

Computer Systems Laboratory

Departments of Electrical Engineering & Computer Science
Stanford University



**Data Converters
for
High Speed CMOS Links**

A PhD Dissertation

William F. Ellersick

August 2001

Copyright © 2001

by

William F. Ellersick

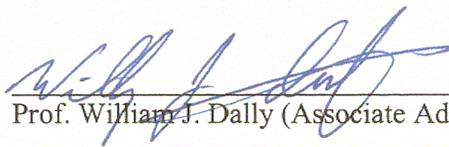
All Rights Reserved

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.



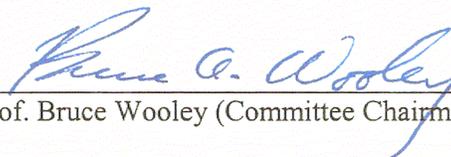
Prof. Mark A. Horowitz (Principal Advisor)

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.



Prof. William J. Dally (Associate 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.



Prof. Bruce Wooley (Committee Chairman)

Approved for the University Committee on Graduate Studies:

Abstract

The long links that interconnect networking and computing systems and boards need high throughput to avoid expensive, massively parallel connections. However, long wires suffer signal losses that increase with frequency. Digital communication techniques can compensate for these losses, but require Analog-to-Digital and Digital-to-Analog converters (ADCs and DACs). To understand the application of these techniques to high speed links, an 8 GSample/sec CMOS transceiver chip is designed to explore the limits of high speed data converter performance.

The transceiver chip provides a high bandwidth signal path and precision clocks, despite the large parasitic capacitances and transistor matching errors of CMOS technology. Small, high bandwidth sample-and-hold amplifiers are used in the ADC, and the resulting large mismatch errors are corrected by small DACs in each ADC comparator. Other circuit and signal degradations such as transmitter nonlinearity, clock coupling, and static phase errors are also digitally corrected. Time interleaving is used to achieve 8GSa/s, and the effects of the increased data converter capacitances are reduced with bond-wire inductors. These inductors distribute the lumped parasitic capacitances at the transceiver input and output to approximate distributed 50Ω transmission lines, reducing attenuation by 10 dB at 4 GHz. The data converters allow the transceiver to use digital equalization to compensate for the 3 GHz transceiver bandwidth to allow 8GSa/s multi-level data transmission. Measured results indicate that digitally corrected data converters will allow digital communication techniques to be applied to high speed CMOS links.

Acknowledgments

I would like to acknowledge the guidance and expertise of Prof. Mark Horowitz and Prof. Bill Dally. Without their knowledge and the foundation of the research of their students, this work would not have been possible.

Ken Yang was always generous with his time and ideas, and was like a third advisor to me. He designed the transmitter with his student Siamak to complete the link transceiver. Ken Chang designed the PLL, delaying his own graduation to help me graduate. Vladimir Stojanovic tirelessly assisted with tapeout, characterization and extensive development of measurement techniques and equalization algorithms, contributing greatly to the thoroughness of this dissertation.

Azita Emami and Dean Liu both worked hard to help make this project successful. Our bonding technician, Pauline Prather, did amazing work to support experimentation with inductors. I'd like to thank Stefanos Sidiropoulos for his insight and keen analysis. While I've listed a few here, everyone in Mark's group was free with their time and ideas and made my PhD fun and easier, as were the students in Prof. Wooley's and Prof. Lee's groups. At Analog Devices, Kimo Tam strongly supported my efforts to complete this dissertation, and Jesse Bankman thoroughly reviewed the manuscript.

I'd like to thank my wife Chris for her support, and Jimmy, John and Stevie for help with the layout of my chips and boards.

This thesis is dedicated to my editor father, Fred, who unfailingly covered my compositions with arcane editing notations, always in his trademark purple ink. Thanks for never accepting less than my best.

L.I.S.

(Let It Stand)

Table of Contents

Abstract.....	v
Acknowledgments	vii
Table of Contents	ix
List of Tables	xi
List of Figures.....	xiii
Chapter 1: Introduction	1
1.1 Organization.....	3
Chapter 2: Background.....	5
2.1 Transceiver Overview	6
2.2 Clock Generation	10
2.3 Transmitter Design	15
2.4 Receiver Design.....	17
2.5 Summary	19
Chapter 3: Receiver Design	21
3.1 Sampler Design Analysis.....	22
3.2 Comparator Architecture	26
3.3 Sample/Hold Amplifier.....	27
3.4 Second Stage Latch.....	31
3.4.1 Offset Correction DACs	32
3.4.2 Offset DAC Bias Generator.....	34
3.4.3 Latch Output Sampling.....	34
3.5 Output Latch	35
3.6 High Speed Receiver Logic	36
3.7 Reference Voltage Generation.....	37
3.8 Simulated Comparator Performance.....	38

3.9	Comparator Array Layout.....	39
3.10	Summary	41
Chapter 4: Transceiver Chip Design		43
4.1	Transceiver Architecture.....	43
4.2	Transmit DAC Design	44
4.3	Inductors to Distribute Parasitic Capacitance	47
4.4	Clock Generation	51
4.4.1	Digital Phase Detector	53
4.5	Synthesized Logic and Memories.....	55
4.6	Transceiver Capabilities	57
4.6.1	Interference Correction	57
4.6.2	Multi-level Modulation.....	58
4.6.3	Linear Equalization.....	60
4.6.4	Decision Feedback Equalization.....	62
4.7	Summary	63
Chapter 5: Experimental Results		65
5.1	Inductors to Distribute Parasitic Capacitances	66
5.2	Timing Noise	68
5.3	ADC Measurements.....	69
5.4	DAC Measurements.....	75
5.5	Equalized Transmitter Results	79
5.6	Full Transceiver Results	81
5.7	Summary	84
Chapter 6: Conclusion.....		85
Bibliography		89
	Serial Links	89
	A/D Converters	90
	Circuits.....	92
	Clock Generation	93
	Communications	94

List of Tables

Table 3.1:	Simulated Comparator Performance.....	38
Table 5.1:	ADC Performance Summary	75
Table 5.2:	DAC Performance Summary	79
Table 5.3:	Transceiver Performance Summary.....	83

List of Figures

Figure 1.1:	Pulse Distortion from Long Wire	1
Figure 1.2:	Equalized Pulse Train	2
Figure 1.3:	Transceiver Based on DAC and ADC	3
Figure 2.1:	Typical Transceiver	6
Figure 2.2:	Time-Interleaved Receiver	7
Figure 2.3:	Timing Errors.....	8
Figure 2.4:	LENOB, ENOB vs $t_{\text{error}}/T_{\text{symbol}}$	9
Figure 2.5:	VCO Architecture	11
Figure 2.6:	Delay Element.....	11
Figure 2.7:	PLL Phase Noise Spectrum	11
Figure 2.8:	Dual Loop Clock Generation	13
Figure 2.9:	Clock Interpolator	14
Figure 2.10:	Open Drain Transmitter	15
Figure 2.11:	Grounded Source Tx.....	16
Figure 2.12:	Passgate/Latch Receiver	17
Figure 2.13:	Flash ADC Architecture	17
Figure 2.14:	Offset Cancellation	18
Figure 3.1:	4-bit Flash ADC.....	21
Figure 3.2:	Time-interleaved ADC	21
Figure 3.3:	Sampled Latch	22
Figure 3.4:	Latch Impulse Response	22
Figure 3.5:	StrongArm Latch	23
Figure 3.6:	Pass Gate Sampler	23
Figure 3.7:	Wideband Pre-Amp	25
Figure 3.8:	Comparator	26
Figure 3.9:	Sample/Hold Amplifier.....	27
Figure 3.10:	Differential Load M2	28
Figure 3.11:	Sampling	29
Figure 3.12:	Sampling Waveforms	30
Figure 3.13:	Matched Clock Generation	31
Figure 3.14:	Latch with Offset Correction	32

