and *down* signals from the phase-only comparator and pumps a proportional amount of charge to its output capacitance. A critical design objective is to ensure that the up and down charges are equal at lock. A charge imbalance is usually caused by transistor mismatches, capacitive charge injection, and channel length modulation which results in up and down current mismatch. Transistor mismatches are reduced by using large transistors; capacitive charge injection are mitigated by moving the switching transistors M1 and M2 to the bottom of the current biasing transistors M3 and M4 and by using large pump current; current mismatch due to channel length modulation is reduced by careful design of the




Figure 6.6: Charge pump employed in the multi-phase DLL.

phase-only comparator and the charge pump. Assuming the phase-only comparator generates equal *up* and *down* pulses at lock, then the up and down charge difference due to channel length modulation is given by:

$$\Delta Q = (I_{up} - I_{down}) \times t_s = \Delta I \times t_s \quad (6.4)$$

where ΔI is the difference between up and down current of the charge pump due to channel length modulation, and t_s is the width of the up and down pulses at steady state. Equation (6.4) indicates that to minimize the phase offset, ΔI and t_s should be minimized. This phase-only comparator design generates *up* and *down* pulses of extremely short duration at steady state due to its simplicity and non-feedback operation. To minimize ΔI resulting from channel length modulation, a high swing cascode circuit is employed in the charge pump.

The charge pump current, I_p , is made proportional to V_{ctrl}^2 with M5 ($I_p = K_1 V_{ctrl}^2$). The inverter delay-line gain K_{dl} is proportional to $1/V_{ctrl}^2$ ($K_{dl} = K_2/V_{ctrl}^2$). The loop bandwidth ω_1 of the DLL is thus fixed at a constant fraction of the input frequency as seen by:

$$\omega_1 = \frac{I_p K_{dl} \omega_{in}}{2\pi C_l} = \frac{(K_1 V_{ctrl}^2) \times (K_2/V_{ctrl}^2) \times \omega_{in}}{2\pi C_l} = \frac{K_1 K_2}{2\pi C_l} \omega_{in} \quad (6.5)$$

where ω_{in} is the input clock frequency and C_l is the loop filter capacitance at the output of the charge pump. The voltage regulator, which has a -3 -dB bandwidth in excess of 300-MHz, does not affect the loop dynamics. Keeping the loop bandwidth at a constant fraction of the input clock frequency ensures the stability of the loop across a wide range of input frequencies as well as process, voltage, and temperature (PVT) corners.




Figure 6.7: Linear voltage regulator employed in the multi-phase DLL.

The purpose of the linear voltage regulator, shown in Figure 6.7, is to drive the delay-line and shield it from the supply noise while having enough bandwidth to ensure the overall stability of the DLL. Stability compensation is employed at the output of the one-pole feedback amplifier. Without the compensation, the pole at the output of the feedback amplifier (V_m) would be very close to the pole at the regulator output (V_{out}), making the circuit unstable. This is because the output resistance of the feedback amplifier is usually much larger than that of the regulator for gain and the output capacitance of the regulator is much larger than that of the feedback amplifier for good power supply rejection. We use a simplified circuit model of Figure 6.8 to analyze the approximate frequency response of this circuit. r_l is the effective output resistance of the feedback




Figure 6.8: Simplified model of the linear voltage regulator.

amplifier; c_2 is the effective capacitance at the output of the feedback amplifier, taking into account the Miller effect; r_o is the output resistance of the regulator (V_{out}), including the effective resistance of the regulated circuits; and c_o is the effective capacitance at the regulator output. Whereas the circuit without compensation has two dominant poles which are close to each other, this circuit has one dominant zero at the feedback amplifier output and three dominant poles, one at the regulator output and two at the feedback amplifier output. In this design, we use the zero to approximately cancel the pole at the regulator output. The remaining two poles are approximately given by

$$p_{1,2} = \frac{-(r_1c_1 + r_1c_2 + r_2c_1) \pm (r_1c_1 + r_1c_2 + r_2c_1) \sqrt{1 - \frac{4r_1r_2c_1c_2}{(r_1c_1 + r_1c_2 + r_2c_1)^2}}}{2r_1r_2c_1c_2} \quad (6.6)$$

$$\approx \frac{-(r_1c_1 + r_1c_2 + r_2c_1) \pm (r_1c_1 + r_1c_2 + r_2c_1) \left(1 - \frac{1}{2} \frac{4r_1r_2c_1c_2}{(r_1c_1 + r_1c_2 + r_2c_1)^2}\right)}{2r_1r_2c_1c_2}$$

$$= \begin{cases} \frac{(r_1c_1 + r_1c_2 + r_2c_1)}{r_1r_2c_1c_2} \left(1 - \frac{r_1r_2c_1c_2}{(r_1c_1 + r_1c_2 + r_2c_1)^2}\right) \\ -\frac{(r_1c_1 + r_1c_2 + r_2c_1)}{r_1r_2c_1c_2} \left(\frac{r_1r_2c_1c_2}{(r_1c_1 + r_1c_2 + r_2c_1)^2}\right) \end{cases}$$

where the approximation from the first equation to the second equation is valid for typical component values. The separation of the two poles is determined by the magnitude of $k = r_1r_2c_1c_2 / (r_1c_1 + r_1c_2 + r_2c_1)^2$, which can be decreased to increase the phase margin. In this design, the worst case phase margin is chosen to be $\sim 50^\circ$ with a bandwidth in excess of 300-MHz. The stabilizing zero also desensitizes the stability of the circuit on the exact regulator output pole location ($1/r_oc_o$). Although r_o , which is mostly determined by the




Figure 6.9: Level shifter employed in the multi-phase DLL.

delay line, is inversely proportional with the frequency of the delay line, the phase margin remains above 45° across a $6\times$ frequency range.

The supply sensitivity of the regulator with a fast (100-ps) supply transition is 0.1 (peak noise) and the steady state error is 0.01. As a consequence, although inverter delay element has a large supply sensitivity of about -0.9 ($4.5\times$ worse than a source-coupled delay element), the voltage regulator reduces this number by $10\times$, resulting in overall supply sensitivity of -0.09 [21].

Figure 6.9 shows the level shifter at the output of the delay-line. It converts a V_{ctrl} level signal to full V_{DD} level. The level shifter employs circuit topologies that have opposite supply sensitivity connected in series to cancel noise from the unregulated power supply and to reduce the steady state phase error. The first stage current mirror amplifier has a positive supply sensitivity since its input swing is fixed by the regulator while the output swing changes with the unregulated supply. The subsequent inverter stages, on the other hand, have a negative supply sensitivity. The relative fan-out of the two stages is tuned so that the combined supply sensitivity of the delay-line and the level shifter is minimized.

The DLL locks in ~ 30 cycles (30-ns). Figure 6.10 shows the simulated delay from a fixed clock source V when the DLL undergoes a 10% supply pulse with 100-ps rise/fall

time. The simulation is done in the worst case corner (slow transistors, 2.25-V, 100°C) and includes the jitter of the output buffers. The p-p jitter is ~ 30 -ps.

6.2 Clock Recovery

6.2.1 Architecture

Figure 6.11 shows the architecture of the clock recovery unit adapted from [12]. The multi-phase DLL described in the previous section generates 8 clock phases. The absolute phase positions of the 8 clock phases are simultaneously adjusted by 4 differential timing verniers, each composed of two phase multiplexers and one phase interpolator sequenced by a phase controller. Each timing vernier selects two adjacent phases using the phase multiplexers and interpolates between them using the phase interpolator to create 8 finer phase steps. Different adjacent clock phases can be selected to achieve infinite phase rotation. Because of this property, this architecture is compatible with plesiochronous clocking between the transmitter and the receiver. Both the phase multiplexer and the phase interpolator are thermometer coded. The 8 phases generated by the timing verniers are used to sample the incoming data as well as the data edges to gather timing information. The data stream is 1:4 demultiplexed by the phases to achieve a

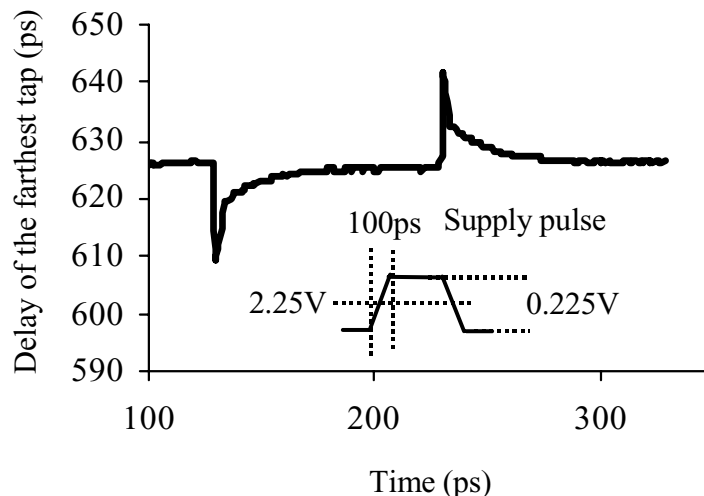


Figure 6.10: Simulated jitter due to a 10% supply pulse with 100-ps rise time.

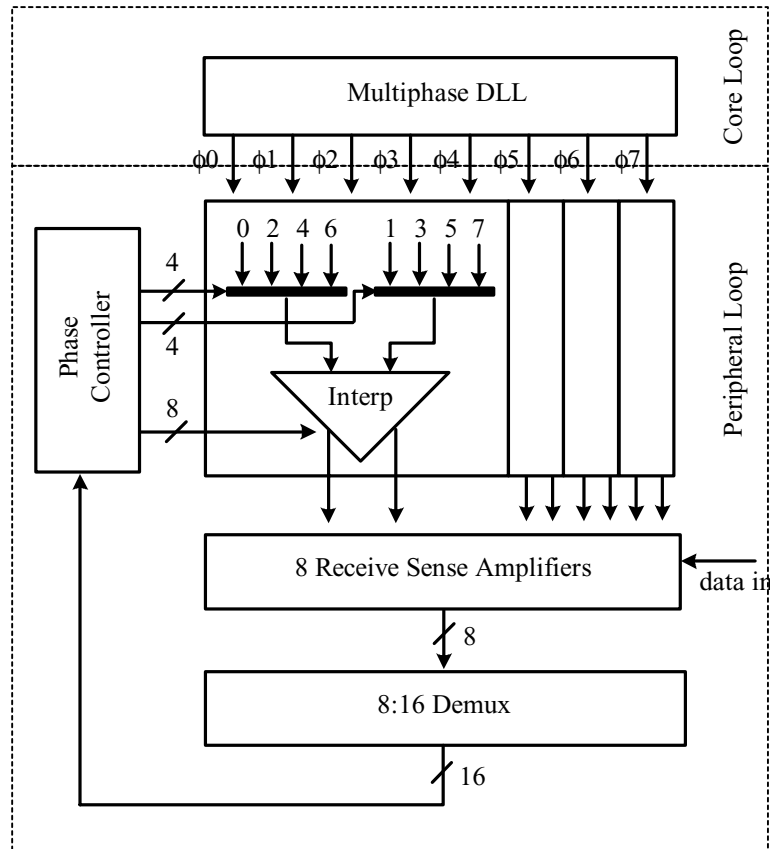


Figure 6.11: Clock recovery architecture.

bit rate which is 4 times the clock frequency, as described in Chapter 5. Before feeding into the phase controller, the resulting data samples are further demultiplexed to half the clock frequency to relax the frequency requirement of the digital logic.

The phase controller architecture is shown in Figure 6.12. It is clocked at half the receive clock frequency (500-MHz at 4-Gb/s). The 16 samples generated every cycle first go through a set of early/late decoders. The early/late decoder determines whether there is a data transition for each bit. If there is, the edge sample is used to decide whether the receive clocks are early or late. Otherwise, there is no timing information contained in that particular bit and the decoder outputs a *no_info*. The resulting 8 *early/late/no_info* signals are then resolved by a majority vote unit. The summarizing signal generated by the majority vote is low-passed filtered by an 8-bit ring counter, which updates a finite state

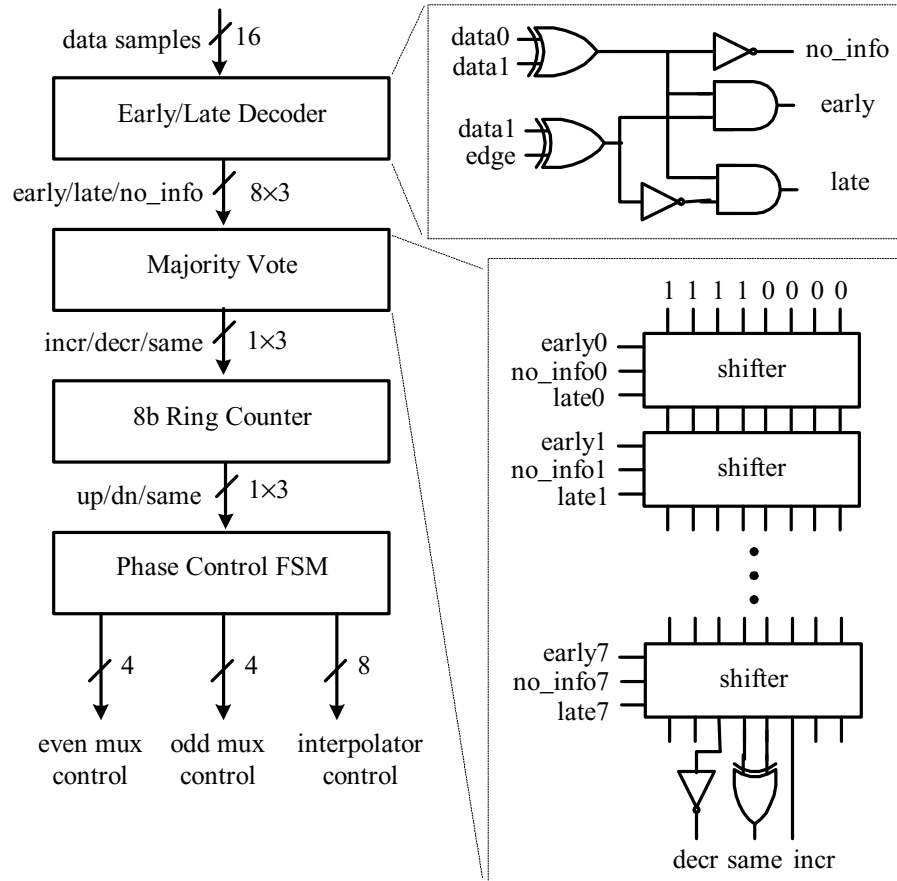


Figure 6.12: Phase controller architecture.

machine to generate the appropriate phase control signals. The purpose of the 8-bit ring counter is two-fold. It is used to decrease the peripheral loop bandwidth so that noisy input signals do not affect the clock jitter. It is also used as a low-pass filter within the peripheral loop to avoid loop instability due to the loop delay.