Figure 3.15:	Offset Correction DACs	32
Figure 3.16:	Offset DAC Bias	34
Figure 3.17:	Output Latch Waveforms.....	34
Figure 3.18:	Output Latch Schematic.....	35
Figure 3.19:	High Speed ADC Circuitry.....	36
Figure 3.20:	Reference Bias Circuit	37
Figure 3.21:	Comparator Array Layout.....	39
Figure 3.22:	Interleaved Comparator Bank Layout.....	40
Figure 3.23:	Shield Troughs	40
Figure 4.1:	Transceiver Chip Architecture	44
Figure 4.2:	Time Interleaved Transmitter	45
Figure 4.3:	Interleaved DAC Design.....	46
Figure 4.4:	Lumped 50 Ohm Transmission Line	47
Figure 4.5:	DAC Output Inductors.....	47
Figure 4.6:	a) Simulated DAC Output Bandwidth..b) Simulated TDR of Output	48
Figure 4.7:	Transceiver Block Diagram with Inductors.....	49
Figure 4.8:	a) Simulated ADC Input Bandwidth..b) Simulated TDR of ADC Input	50
Figure 4.9:	Enclosed Dual Loop Clock Generation	51
Figure 4.10:	Transition Types	54
Figure 4.11:	Synthesized Receive Logic	56
Figure 4.12:	Transmit Logic.....	56
Figure 4.13:	Transmit Sequence Processing	58
Figure 4.14:	2-PAM and 4-PAM Modulation on 18m RG59 cable (6 tap equalizers) ..	58
Figure 4.15:	2-PAM and 4-PAM Wire Attenuation vs. Frequency	59
Figure 4.16:	Frequency Response of Channel and Equalizer	60
Figure 4.17:	5-tap FIR filter	60
Figure 4.18:	Receive Equalization	61
Figure 4.19:	Transmit Equalization.....	61
Figure 4.20:	Decision Feedback Equalizer.....	62
Figure 5.1:	Transceiver Die Photo	65
Figure 5.2:	Bond Wire Inductors.....	66
Figure 5.3:	ADC Input TDR Results.....	67
Figure 5.4:	TDR of DAC Output.....	67
Figure 5.5:	PLL Phase Noise vs. Noise Frequency	68
Figure 5.6:	ADC Input Offset Histogram.....	69
Figure 5.7:	Offset Correction DAC Step Sizes	69
Figure 5.8:	ADC INL After Calibration	70
Figure 5.9:	ADC Error Distribution	71
Figure 5.10:	ADC Frequency Response.....	72
Figure 5.11:	ADC SNDR	72

Figure 5.12:	ADC LSNDR	74
Figure 5.13:	Uncalibrated Pulses.....	76
Figure 5.14:	Calibrated Pulses.....	76
Figure 5.15:	Transceiver Clock Coupling	77
Figure 5.16:	DAC INL and DNL vs. Code	77
Figure 5.17:	DAC Noise vs. Code.....	78
Figure 5.18:	Raw DAC Pulse	80
Figure 5.19:	Equalized DAC Pulse	80
Figure 5.20:	8 GSa/s Transmitter Pulse and Frequency Response.....	81
Figure 5.21:	8 GSa/s Transceiver Pulse and Frequency Response	82
Figure 5.22:	Equalized Binary and 4-Level Transceiver Schmoos	82
Figure 5.23:	Transceiver 10-10 BER Schmoos	83

Chapter 1

Introduction

High speed links improve the size, cost and performance of systems by reducing the number of wires needed. Link data rates can be increased either by encoding more bits in each symbol, or by running at a faster symbol rate. While running faster is generally the simplest approach, implementation is complicated by the low pass filters inherent to wires and semiconductors. Wires often form the dominant low pass filters in long links, because the conductor resistance and the dielectric (insulation) loss both increase with frequency. Parasitic capacitances and inductances also tend to form low pass filters (otherwise they might be touted as “features” instead of parasitics).

Low pass filters distort signal waveforms, attenuating high frequency components which rounds fast edges and causing interference, as shown in Figure 1.1. The nearly ideal input pulse on the left is rounded and spread out after transmission on a long wire. The

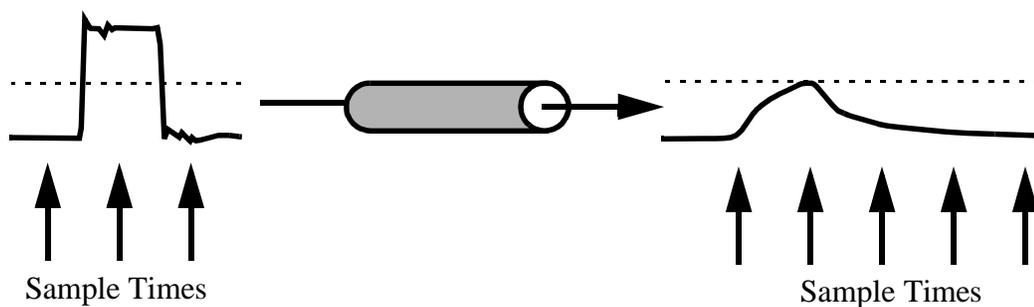


Figure 1.1: Pulse Distortion from Long Wire

input pulse has a strong peak, and is almost zero at adjacent sample times, while the distorted pulse has a smaller peak, and affects the received value at several sample times.

Thus, a train of pulses carrying digital data will interfere with each other, causing inter-symbol interference (ISI). Unequalized binary data transmission fails when pulses fail to reach half-swing at the receiver, since positive and negative pulses (due to 00100 and 11011 data patterns, respectively) cannot be received by comparing to a fixed (mid-swing) voltage.

Fortunately, wire losses and many parasitics are linear, and thus can be corrected with adjustable high-pass filters called linear equalizers [10]. Figure 1.2 shows a pulse train transmitted on a long wire, with the received signal on the upper right distorted beyond reception by a simple comparator. The step response of a transmit equalizer can be

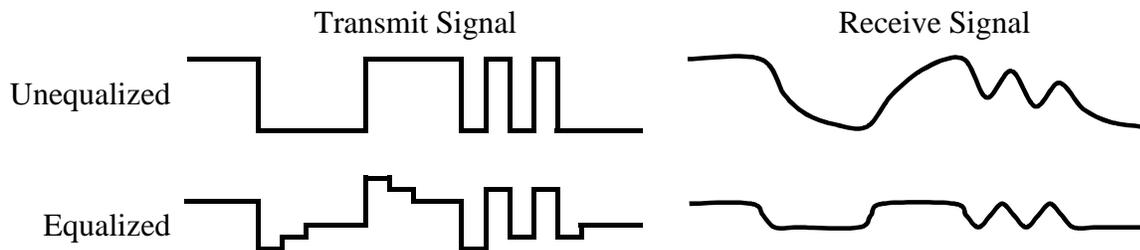


Figure 1.2: Equalized Pulse Train

seen in the first few symbol times, emphasizing the high frequency edges, then decaying to a smaller value. Because link transmitters are usually limited in voltage swing, a transmit equalizer *attenuates* low frequency signal components to match the high frequency wire and parasitic losses. The equalized received signal reaches the same high and low values regardless of the adjacent symbols, but has been significantly attenuated. Linear equalization becomes increasingly difficult as the parasitic attenuation increases. For example, when pulses are attenuated by a factor of 3, equalization errors must be less than 1/6th of the signal swing to maintain a 50% eye opening, and more complex digital communication techniques become attractive.