6.2.2 Circuit Implementation

Figure 6.13 shows the schematic of the phase interpolator. Control signals $w_{\{0-7\}}$ and $w_{\{0-7\}_b}$ are the interpolation weight. The circuit operates by assigning variable amount of strength to the $clk1$ and $clk2$ branches, thereby creating an adjustable delay that spans from $clk1+d$ to $clk2+d$, where d is the delay of the interpolator circuit. Since the interpolator adjusts the phase of the receive clock in discrete steps, it is important to minimize the maximum phase step to avoid excessive dithering at steady state. To do so




Figure 6.13: Peripheral loop interpolator.




Figure 6.14: Phase position vs. the interpolation step for current-mirror interpolator.

under a fixed number of phase steps requires a phase interpolator which sweeps phases in linear steps. Figure 6.14 shows the simulated phase position versus the interpolation step for this circuit. This plot shows the *inherent phase linearity* of the circuit without taking into account transistor or layout mismatches. The differential non-linearity (DNL) is 0.24-LSB of the interpolating interval. We compare it with a more straightforward




Figure 6.15: Tri-state inverter based interpolator.

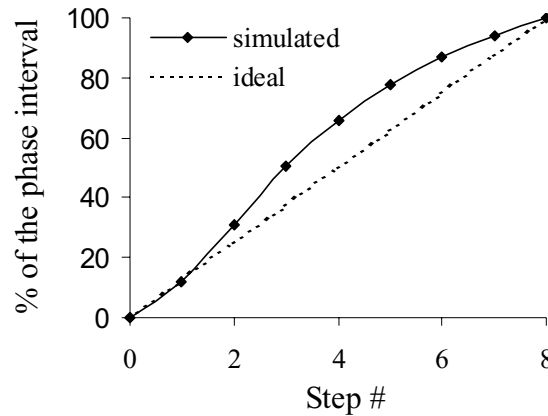


Figure 6.16: Phase position vs. the interpolation step for tri-state inverter based interpolator.

implementation based on CMOS tri-state inverters, as shown in Figure 6.15. Figure 6.16 shows the inherent phase linearity of this implementation. The DNL is 0.55-LSB. The above simulations were done at 1-GHz with a 150-ps rise time. Because both the current-mirror interpolator and tri-state CMOS interpolator are based on *phase mixing*, phase linearity invariably becomes worse at lower frequencies as the interpolating phase spacing

becomes larger but the clock rise time remains approximately fixed. At 500-MHz and the same rise time, for example, the phase DNLs of current mirror interpolator and tri-state CMOS interpolator are 1.92-LSB (60-ps) and 2.24-LSB (70-ps), respectively. If good phase linearity is desired across a wide frequency range, input clock signals should be shaped so that the rise time remains a fixed fraction of the clock cycle [12] [30].

To avoid glitches due to the phase multiplexer switching, the interpolation weight is sequenced all the way to the extreme before the phase multiplexer changes its selection. An extra phase step which nominally has no effect on the phase position of the receive clock is created at the boundary of the interpolation interval as a result. This translates into a reduction in the frequency tolerance of plesiochronous clocking. For our implementation with a 1-GHz clock, the frequency tolerance of the clock recovery can be expressed as

$$\begin{aligned} \frac{\Delta f}{f} &\cong \frac{\text{average phase step}}{\text{phase update interval} \times (\# \text{ of intervals without a transition} + 1)} & (6.7) \\ &= \frac{1 \text{ ns} / 72}{2 \text{ ns} \times 8 \times (d + 1)} = \frac{868}{(d + 1)} \text{ ppm} \end{aligned}$$

where d is the average number of intervals (where one interval is 8 bit time) without a data transition. If we assume an average d of 1 for a 20-bit PRBS, the expected frequency tolerance is $\sim \pm 434$ -ppm. This number is significantly higher than the frequency tolerance of most commercial oscillators (usually ± 50 -ppm) and should be adequate for plesiochronous clocking.

Similarly, the extra step reduces the bandwidth of the clock recovery in tracking the phase noise of the input data. However, bandwidth reduction is not critical as minimizing bandwidth subject to frequency tolerance requirement is good for filtering out noisy input signal. Since the peripheral loop is a non-linear system, its bandwidth depends on the magnitude of the input jitter. To estimate its bandwidth corresponding to a specific input jitter magnitude, we apply a sinusoidal jitter with amplitude A_j to its input and calculate the maximum jitter frequency that can be handled by the loop. The maximum slope of such input jitter should be less than the slew rate of the clock recovery below the bandwidth f_{bw} of the loop.

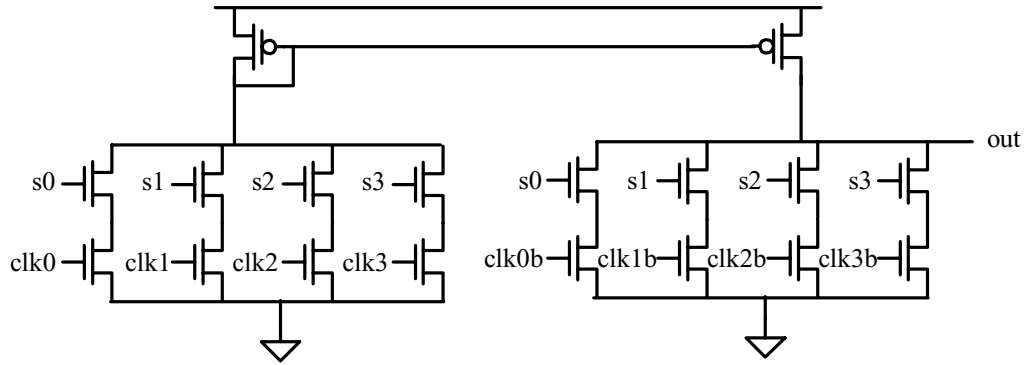


Figure 6.17: Phase multiplexer employed in the peripheral loop.

$$2\pi f_{bw} A_j = \frac{1 \text{ ns} / 72}{2 \text{ ns} \times 8 \times (d + 1)} \quad (6.8)$$

$$f_{bw} A_j = \frac{1}{2\pi(d + 1) \times 1152}$$

Assuming an average d of 1, the bandwidth of the loop is 1.38-MHz for 50-ps of jitter and 138-kHz for 500-ps of jitter.

The current mirror topology of Figure 6.13 is also used for the phase multiplexor. The schematic is shown in Figure 6.17. Besides selecting one of the 4 phases, it also performs level conversion between the regulated DLL and the unregulated peripheral circuits, obviating the need for an extra stage of level conversion. As a reference, we again compare it to a more straightforward digital implementation based on tri-state inverters, shown in Figure 6.18. Figure 6.19 shows the pulse-amplitude-closure as a function of




Figure 6.18: Tri-state inverter based phase multiplexer.