Advanced techniques such as Decision Feedback Equalization (DFE) or Tomlinson precoding, multi-level modulation and adaptive interference cancellation [59][60][62] can increase performance on long links with severe losses. However, these techniques all require extra resolution in the transmitters or receivers. Thus, their effectiveness depends on the performance of very high speed ADCs and DACs. This

this thesis investigates the applicability of digital communication techniques to high speed CMOS links by exploring the limits of ADC and DAC performance in a prototype transceiver shown in Figure 1.3.

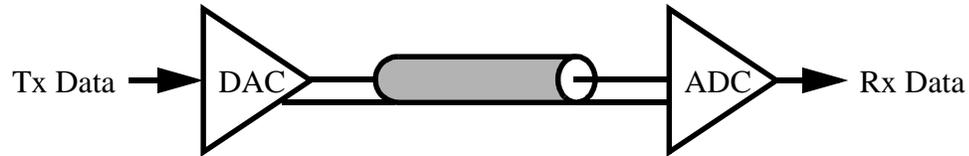


Figure 1.3: Transceiver Based on DAC and ADC

To increase data rates as well as handle more wire losses, the ADC and DAC need to realize the high sample rates and bandwidths of binary transceivers while achieving multiple bits of resolution. The data converters need Nyquist bandwidth (half the sampling frequency) to prevent the circuits from limiting the link performance. Roughly 4 bits of resolution are required, with 2 bits to support multi-level signalling, and 2 more bits for additional signal processing. Extra resolution is desirable to correct for systematic nonlinearities and interference, and to reduce quantization errors in equalization.

1.1 Organization

The following chapters examine the performance limits of ADCs and DACs that run at binary transceiver rates to understand the applicability of digital communication techniques to high speed links. Because the performance of the communication techniques is well understood, the focus of this dissertation is on the implementation and performance of the receiver, transmitter and clock generation circuits.

Chapter 2 develops a framework for comparing link circuit performance, and uses that framework to derive process independent targets for the transceiver circuit blocks. The major circuit blocks in a transceiver are then surveyed to identify the development needed to design accurate, high speed transceivers that leverage previous research on high speed binary transceivers and high accuracy ADCs.

Chapter 3 focuses on the design of the receiver, beginning with an analysis of several types of sampling circuits. The remainder of the chapter presents the circuit design of the comparators that comprise the ADC, and the bias circuits and matched clock generators that support the comparators.

Chapter 4 presents the architecture of the full transceiver and then examines the design of the transmitter, clock generation circuitry, inductive network, logic and memories. Communication techniques that the transceiver can support are also discussed, setting up the experimental results in Chapter 5.

Chapter 5 presents experimental data from a transceiver chip fabricated in 0.25 μm CMOS. TDR results are used to verify inductance values, and timing noise is characterized, to allow frequency performance results to be correctly interpreted. Results are then presented on the calibration, performance over frequency, and accuracy of the receiver and transmitter. Equalization results are shown first on the transmitter because they are easy to understand and measure, and then equalized link results are presented for the full transceiver. Chapter 6 summarizes the results and concludes the dissertation.

Chapter 2

Background

The design of ADCs and DACs at the high sample rates of link transceivers builds on previous binary link and high speed ADC work. While link transceivers achieve the high sample rates and bandwidth needed, the circuits used lack the accuracy to be directly used in ADCs and DACs. Conversely, high speed ADCs routinely exceed 8 bits of resolution, but do not operate at link speeds.

To provide a framework for comparing circuit performance, this chapter begins with a review of a typical transceiver architecture and the fanout-of-4 gate delay metric that is used to normalize circuit speed across CMOS processes [11][44]. This discussion is followed by an analysis of link performance that derives the number of bits a link can transport from the peak-to-peak signal to noise ratio. Process independent targets are derived for timing noise, sample rate and bandwidth using the analytical framework.

The major circuit blocks in a transceiver are surveyed, starting with a discussion of clock generation and timing recovery circuits that shows that existing circuits have adequate performance to allow 3-bit transceivers to be built. State-of-the-art transmitter, receiver and ADC designs are then reviewed, with an eye toward the development needed to extend transceiver designs to multiple bits of resolution to support signal processing, while maintaining high sample rate and bandwidth.

2.1 Transceiver Overview

A typical high speed transceiver (shown in Figure 2.1) includes a transmitter to drive data onto a terminated transmission line, and a receiver to sample and amplify the incoming signal. Clock generation circuitry provides stable clocks for the transmitter, and timing recovery circuitry provides receiver clocks that are synchronized to the input signal.

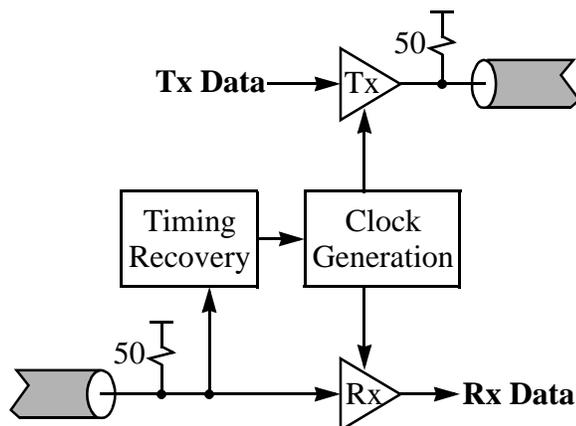


Figure 2.1: Typical Transceiver

This dissertation explores the use of digital communication techniques in high speed links by replacing the binary receiver and transmitter with an ADC and a DAC. In addition to supporting signal processing, the ADC and DAC allow the commonly used 2-level Pulse Amplitude Modulation (2-PAM) to be extended to multi-level PAM. While phase modulation and other more complex schemes are used in some applications, amplitude modulation is simple to implement, bandwidth efficient¹ and tolerant of phase noise, and thus is the focus of this work.

Comparing transceiver designs is difficult because circuit speeds depend strongly on the technology used. Thus, a metric called a fanout-of-4 delay (FO4) is used to normalize performance over CMOS processes. A FO4 is the delay of one inverter in a chain with a fanout of 4, such that each inverter drives a load four times as large as its input capacitance. The performance of a wide range of CMOS logic and transceiver circuits varies over process, temperature and supply by approximately the same factor as the FO4 delay. Thus, performance measured in FO4 delays is used both to compare results in disparate processes, and to extrapolate them to future processes [11][44].

1. PAM has the same bandwidth efficiency as carrier-based Quadrature Amplitude Modulation (QAM). Because modulation on a carrier creates double sidebands, it doubles the bandwidth used but supports two data streams on $\sin(\omega t)$ and $\cos(\omega t)$ carriers in quadrature. QAM is used by telephone and xDSL modems because they cannot use the low frequency portion of the spectrum. Since CMOS links have use of the entire spectrum of the wire, there is no apparent benefit to carrier-based modulation unless up- and down-converters can be implemented at higher frequencies than baseband receivers.