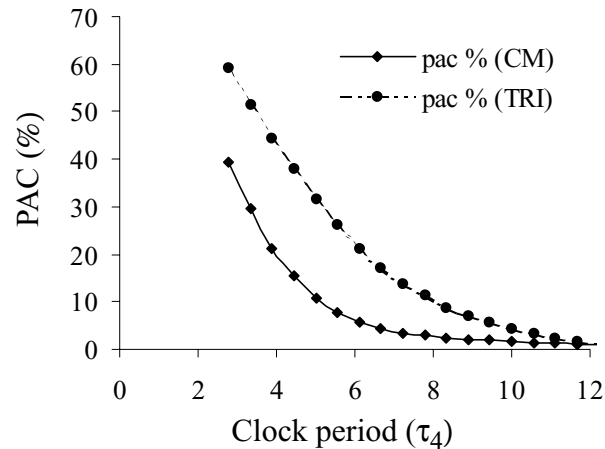


Figure 6.19: PAC for current-mirror phase multiplexer (CM) and tri-state inverter phase multiplexer (TRI).

frequency for both topologies. With a 5% maximum allowed PAC, the current-mirror topology operates at $< 6\tau_4$ (1.33-GHz in 0.25- μm) while the tri-state inverter topology operates at $> 8\tau_4$ (< 1 -GHz in 0.25- μm), which does not meet the target speed of this design. Although multiple stages can be used to increase the speed, the jitter tends to increase proportionally.

Finally, both the phase multiplexer and the interpolator are directly connected to the power supply without any regulation. The simulated p-p jitter of the combined phase multiplexer and interpolator circuit in the worst corner (slow transistors, 2.25-V, 100°C) is ~ 60 -ps for a 10% supply step with 100-ps rise/fall time. Lower jitter is possible by regulating the circuit either with a linear regulator (for example as described in this work) or with a switch regulator [30].

6.3 Summary

In this chapter, the timing circuits, which include the multi-phase delay-locked loop and clock recovery, are described and analyzed. Supply regulated CMOS inverter delay line is used in the multi-phase delay locked loop to save power and reduce jitter. By varying the supply of an inverter delay line, delay adjustment and supply rejection are simultaneously achieved. The supply regulator decreases the supply sensitivity of the

inverter delay line by $10\times$ to ~ 0.09 , making an inverter delay line a feasible design in high performance serial links.

The clock recovery uses a dual-loop architecture described in [12] due to its infinite phase range and compatibility with plesiochronous clocking. The phase controller uses a majority vote unit to combine high frequency (4Gb/s) early/late information. To ensure the stability of the peripheral loop and to filter out the noisy data input, an 8-bit ring counter is also used to deter the phase control update. The phase interpolator is implemented using a current-mirror topology for its phase linearity. The 4:1 phase multiplexer also uses the same topology. Besides operating at above the targeted frequency with a single-stage design, it also performs level conversion directly from the regulated DLL supply to the full supply, thus obviating the need for an extra stage in the clock path. The current mirror topology helps the overall timing budget by reducing the receiver clock jitter and dithering with its high bandwidth and good phase linearity.

Chapter 7

Experimental Results

Two prototype chips were fabricated in a 0.25- μm CMOS technology to verify the techniques introduced in this work. This chapter describes the experimental setups and experimental results. Section 7.1 gives a description of the two prototype chips and the experimental setups. The measurement results are given in Section 7.2, followed by a summary at the end.

7.1 Prototypes

A prototype chip containing the I/O circuits and two multi-phase delay-locked loops, one for the transmitter and one for the receiver, was fabricated first to evaluate the majority of the research ideas presented in this work. A second prototype chip added the receiver clock recovery unit to complete the transceiver design. The two chips were fabricated in a 0.25- μm CMOS technology. The active area of the transmitter is 0.08- mm^2




Figure 7.1: Die photomicrograph of the first prototype chip.

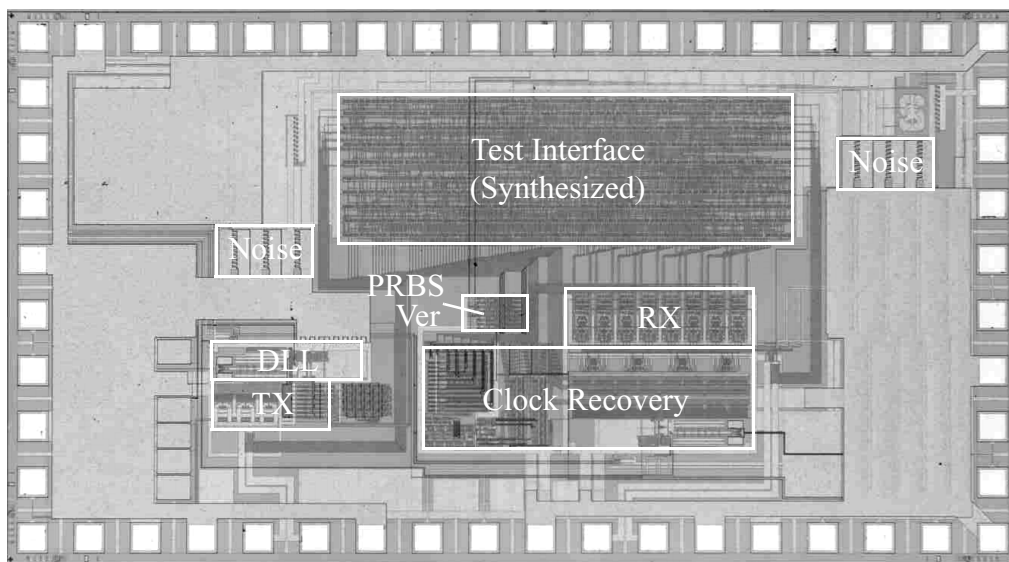


Figure 7.2: Die photomicrograph of the second prototype chip.

and that of the complete clock and data recovery (CDR) circuits is 0.3-mm^2 . Figure 7.1 and Figure 7.2 show the die photomicrographs of the first and second prototype chip respectively. The die sizes of the first and second prototypes are $2\times 2\text{ mm}^2$ and 2.6×1.4

mm². Besides the transceiver circuits, each chip contains a 20-bit PRBS generator and a 20-bit PRBS checker. Two noise generators are also included on the second chip to test the noise sensitivity of the circuits. The schematic of the noise generator is shown in Figure 7.3 [12]. A 300- μm NMOS transistor is connected between the supply and ground of the




Figure 7.3: On-chip noise generator and noise monitor.




Figure 7.4: Experiment setup of the serial link prototype.

transmitter and the CDR to create supply noises. It is capable of creating supply noises with a rise time <1 -ns. A monitoring device connected to the internal supply provides visibility to the supply noise. Both chips were packaged in 52-pin ceramic leaded chip carrier (LDCC) packages with internal power planes for impedance control.

Figure 7.4 shows a picture of the test board and the test setup. It has four signal layers and two power planes. The layer stackup is signal, power, signal, signal, power, signal from top to bottom. The thickness of the line trace is 0.5-oz. and the separation of the layers is adjusted such that all 7-mil traces have $50\text{-}\Omega$ characteristic impedance. Nelco-13, a low-loss dielectric material with a loss tangent of 0.01 was used (Standard FR-4 has a loss tangent of 0.035). The transmitter and receiver have separate clock sources in order to measure timing margins and plesiochronous clocking. The clock generator produces frequencies from 100-kHz to 3-GHz with typically <1 -ps of RMS jitter. The transmitter outputs are connected back to the receiver inputs of the same chip to ease testing. The power supplies are heavily bypassed, both on chip and off chip. A total of 1.3-nF of bypass capacitor is placed on chip. Separate power supplies are used for the transmitter, the CDR, and the test interface in order to minimize noise coupling and ease power measurement. Off chip, small 1-nF surface mount capacitors are used in the vicinity of the chip since they have higher cut-off frequency. The reason for the higher cut-off frequency for smaller capacitor values is that the parasitic L and R are usually fixed for discrete components with the same geometry and form factor. The capacitor values are gradually increased away from the chip. Big aluminum electrolytic capacitors with a cut-off frequency of a few hundred kHz are placed in the vicinity of the power supply connectors. A separate board with 1-m of serpentine PCB traces was also fabricated to test the performance of equalization.