Using the FO4 metric, transceiver designs can be sorted into speed classes, based on the number of FO4 delays per bit time. Most high speed transceivers run at about 4 FO4 delays per bit, and transmit data on both the rising and falling edges of an internal clock [1][6][15]. Time interleaving allows bit rates faster than the internal clock rate [16][18], by

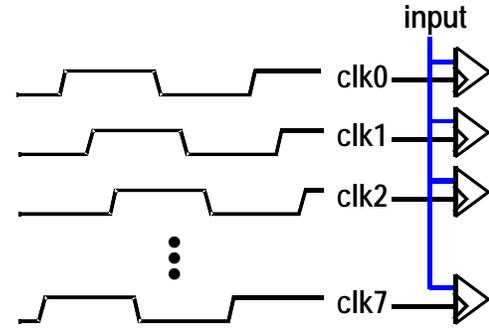


Figure 2.2: Time-Interleaved Receiver

using parallel transmitters and receivers which are clocked with phase shifted versions of the internal clock. An interleaving ratio of 8 is shown in Figure 2.2, and has been used to achieve bit times of approximately 1 FO4 [3][10][52]. Time interleaving is needed to achieve such short bit times, since CMOS clocks with periods less than about 8 FO4 delays are not practical to generate or use. Generation of clocks with shorter periods is difficult, because the clocks must be buffered by low-fanout inverter chains that increase power dissipation. Clocks with periods shorter than 8 FO4 are impractical to use because they result in excessive pipelining of logic, since there are roughly 3 FO4 delays of timing overhead in static flip-flops [11].

To match the speed of binary transceivers, the sample time goal for the transceiver in this work is 1 FO4. To prevent the circuits from limiting performance, the ADC and the DAC should have Nyquist bandwidth $B = \frac{1}{2FO4}$ (half the sample rate). The time constant $\tau = \frac{1}{2\pi B}$ is easier to compute and express in terms of FO4 gate delays, and is used as a process independent goal (with $1/\pi$ approximated by 0.3):

$$\tau = 0.3 \text{ FO4} \quad [2.1]$$

Given the FO4 delay of 125 ps in the 0.25 μ m CMOS process used in this work (worst case process, temperature and voltage), the goal is 4GHz bandwidth at an 8GSa/s sample rate.

To analyze the performance of a link, the noise and distortion from a variety of sources needs to be taken into account (see Chapter 6 in [61] for a thorough, accessible discussion of noise in CMOS systems). Since a bit error rate of less than 10^{-9} is typically required, the worst case total noise with a probability greater than 10^{-9} sets the performance limits of a link. This work takes the conservative approach of using

worst-case noise analysis for most of the dominant noise sources, such as supply noise, transistor mismatches and nonlinearities, and interference from other clocks and signals. This worst case noise is added to the worst-case statistical noise (probability greater than 10^{-9}). Thermal noise causes statistical voltage noise in high bandwidth CMOS circuits [45][46][47], and causes statistical phase noise in clock generation circuits.

To understand the performance of a link with timing errors, phase noise is mapped to voltage noise with a simple model. This model approximates the received signal in a link with a sine wave at half the symbol rate, as shown in Figure 2.3. The sine wave is quite similar to the signal that is received for a 101010... data pattern (and to the transitions in other data patterns), since higher frequency edge components are attenuated by the parasitic low-pass filters in wires, packages and circuits. Since a

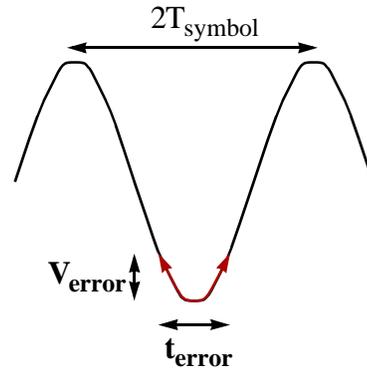


Figure 2.3: Timing Errors

synchronized link nominally samples the received data at the peaks¹, the p-p voltage error caused by a p-p timing error on a sine wave with a 1V swing is given by:

$$V_{\text{noise,p-p}} = \frac{1}{2} - \frac{1}{2} \cos\left(2\pi \frac{t_{\text{error}}/2}{2T_{\text{symbol}}}\right) \cong \frac{1}{4} \left(\frac{\pi t_{\text{error}}}{2T_{\text{symbol}}}\right)^2 \quad \text{for } \frac{t_{\text{error}}}{T_{\text{symbol}}} < 0.5 \quad [2.2]$$

Note that the synchronous $V_{\text{noise,p-p}}$ is much smaller than peak asynchronous errors at the steepest slope of the signal, which vary with t_{error} instead of $t_{\text{error}}^2/4$ (for small t_{error}).

The ratio of the received signal to the total worst case noise (with probability greater than 10^{-9}) determines the robustness of a link and the number of bits that it can transport (with a bit error rate $< 10^{-9}$). A link can receive N bits if the p-p signal is 2^N times greater than the p-p noise, since 2^N signal levels with noise are needed. Thus the Link Effective Number of Bits is defined as a function of the p-p signal to p-p noise voltage ratio, $V_{\text{SNR,p-p}} = \frac{V_{\text{signal,pp}}}{V_{\text{noise,pp}}}$:

1. This analysis assumes that the link uses a linear equalizer. More sophisticated signal processing techniques such as Decision Feedback Equalization (see Section 4.6) can result in non-zero slope at the nominal sampling time and thus have greater sensitivity to timing errors.

$$\text{LENOB} = \log_2(V_{\text{SNR,p-p}}) \quad [2.3]$$

The LENOB is similar to the Effective Number of Bits (ENOB) used to measure ADC and DAC resolution, defined in Equation 2.4 as a function of the Signal to Noise and Distortion Ratio (SNDR) on an asynchronous full scale sine wave [24][26]. The primary difference is that the LENOB uses synchronous p-p noise samples at the peaks of the signal, while the ENOB uses asynchronous noise power. A minor difference is the factor of 1.22 in Equation 2.4, which represents the ratio of the RMS value of ideal quantization noise to the RMS value of a sine wave. The LENOB more accurately measures the larger effect of random voltage noise on a link, and the smaller effect of phase noise on a synchronized link, and allows dissimilar noise sources to be correctly combined, as discussed below.

$$\text{ENOB} = (\text{SNR}_{\text{RMS}} - 1.76) / 6.02 = \log_2(V_{\text{SNR,RMS}}/1.22) \quad [2.4]$$

Both the LENOB and ENOB are plotted in Figure 2.4 vs. the p-p timing error, t_{error} , for an otherwise ideal link and ADC (or DAC). For the ENOB curve, $t_{\text{error,RMS}} = t_{\text{error,pp}}/8.5$, assuming normally distributed phase noise with probability $> 10^{-9}$. As expected, binary data can be received with peak-to-peak $t_{\text{error}} < 1T_{\text{symbol}}$, while 2 bits can be received with $t_{\text{error}} < 0.67T_{\text{symbol}}$, 3 bits with $t_{\text{error}} < 0.46T_{\text{symbol}}$, and 4 bits with $t_{\text{error}} < 0.32T_{\text{symbol}}$. The LENOB is larger for small t_{error} because small timing errors result

in smaller p-p voltage errors at the signal peaks than asynchronous RMS noise.

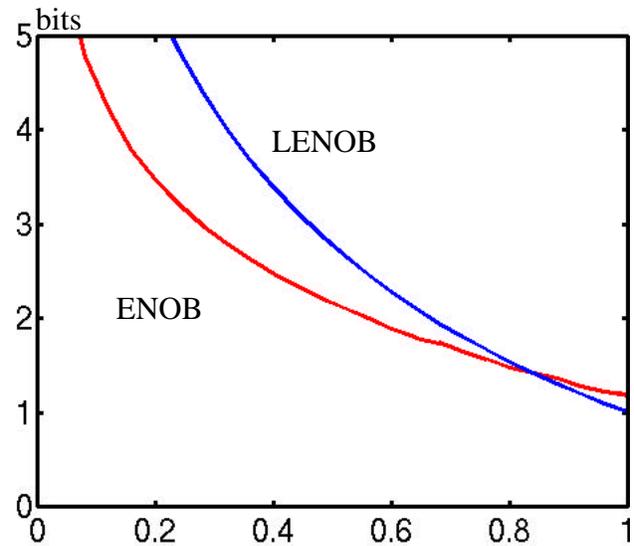


Figure 2.4: LENOB, ENOB vs $t_{\text{error}}/T_{\text{symbol}}$

To complete the performance analysis of a link, statistical noise is root-sum-of-squares combined¹ and the p-p value (with probability $< 10^{-9}$) added to other p-p sources of voltage errors. To set a timing noise requirement for this work, half the p-p voltage noise is budgeted for timing noise and half for other noise, so that the timing error budget for an N-bit link is equal to the timing error that achieves an LENOB of N+1 with no other noise. Also note that t_{error} includes timing errors from both the transmitter and receiver, which again add in root-sum-squared fashion since they are uncorrelated and normally distributed. Thus, an N-bit link needs clocks with $\frac{1}{\sqrt{2}}$ the timing noise that corresponds to an LENOB of N+1 (from Figure 2.4). Given the target of $T_{\text{symbol}} = 1 \text{ FO4}$, the p-p timing noise budget for a 2 bit link is 0.33 FO4, and for a 3-bit link is 0.23 FO4.

Given this framework for analyzing link performance, each of the circuit blocks in a link transceiver is examined, beginning with the clock generation circuitry.

2.2 Clock Generation

Clocks with very low timing noise are needed for ADCs and DACs to be possible at a symbol-time of 1 FO4. In addition, the ADC clocks in a link receiver must track the phase of a noisy received signal. To support time-interleaving, multi-phase clocks are required with timing noise that is a small fraction of the clock period. With 8-way time interleaving, the 3-bit link timing noise budget of $0.23T_{\text{symbol,pp}}$ corresponds to only 3% of a cycle of the multi-phase clocks. Ring oscillators can generate clocks with low timing noise, but work best when locked to stable timing sources. Since the received signal is noisy, a dual-loop architecture is used [49] to generate clocks with low jitter and to phase-lock to the received signal. The primary loop locks to a stable, low noise reference clock for low jitter, while the second loop aligns the clocks to the input data.

1. The approximate square law mapping of timing to voltage noise (see Equation 2.2) increases large amplitude (low probability) noise so that the p-p noise voltage (with probability $< 10^{-9}$) is larger than 6σ (where σ is the RMS noise voltage). To compute an upper bound on the noise, the p-p noise voltage due to timing noise is root-sum-of-squares added to p-p thermal noise voltage sources.

The primary phase-locked loop (PLL) is based on a Voltage Controlled Oscillator (VCO), shown in Figure 2.5. The output frequency and phase of the VCO are controlled by V_{ctl} , which varies the delay in the four differential delay elements. The VCO in this work generates a clock with a period of 8 FO4 delays, so that each delay element has a nominal delay of 1 FO4 (note the wiring crossover to invert the signal after the rightmost buffer). Four differential clocks are available, allowing 8 clocks with a nominal phase spacing of 45° to be generated.

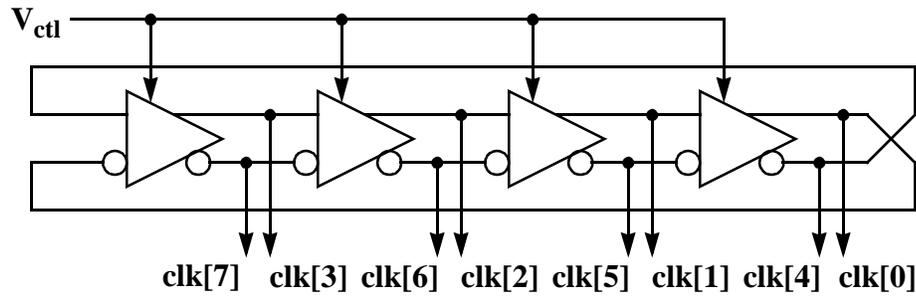


Figure 2.5: VCO Architecture

The differential delay elements in the VCO (see Figure 2.6) use replica-biased symmetric loads for high supply noise rejection [48]. These loads present a resistance that is symmetric over the signal swing from v_{cp} to V_{dd} . The effective load resistance is adjusted to control the RC time constant and thus the delay of the buffer. A replica bias cell adjusts v_{cn} in concert with v_{cp} so that the decreasing resistance of the PMOS diode load with output voltage is symmetric with the increasing resistance of the PMOS load biased with v_{cp} .

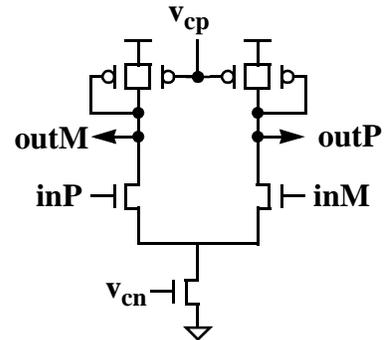


Figure 2.6: Delay Element

Random thermal noise causes variations in the delay of the VCO buffers both directly and by coupling to v_{ctl} . An analysis of the effect of random noise in VCOs shows that the generated phase noise spectrum from a well-designed ring oscillator

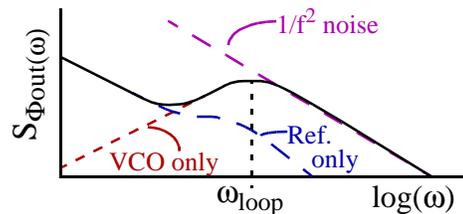


Figure 2.7: PLL Phase Noise Spectrum

VCO is dominated by upconverted white noise with a $1/f^2$ spectrum. Figure 2.7 shows a

typical plot of the output phase-noise power spectral density of a VCO with a noisy input reference clock [51]. Low frequency noise below the PLL loop bandwidth ω_{loop} does not cause a significant phase error before the loop responds and tracks it out. Thus, low frequency $1/f^2$ noise is reduced and the random VCO noise generated in a PLL peaks at the loop bandwidth. However, low frequency noise on the input reference is tracked by the loop and dominates the output noise spectrum at low frequencies. This plot illustrates how higher loop bandwidth ω_{loop} reduces $1/f^2$ phase noise in a ring oscillator.

To minimize phase noise due to thermal, supply and substrate noise, the main PLL loop bandwidth ω_{loop} should be as high as possible, but must remain well below the reference clock frequency for stability. The PLL is self-biased so that its loop bandwidth tracks the PLL operating frequency and its damping factor remains constant. This allows the PLL to be optimized for low phase noise over a wide frequency range [48].