7.2 Measurement Results

Figure 7.5 shows the differential eye diagram at the transmitter output with a 4-Gb/s 2^{20} -1 PRBS pattern. The waveform is sampled with a 250-MHz trigger clock and thus repeats itself every 4 bits. The plot indicates a phase offset of ~ 15 -ps (The difference between the largest eye and the smallest eye is 30-ps). Figure 7.6 shows the four eyes laid on top of each other. The margin rectangle shown in the middle is 100-mV by 170-ps (more timing margin can be obtained if less swing is required). Figure 7.7 shows the differential eye diagram at the output of a 1-m PCB trace without any equalization. There




Figure 7.5: Differential eye diagram at the transmitter output.




Figure 7.6: Overlap of the bit pattern in Figure 7.5 to show the effective margin. The rectangle shown in the middle is 100-mV by 170-ps.

is no observable eye opening. Even with offset cancellation at the receiver, this unfiltered data pattern is undetectable. Figure 7.8 and Figure 7.9 show the differential eye diagrams at the transmitter output and at 1-m of PCB trace output after equalization is turned on.




Figure 7.7: Differential eye diagram after 1m of PCB trace without equalization.




Figure 7.8: Differential eye diagram at the transmitter output with equalization.




Figure 7.9: Differential eye diagram after 1m of PCB trace with equalization.

The strength of the main tap is kept the same, while the equalization tap is adjusted to be 40% of the main tap. Figure 7.9 shows that the two-tap transmitter pre-emphasis is very effective in opening up the eye. Figure 7.10 again shows the overlap the bit pattern in Figure 7.9. The margin rectangle shown in the middle is 100-mV by 120-ps.

Figure 7.11 shows the quiet supply jitter histogram at the transmitter output with a 1010... pattern. The waveform is again sampled with a 250-MHz clock trigger. The histogram only measures the random jitter and does not include deterministic jitter such as




Figure 7.10: Overlap of the bit pattern in Figure 7.9 to show the effective margin. The margin rectangle shown in the middle is 100-mV by 120-ps.




Figure 7.11: Jitter histogram of the differential transmitter output.

phase offset and ISI. The p-p jitter is 22.2-ps. When 1-MHz 200-mV p-p supply pulses are applied with the noise generator, the p-p jitter increases to 39-ps, corresponding to a supply noise sensitivity of 0.088-ps/mV. The noisy supply jitter histogram is shown in Figure 7.12. The two histogram peaks correspond to the phase positions of the clock at the two alternating power supply levels. This phase variation is mostly due to steady state




Figure 7.12: Jitter histogram of the differential transmitter output with 1-MHz 200-mV p-p pulses superimposed on the supply.




Figure 7.13: Jitter histogram of the receiver sampling clock with automatic phase control turned on (for input data of Figure 7.5).

error of the voltage regulator and the supply sensitivity of the clock buffer, which is outside the multi-phase DLL.

Figure 7.13 shows the jitter histogram of the receiver sampling clock with quiet supply. The transmitter output is connected to the receiver input without significant channel attenuation (eye diagram in Figure 7.5). The receiver sampling clock is brought out to the pin with a PMOS open-drain driver shown in Figure 7.14. The two larger peaks and one smaller peak shown in the histogram are a result of dithering between three phase settings at steady state, a property of the digital bang-bang control. The p-p jitter is 39-ps.




Figure 7.14: PMOS open-drain driver for the on-chip clock signals under observation.

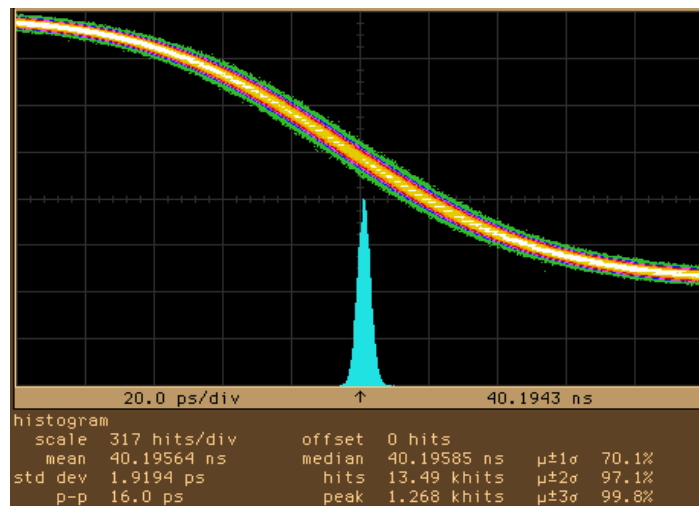


Figure 7.15: Jitter histogram of the receiver sampling clock with automatic phase control turned off (for input data of Figure 7.5).

Figure 7.15 shows the jitter histogram when automatic phase control is turned off. The p-p jitter is 16-ps. It shows that about half of the p-p jitter in Figure 7.13 comes from phase dithering, which can be reduced by decreasing the maximum phase step size. Figure 7.16 shows the jitter histogram with equalized input data after 1-m of PCB trace. The pk-pk jitter is 41-ps. As in Figure 7.13, the recovered clock dithers between three phase positions. Figure 7.17 shows the jitter histogram when 1-MHz 200-mV p-p supply pulses are applied. The p-p jitter increases to 107-ps, corresponding to a supply sensitivity of

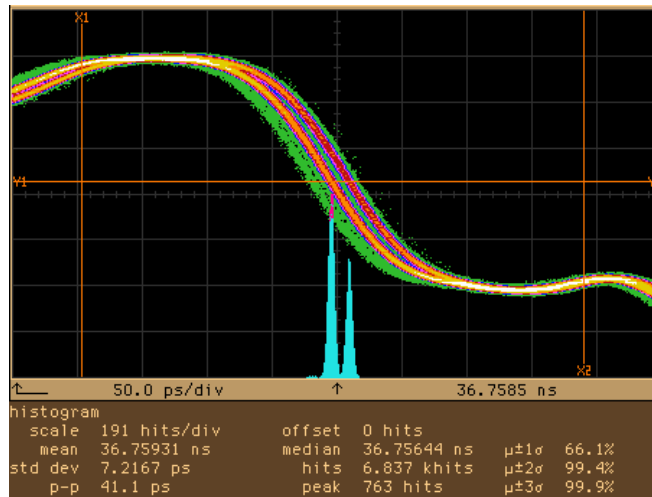


Figure 7.16: Jitter histogram of the receiver sampling clock with automatic phase control turned on (for input data of Figure 7.9).




Figure 7.17: Jitter histogram of the receiver sampling clock with automatic phase control turned on and with 1-MHz 200-mV p-p pulses superimposed on the supply.

0.34-ps/mV. The phase multiplexer, phase interpolator and clock buffers of the clock recovery introduce approximately 70-ps of additional jitter beyond the core DLL jitter.

To test the effectiveness of offset cancellation the receiver margin is measured with and without offset cancellation. The receiver sampling clock is manually swept across a bit time to generate a PASS/FAIL plot. Here “PASS” means successfully receiving a 20-bit PRBS for more than 5 minute. At 4-Gb/s, this corresponds to a BER of $\sim 10^{-12}$. This experiment is repeated for various signal swings. Figure 7.18 shows the single-ended swing versus the sampling clock position. Offset cancellation increases the window width from 170-ps to 200-ps (0.8-UI) and reduces the minimum resolvable single-ended swing from 20-mV to 8-mV. The plot shows that the uncanceled offset of the receivers is around 20-mV, which is approximately the calculated 1σ offset.

With offset calibration, the minimum differential swing required for $< 10^{-14}$ BER is around 20-mV. With noise generator creating 1-MHz 200-mV p-p supply pulses, the minimum differential swing increases to 50-mV. Without offset calibration, the minimum differential swing increases by about 20-mV, which corresponds to the measured offset for

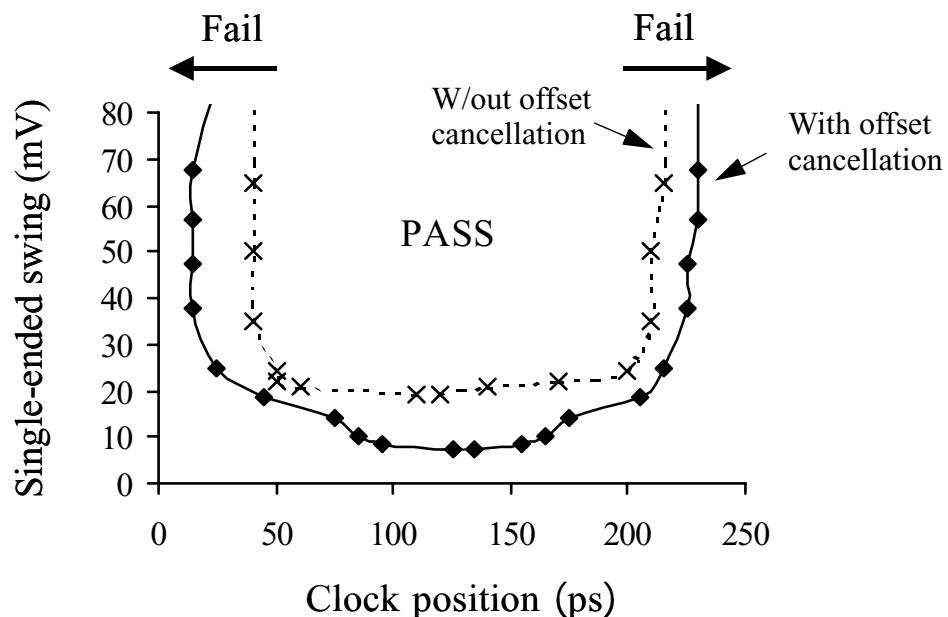


Figure 7.18: Receiver single-ended swing versus clock position window. The PASS region has a BER $< 10^{-12}$.

the test chip. As offset increases, the gain with offset calibration would become more significant.

To verify the phase linearity of the timing vernier, the phase settings are manually swept across a full clock cycle and the delay is measured for each setting. Figure 7.19 shows the phase position versus the phase step over a full clock cycle (72 total steps) for the 90° phase. Figure 7.20 shows the phase step size variation. The measured maximum




Figure 7.19: Phase position versus the phase step of the clock recovery phase adjustment over a full clock cycle.




Figure 7.20: Phase step size. The numbers at each interval indicate the core DLL phase interval. For example, 0-1 indicates the phase interval between 0° and 45° .

step overall is 19-ps with a receive clock frequency of 1-GHz, corresponding to a DNL of 0.37-LSB when the boundary glitch guarding steps are considered and 0.2-LSB when they are ignored.

Whereas the phase steps at the even boundaries, where the even clock phases are switched, show the expected behavior of reduced sizes, the steps at the odd boundary exhibit comparable or even larger sizes than the regular steps. This is due to layout asymmetries, as shown in Figure 7.21. Figure 7.20 shows the *positive edge* step size of the 90° phase, since it is the edge which samples the incoming data. The positive edge of the 90° phase is derived from the *even* and *odd* clock inputs in Figure 7.21. When the phase is at the odd boundaries, both *odd* and *odd_b* inputs change by 90° (250-ps for 1-GHz clock). Since there is significant inter-wire capacitance between *odd* and *even* clocks, the phase of *even* clock also changes slightly, resulting in larger-than-expected step size at the odd boundaries. Since the *odd* clock input is sandwiched by *even* and *even_b*, the capacitive couplings more or less cancel each other, resulting in the expected behavior of reduced step size at the even boundaries. Although the 135° , 225° , 270° , and 315° phases are not visible at the pins, it is expected that they would exhibit the opposite effect (i.e. larger step size at even boundaries) since they are derived from the *even_b* and *odd_b* inputs of the interpolator. Although unexpected, this coupling behavior, which can be removed by clock shielding if desired, turns out to be beneficial since it reduces the average step size by creating extra effective phase steps.




Figure 7.21: Layout of complement phases of interpolators.

Figure 7.22 shows the power consumption of the transmitter (including a 20-bit PRBS generator), the CDR (including a 20-bit PRBS checker), and the total at 2.5-V supply and at the minimum operating supply as a function of clock frequency (the bit rate is 4 times the clock frequency). The transmitter output swing is 100-mV differential (2-mA of current). The maximum speed of the transceiver is 5.32-Gb/s with a 2.5-V supply. At 4-Gb/s, the power consumption of the transceiver is 127-mW at minimum supply and 180-mW at 2.5-V. This plot indicates that significant power saving can be obtained by operating the link at the minimum operating supply. The method by [30], for example, where the whole transceiver is regulated according to the supply voltage of the inverter delay line, can be used to obtain this extra power saving.

Plesiochronous clocking has been applied to the transceiver by using different clock sources for the transmitter and the receiver at slightly different frequencies. The

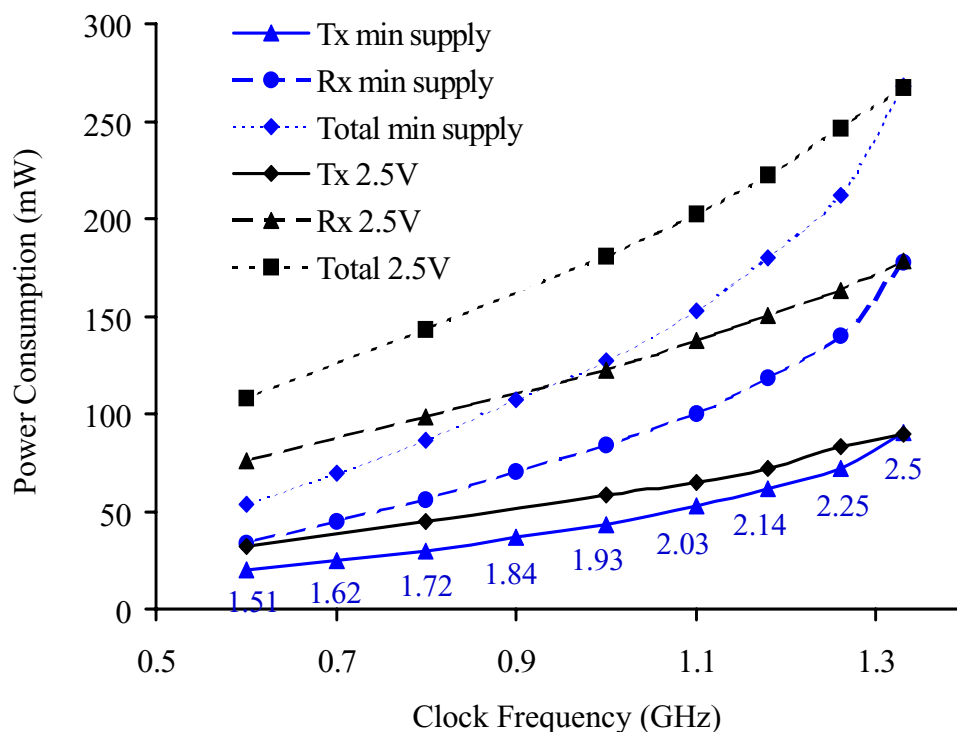


Figure 7.22: Power versus bit rate for the transceiver with minimum operating supply.

maximum frequency tolerance was verified to be $\sim\pm 400$ -ppm. Table 5 summarizes the performance of the transceiver.