To analyze the impact of phase noise on a high speed link, the power spectral density of phase noise can be converted to time-domain jitter by integrating phase noise over frequency, or direct time-domain analysis of jitter can be performed [57][58]. Empirical data on time-domain jitter is also available from several chips with the same PLL design used in this work. The quiescent jitter data compiled in Table 2.1 indicates that the jitter design goal of 3% of a clock cycle is possible to attain. The supply sensitivity data indicates that the clock generation supply must be isolated carefully to allow integration with noisy digital processors and logic.

Description	Multiplier	Quiescent Jitter, p-p	Supply Sensitivity, p-p
900 MHz PLL [53]	1	2.3%	1.3% / %
900 MHz DLL [53]	1	2.3%	0.9% / %
1 GHz PLL [12]	4	1.8%	1.8% / % (simulated)
500 MHz PLL [2]	2	--	0.7% / %
1 GHz PLL [7]		1.9%	--

Table 2.1: Measured PLL and DLL Jitter

The high supply noise sensitivity comes from two sources. First, supply noise above ω_{loop} causes variations in the delay of the VCO buffers that are proportional to the clock period. Therefore, supply noise sensitivity is expressed as the percentage of a clock cycle of jitter caused by square wave supply noise. Second, some of the supply sensitivity is due to the clock buffers after the PLL. Simulation shows that the delay of a 0.25 μ m CMOS inverter varies 0.9% per % supply change. Thus, a short clock buffer chain (after the PLL) of 4 inverters with a delay of 4 FO4 has a supply sensitivity of $0.36FO4=0.45\%$ of an 8 FO4 clock cycle per % supply change. By isolating the clock generation supply and minimizing supply sensitivity, the timing noise budget of 3% of a clock cycle derived in Section 2.1 can be met for a 3-bit transceiver with a 1 FO4 symbol time.

Thus, clocks with low jitter and supply sensitivity can be generated in a ring oscillator locked to a stable clock source, but the clocks have no phase relationship to the received signal, and could be slightly off in frequency. To lock the clocks to the received signal, a second, delay-locked loop (DLL) is added to the primary PLL, as shown in Figure 2.8 [49]. The primary PLL (**upper loop**) phase-locks a multi-stage VCO to a stable reference clock, with high loop bandwidth to minimize random and coupled noise. The secondary DLL (**lower loop**) generates receiver clocks from the VCO clocks, adjusting their phase to the receive signal, using low bandwidth digital filters to decrease noise from the receive signal.

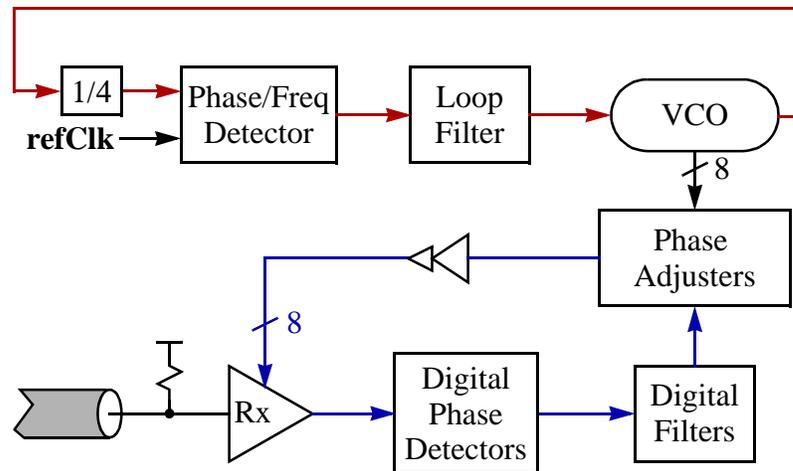


Figure 2.8: Dual Loop Clock Generation

The DLL uses phase adjusters to generate clocks of arbitrary phase from the VCO clocks. The phase adjusters are based on clock interpolators that use the same technology as the VCO delay elements. Each clock interpolator synthesizes a clock from two input clocks by digitally varying the current of two symmetric load buffers (same circuit as in Figure 2.6) with their outputs connected to a shared load. At the extremes, when the current all goes to the upper or the lower buffer, the output clock is a delayed version of the upper or the lower input clock, respectively. When the current is divided between the two buffers, the output clock phase varies between the extremes.

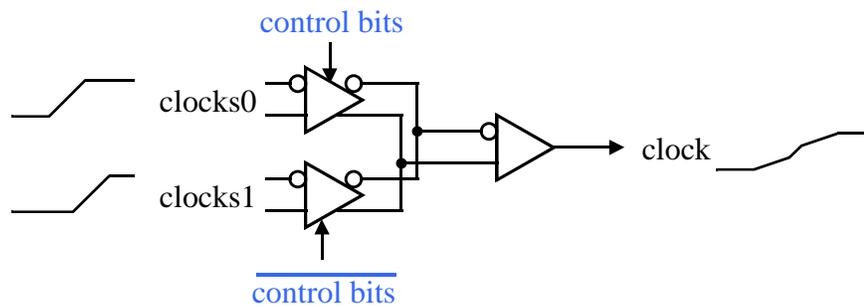


Figure 2.9: Clock Interpolator

Consider the generation of a clock that is phase locked to a receive signal of slightly lower frequency than **refClk**. As the receive signal drifts earlier in phase, the DLL adjusts so that the clock interpolator buffer with the earlier clock input has more current. When all the current is switched to the earlier clock buffer, the later input clock of the clock interpolator must be changed to an even earlier clock to accommodate further receive signal phase drift. Thus, the phase adjusters use multiplexors to select pairs of VCO clocks (with 45° phase spacing) for each clock interpolator. Digital control of the current in the clock interpolator buffers allows the current control to be coordinated with the switching of the multiplexors, permitting endless phase adjustment so that frequency differences can be tracked. Good linearity has been measured with segmented tail current sources, achieving more than 3 effective bits of delay resolution with a 4 bit control word [1][6][12].

Thus, a dual-loop architecture allows a high bandwidth, low phase-noise PLL to be combined with a lower bandwidth DLL that tracks the receive signal but rejects high frequency receive phase noise. However, transistor mismatches and layout asymmetries

result in static timing errors between otherwise identical VCO and DLL delay cells and subsequent clock buffer chains. This results in static phase errors that shift the multiple clock phases in a time-interleaved transceiver away from the nominal spacing (45° with 8-way interleaving). Static errors did dominate the timing noise in [6][11], but can be corrected with timing calibration circuitry based on clock interpolators (described in Section 4.4 on page 51).

The dual-loop multi-phase clock generator described in this section, with static timing error correction, phase-locks to the receive signal with timing errors that approach that of a PLL locked to a stable reference, meeting the timing noise budget of 3% of a clock cycle derived in Section 2.1. The next two sections explore the accuracy and bandwidth of the remaining transceiver circuit blocks, the transmitter and receiver.

2.3 Transmitter Design

A high speed DAC can be based on the open drain transmitters used in high speed links, with their high output impedance allowing accurate transmission line termination with a parallel resistor. The simplest and highest performance drivers use NMOS current sources terminated to V_{dd} , avoiding the use of lower mobility PMOS transistors. The output current is converted to voltage by the 25Ω impedance formed by the 50Ω line and the 50Ω termination. The output impedance of the open drain drivers remains well above 50Ω as long as they remain saturated. Because the output transistors enter the saturation region at reduced V_{ds} due to velocity saturation [39], an output swing of 1V is possible in $0.25\mu\text{m}$ CMOS ($V_{dd}=2.5\text{V}$).