Table 5: Test chip performance summary.

Active area	Transmitter: 0.08-mm ² CDR: 0.3-mm ²
Power consumption at 4-Gb/s, 50-mV differential swing, 2-V power supply	127-mW at minimum supply 180-mW at 2.5-V supply
Maximum speed	5.32-Gb/s
Transmitter clock quiet supply jitter	22.2-ps p-p
Transmitter clock supply sensitivity of jitter	0.088-ps/mV
CDR clock quiet supply jitter	38.9-ps p-p
CDR clock supply sensitivity of jitter	0.34-ps/mV
Minimum differential swing for 10^{-14} BER with quiet supply	20-mV
Minimum differential swing for 10^{-14} BER with 1-MHz 200-mV noise pulses superimposed on the supply	50-mV
Frequency tolerance	± 400 -ppm

7.3 Summary

Prototype chips employing the techniques introduced in this work are described and the measurement results detailed in this chapter. In a 0.25- μ m CMOS technology, the active area of the transceiver is 0.31-mm². The transceiver operates with $<10^{-14}$ BER with a 20-mV differential swing and dissipates 127-mW on minimum operating supply and 180-mW on 2.5-V at 4-Gb/s. With a noise generator creating 200-mV of supply pulses to both the transmitter and the CDR, the link operates with $<10^{-14}$ BER with a 50-mV differential swing. The quiet supply jitters of the transmitter and receiver clock are 22.2-ps and 38.9-ps p-p with supply sensitivities of 0.088-ps/mV and 0.34-ps/mV respectively. The maximum phase step of the timing vernier is <20 -ps with a clock frequency of 1-GHz. Offset cancellation increases the receiver timing window from 0.68-UI to 0.8-UI and

reduces the minimum resolvable swing from 20-mV to 8-mV (with BER of 10^{-12}), showing that both the voltage and timing margins are improved. Experiments with 1-m of PCB trace show that a two-tap pre-emphasis filter is very effective in canceling out the ISI of the channel and improving the signal integrity of the link.

These experimental results show that the transceiver design presented in this work enables a high chip throughput by allowing for high speed, low power, low area, and noise immune I/Os to be massively integrated on the same die. 125 of these I/Os would achieve 1-Tb/s of total I/O bandwidth but require only 37-mm² and 22-W on a 2.5-V supply in the 0.25- μ m CMOS technology.

Chapter 8

Conclusion

This work looks at the problem of large-scale I/O integration and describes techniques both on the circuit level and architecture level to improve the power, area, and noise immunity of high speed I/Os. With many innovations introduced in the past years, inter-chip communications over 1-m of PCB trace or 10 – 20-m of coaxial cable at multiple Gb/s speed had just become possible at the onset of this research. A key problem with many of the previous designs, however, is that they consume too much power and area to be cost effective in applications requiring hundreds of high-speed I/Os on the same chip.

The first block we examined was the transmitter. The key problem here was how to serialize parallel data on-chip into high speed serial data off-chip at our target speed (4-Gb/s in 0.25- μm CMOS) while minimizing area and power. In addition, one of the drawbacks of the previous design is the significant capacitive loading at the transmitter output, degrading the quality of transmitter termination. We use a low-swing input-multiplexed architecture to mitigate these shortcomings while achieving our speed requirement.

For channels that have significant frequency-dependent attenuation, data need to be filtered, usually with a finite-impulse-response (FIR) filter, to be received reliably. The complexity of the filter often needs to be constrained to minimize power and area. In this thesis, we developed a mathematical tool which allows a quick quantization of bit error rate variation with filter complexity. The analysis was used to validate our use of a simple two-tap filter on a backplane channel.

One of the major drawbacks of the previous receiver designs is that they operate with uncanceled offset. We introduced a capacitive offset trimming method which reduces the receiver offset to $< 8\text{-mV}$ while degrading the aperture time of the receiver by only 6%. This scheme improves both the voltage margin and the timing margin and saves power and area by requiring less swing and smaller receivers.

To reduce the power of the timing circuits while maintaining a good immunity to the power supply noise, we use a supply regulated CMOS inverter delay line for the multi-phase generation. Compared to a source-coupled delay element, this design saves approximately 30% of the power for 4 phases and 60% of the power for 8 phases. It also reduces the supply sensitivity of the CMOS delay element to about -0.09, which is half that of a source-coupled delay element.

For the clock recovery, we adopted the Sidiropoulos dual-loop architecture. The phase multiplexer and phase interpolator are implemented with a current-mirror topology to obtain a high bandwidth and a good interpolation linearity. This topology helps the overall timing budget by reducing the receiver clock jitter and dithering.

A 4-Gb/s I/O which incorporates the above techniques has been built in a 0.25- μm CMOS technology and verified in the lab. It consumes 180-mW of power on a 2.5-V supply and occupies 0.3-mm² of area. It is also able to withstand a 200-mV supply noise generated on-chip with $< 10^{-14}$ BER with only 50-mV of differential swing. This work shows that it is possible to achieve a bandwidth on the order of Tb/s on a single chip with a reasonable amount of power and die area in the current CMOS technology. For example, to achieve an aggregate 1-Tb/s I/O bandwidth requires 125 copies of our I/O, 22-W of power, and 37-mm² of die area.

8.1 Future Work

As process technology scales, transistor mismatch, which introduces both voltage and phase offsets, becomes more significant. In this work, we use a capacitive trimming method to cancel out the receiver input voltage offset. As shown by the experimental results of this work, phase offset is also becoming a significant limiting factor. Perhaps a similar approach, in which the phase offset is measured and corrected statically at startup [31], can be used to cancel out the phase offset in the multi-phase clocking scheme.

Another critical area of high speed I/O design is channel equalization. As the signaling rate increases, not only do the skin effect and the dielectric loss become more significant, but reflections due to connectors and impedance mismatches also worsen. More complicated filter designs which cancel out all of the above effects and better connector and material designs need to happen simultaneously to push the bit rate higher.

Frequency synthesis for high-speed I/Os is commonly done with a ring-oscillator PLL. This type of design is very sensitive to power supply noise due to phase noise accumulation, as pointed out in Chapter 2. Although better process technology helps reduce the jitter, a big part of it is fundamental to the architecture and the circuit topology. As the bit time continues to shrink, innovations both on the architectural level and on the circuit level must happen to overcome the clock jitter limitation. One approach which shows a great promise is LC oscillator based frequency synthesis [32]. The limitations here are the large area required for the on-chip spiral inductor (to achieve a good quality factor) and the small tuning range.