Differential pairs are typically used to generate differential transmitter signals to reduce coupling and supply noise, even if a single ended signal is ultimately transmitted on the wire, as shown in Figure 2.10. However, the **tail** node of the differential pair must remain above $V_{bias}-V_{th}$ to keep the tail current source

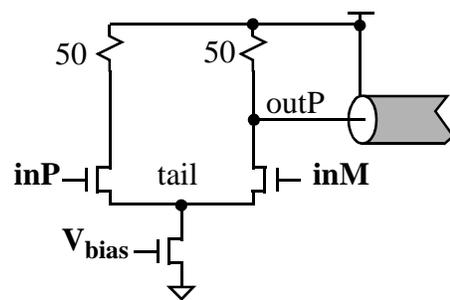


Figure 2.10: Open Drain Transmitter

saturated. This reduces the V_{gs} of the output transistors and requires larger transistors for the same output swing.

The capacitance of the large output transistors in a time-interleaved transmitter (with architecture similar to the receiver in Figure 2.2) can limit its bandwidth because many output transistors are connected in parallel. However, the transmitter

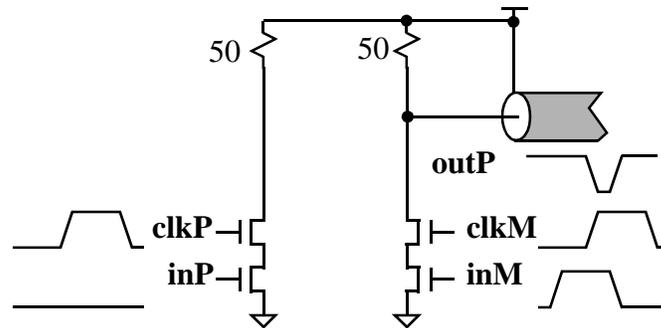


Figure 2.11: Grounded Source Tx

in [11] minimizes the size of the output transistors by eliminating the tail current source as shown in Figure 2.11, and achieves a 1 FO4 bit time. The output current is adjusted by controlling the swing of the predriver outputs **inP/M** and **clkP/M**. A single output transistor would require a narrow clock pulse, which is difficult to generate and buffer accurately. Instead, each interleaved transmitter is enabled by the overlap of two pulses with 50% duty cycle. The bottom clock pulses **inP/M** are qualified with the data to be transmitted, and go high first. The output pulse starts when the top clock goes high, and ends when the bottom clock goes low.

Providing more than one bit of resolution is conceptually simple in an open-drain transmitter, as the large output transistors can be divided into smaller transistors that are individually controlled. The outputs of the smaller transistors are then summed onto the 25Ω output impedance. A 4-bit transmitter has been constructed in this way to implement a linear equalizer [10]. The transmitter presented in Chapter 4 uses a similar approach with 8-bit resolution, with time-interleaving to increase the transmitted symbol rate. The higher resolution reduces quantization errors in equalization and in correction of nonlinearities and interference, and avoids the need for more than 4 bits of resolution in the receiver, where extra bits are more difficult to implement.

2.4 Receiver Design

High speed CMOS receivers typically use an NMOS passgate transistor to sample the input, and a regenerative latch [1][11][15][17] to provide high gain. The receiver design used in [11] is shown in Figure 2.12. This binary receiver has the high bandwidth that NMOS passgate samplers provide. The input common mode is near $V_{dd}/2$ to allow reasonable V_{gs} on the sampler transistors, and to keep the input transistors in the comparator saturated. Capacitors are used to couple the high common mode output of the transmitter (see Section 2.3) to the receiver. An alternative approach has also achieved high performance by using a regenerative latch to directly sample and amplify the input [10].

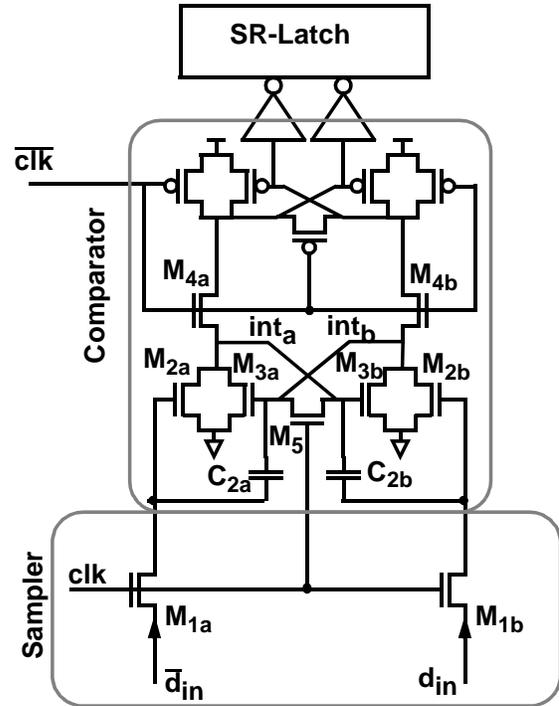


Figure 2.12: Passgate/Latch Receiver

Again, extending binary receiver circuits to provide more than one bit of resolution is conceptually simple - the input merely needs to be compared against multiple reference voltages. Many ADC architectures have been proposed, but the flash architecture shown in Figure 2.13 is the simplest and fastest, and is appropriate for low resolution ADCs. In comparison, multiple-step architectures reduce the

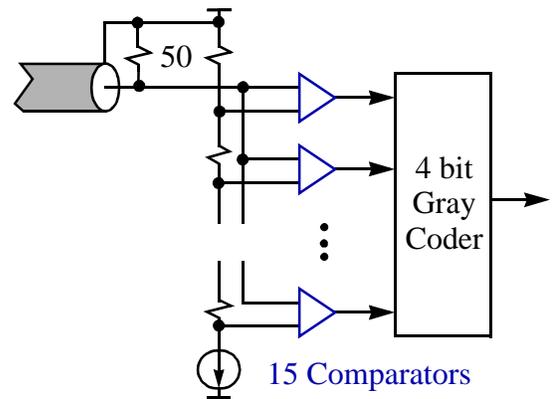


Figure 2.13: Flash ADC Architecture

number of comparators needed, but require sample/hold amplifiers that load the input, as well as DACs and adders, so they do not have clear advantages over flash architectures at 4 bits of resolution. Averaging architectures [28] hold promise for reducing comparator

offsets, but the resistors used to average the outputs of the pre-amplifier stages reduce input bandwidth. Folding and interpolating architectures [21][29] use somewhat similar resistor networks to reduce the number of pre-amplifiers, and the resistors again reduce input bandwidth¹. This work focuses on the simplest and highest bandwidth choice, the flash architecture.

The key circuit in a flash converter is the comparator, which dictates the performance of the ADC. The accuracy of CMOS comparators is limited by random variations in transistors, which causes mismatches between nominally identical transistors that to first order vary with $\frac{1}{\sqrt{WL}}$ [31]. Thus, the small transistors needed for low input capacitance and high bandwidth suffer large mismatches. The effect of random transistor mismatches on the switching point of a comparator is expressed as an input-referred offset voltage that accounts for the variations in all the transistors in the comparator.