Finally, the clock recovery circuits implemented in this work still require a significant amount of power and area compared to other components of the transceiver. This is especially true for the interpolation circuits (with multi-phase generation). A more power and area efficient scheme, such as the one based on the idea of coupled oscillators [33], might be better suited for highly integrated applications.

Bibliography

- [1] M.-J. E. Lee, W. J. Dally, and P. Chiang, "A 90-mW 4-Gb/s equalized I/O circuit with input offset cancellation," in *ISSCC Dig. Tech. Papers*, pp. 252-253, Feb 2000.
- [2] M.-J. E. Lee, W. J. Dally, and P. Chiang, "Low-power area efficient high speed I/O circuit techniques," *IEEE J. Solid-State Circuits*, vol. 35, pp. 1591-1599, Nov 2000.
- [3] M.-J. E. Lee *et al.*, "An 84-mW 4-Gb/s clock and data recovery circuit for serial link applications," in *Proc. IEEE Symposium on VLSI Circuits*, June 2001.
- [4] W. J. Dally *et al.*, "High-performance electrical signaling," MPPOI98, 1998.
- [5] R. Marbot *et al.*, "Integration of multiple bidirectional point-to-point serial links in the gigabit per second range," in *Proc. Hot Interconnect*, Aug 1993.
- [6] W. J. Dally and J. Poulton, "Transmitter equalization for 4Gb/s signaling," in *Proc. Hot Interconnect*, pp. 29-39, Aug 1996.
- [7] J. Poulton, W. J. Dally, and S. Tell, "A tracking clock recovery receiver for 4 Gb/s signaling," in *Proc. Hot Interconnects*, pp. August 21-23, 1997, Palo Alto, CA

- [8] C. K. K. Yang, R. Farjad-Rad, and M. Horowitz, "A 0.5 μ m CMOS 4Gbps transceiver with data recovery using oversampling," *IEEE J. Solid-State Circuits*, vol. 33, pp. 713-722, May 1998.
- [9] R. Farjad-Rad, C. K. K. Yang, M. A. Horowitz, and T. H. Lee, "A 0.4 μ m CMOS 10-Gb/s 4-PAM pre-emphasis serial link transmitter," *IEEE J. Solid-State Circuits*, vol. 34, pp. 580-585, May 1999.
- [10] K.-Y. K. Chang *et al.*, "A 2Gb/s/pin asymmetric serial link," in *Proc. IEEE Symposium on VLSI Circuits*, pp. 216-217, June 1998.
- [11] J. G. Maneatis, "Low jitter process-independent DLL and PLL based on self-biased techniques," *IEEE J. Solid-State Circuits*, vol. 31, pp. 1723-1732, Nov 1996.
- [12] S. Sidiropoulos and M. Horowitz, "A Semidigital Dual Delay-Locked Loop," *IEEE J. Solid-State Circuits*, vol. 32, pp. 1683-1692, Nov 1997.
- [13] R. Gu, *et. al.*, "A 0.5-3.5Gb/s low-power low-jitter serial data CMOS transceiver," in *ISSCC Dig. Tech. Papers*, pp. 352-353, Feb. 1999.
- [14] S. Sidiropoulos and M. Horowitz, "A 700-Mb/s/pin CMOS signaling interface using current integrating receivers," *IEEE J. Solid-State Circuits*, vol. 32, pp. 681-690, May 1997.
- [15] J. Montanaro, *et al.*, "A 160MHz, 32b, 0.5W CMOS RISC microprocessor," *IEEE J. Solid-State Circuits*, vol. 31, pp. 1703-1714, Nov 1996.
- [16] R. Farjad-Rad *et al.*, "A 0.3- μ m CMOS 8-Gb/s 4-PAM serial link transceiver," *IEEE J. Solid-State Circuits*, vol. 35, pp. 757-764, May 2000.
- [17] E. Yeung and M. A. Horowitz, "A 2.4 Gb/s/pin simultaneous bidirectional parallel link with per-pin skew compensation," *IEEE J. Solid-State Circuits*, vol. 35, pp. 1619-1628, Nov 2000.
- [18] B. Kim, T. Weigandt, and P. Gray, "PLL/DLL system noise analysis for low jitter clock synthesizer design," *Proc. International Symposium on Circuits and Systems*, vol. 4, pp. 31-38, 1994.

- [19] A. Fiedler *et al.*, "A 1.0625 Gbps transceiver with 2X oversampling and transmit pre-emphasis," in *ISSCC Dig. Tech. Papers*, pp. 238-239, Feb 1997.
- [20] <http://www.velio.com>
- [21] W. J. Dally and J. Poulton, *Digital Systems Engineering*, Cambridge University Press, 1998.
- [22] *Star-Hspice User's Manual*, Avant! Corporation, 1998.
- [23] L. R. Dennison, W. S. Lee, and W. J. Dally, "High performance bi-directional signaling in VLSI systems," in *Proc. Symposium on Research on Integrated Systems*, Jan. 1993.
- [24] A. Chandrakasan, W. Bowhill, and F. Fox, *Design of high performance microprocessor circuits*, Chap 19, IEEE Press, 2000.
- [25] M. Horowitz, C. K. K. Yang, and S. Sidiropoulos, "High-speed electrical signaling: overview and limitations," *IEEE Micro*, vol. 18, pp. 12-24, Jan-Feb 1998.
- [26] W. Ellersick *et al.*, "GAD: a 12-GS/s CMOS 4-bit A/D converter for an equalized multi-level link," in *Proc. IEEE VLSI Circuits Symp.*, pp. 49-52, June 1999.
- [27] M. Pelgrom, A. Duinmaijer, and A. Welbers, "Matching Properties of MOS Transistors," *IEEE J. Solid-State Circuits*, vol. 24, pp. 1433-1440, Oct 1989.
- [28] V. von Kaenel *et al.*, "A 320 MHz, 1.5 mW @ 1.35 V CMOS PLL for microprocessor clock generation," *IEEE J. Solid-State Circuits*, vol. 31, pp. 1715-1722, Nov 1996.
- [29] Y. Moon *et al.*, "An all-analog multiphase delay-locked loop using a replica delay line for wide-range operation and low-jitter performance," *IEEE J. Solid Circuits*, vol. 35, pp. 377-384, Mar. 2000.
- [30] G.-Y. Wei *et al.*, "A variable-frequency parallel I/O interface with adaptive power-supply regulation," *IEEE J. Solid State Circuits*, vol. 35, pp. 1600-1610, Nov. 2000.

- [31] D. Weinlader *et al.*, "An eight channel 36GSample/s CMOS timing analyzer," in *ISSCC Dig. Tech. Papers*, pp. 170-171, Feb 2000.
- [32] Thomas H. Lee, *The design of CMOS radio-frequency integrated circuits*, Cambridge University Press, 1998.
- [33] J. G. Maneatis and M. A. Horowitz, "Precise delay generation using coupled oscillators," *IEEE J. Solid-State Circuits*, vol. 28, pp. 1273-1282, Dec 1993.
- [34] W. J. Dally, "Digital transmitter with equalization," US Patent 6266379.
- [35] N. Kushiyama *et al.*, "A 500-megabyte/s data rate 4.5M DRAM," *IEEE J. Solid-State Circuits*, vol. 28, pp. 490-498, April 1993.