Switched capacitor techniques are frequently used to correct the input offset voltage of a CMOS comparator. Figure 2.14 shows a circuit that subtracts the input offset of a comparator from the input. The input offset is stored on capacitor C_h during a reset cycle with the switches flipped. Unfortunately, C_h must be significantly larger than the nonlinear input capacitance of the amplifier to achieve good linearity. The parasitics of the large C_h and the required large switches reduce the sampling bandwidth of the comparator.

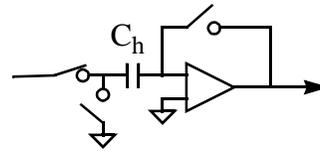


Figure 2.14: Offset Cancellation

Both the use of large transistors and series offset cancellation techniques limit the bandwidth of comparators. Increasing the sample rate through time-interleaving exacerbates the problem by increasing input and output capacitance. Perhaps as a result, previous work in time-interleaved ADCs [18][19][20] has focused on area and power savings at much lower sample rates. The receiver presented in Chapter 3 increases

1. Sampling before averaging or interpolating would allow high bandwidth and may avoid offset calibration or reduce input capacitance, respectively.

accuracy while maintaining bandwidth through the use of calibration techniques that typically are used to increase the resolution of ADCs beyond 10 bits [26].

2.5 Summary

This chapter begins with the development of a framework for analyzing link performance that uses the FO4 delay as a metric for comparing across CMOS processes. A time interleaved transceiver is presented that shows that 8 way time-interleaving allows the symbol time target of 1 FO4 to be achieved. So that the circuits don't limit the link performance, a target of $\tau = 0.3\text{FO4}$ is set for the input and output time constants of the ADC and DAC, corresponding to Nyquist bandwidth. Peak noise with probability greater than the bit-error-rate target of 10^{-9} is used to derive the Link Effective Number of Bits, and a simple relationship between timing noise and voltage noise is derived, showing that a p-p timing noise target of 0.23 FO4 allows a 3-bit transceiver to be built.

The remainder of the chapter addresses the challenges in each of the major circuit blocks of a link to match the raw speed that time-interleaving makes possible. Previous clock generation circuitry has achieved excellent performance that meets the timing noise budget of 0.23 FO4, p-p. A dual loop architecture provides low jitter clocks with a high bandwidth inner loop to reduce random noise from the VCO, and phase locks the clocks to the received signal with lower bandwidth to reduce phase noise in the outer timing recovery loop. Static timing offsets due to random transistor mismatches have dominated the timing errors in high speed links, but clock interpolators were presented that can be used to correct static and low frequency timing errors.

A review of high speed transmitter designs indicates that multiple bits of resolution can be achieved by dividing up and individually controlling the open-drain output transistors. This should not significantly reduce the bandwidth of the binary transmitter design which meets the 1 FO4 symbol time target.

A look at receiver and ADC design shows that multiple bits of resolution, high sample rate and high bandwidth are difficult to achieve simultaneously. Link receivers achieve the high sample rates and bandwidth needed, but the circuits used lack the accuracy to be directly used in this work. Conversely, high speed ADCs routinely exceed

8 bits of resolution, but do not operate at link speeds. The design of the receiver presents the most difficult challenges, so Chapter 3 focuses on the circuit design of a receiver to meet these challenges, using calibration to break the tradeoff between accuracy and bandwidth.

Chapter 3:

Receiver Design

The challenge in the receiver is to design a high sample-rate flash ADC with wide bandwidth and 4-bit accuracy. This challenge is met with comparators designed with a simple sample/hold amplifier, small transistors and digital calibration. The receiver described in this chapter is based on a 4-bit flash ADC, as shown in Figure 3.1. The single-ended input signal is coupled into fully differential comparators, as shown in Figure 3.1, with the reference ladder connected to the signal return.

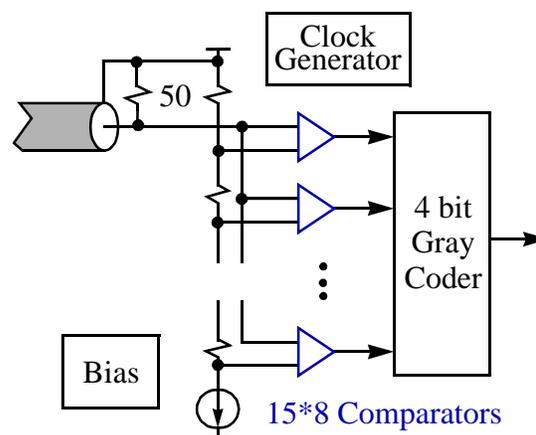


Figure 3.1: 4-bit Flash ADC

Time interleaving is used to achieve the sample period goal of 1 FO4 delay, as shown in Figure 3.2. Since each of the 8 interleaved ADCs comprises 15 comparators, a total of 120 comparators load the analog input. The design of these comparators is key to achieving the

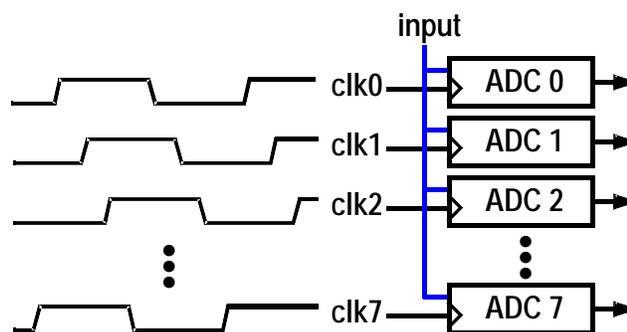


Figure 3.2: Time-interleaved ADC

bandwidth and accuracy goals of the converter. Furthermore, the sampling circuit is the key component of a comparator design. Therefore, this chapter first analyzes several types

of sampling circuits, setting a framework for the comparator design. The rest of the chapter focuses on the circuit design of the comparators, including the bias circuits and matched clock generators that support the comparators.

3.1 Sampler Design Analysis

A comparator needs to both sample and amplify the input to full swing. A regenerative latch is the fastest, lowest power method of providing large amplification of a sampled signal. Regenerative latches can also act as samplers, because the regeneration causes the sensitivity of the latch to peak when first enabled. This is a result of the exponential increase with time of the small-swing output voltage of an enabled regenerative latch:

$$V = V_0 e^{t/\tau} \quad [2.5]$$

Thus, if a latch has a time constant $\tau = 0.5$ FO4 gate delay, an impulse applied at the beginning of an 8 FO4 clock cycle, and sampled 4 FO4 delays later, as in Figure 3.3, is amplified by $e^{4/0.5} = 3000$. An impulse applied 1 FO4 gate delay after the beginning of the clock cycle and sampled 3 FO4 delays later is amplified by $e^{3/0.5} = 403$.

For small input signals, a regenerative latch followed by a sampler has a linear, but time variant impulse response, as shown in Figure 3.4. The voltage sampled at the latch output after 4 FO4 delays of amplification is more affected by impulses just after the latch is enabled than by later impulses. This exponential closing of the latch results in a sampling time-constant of τ , and the latch can be modeled as a low pass filter with time constant τ followed by an ideal switch, just like the model for a pass gate sampler.

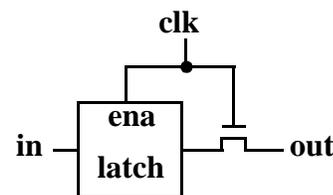


Figure 3.3: Sampled Latch

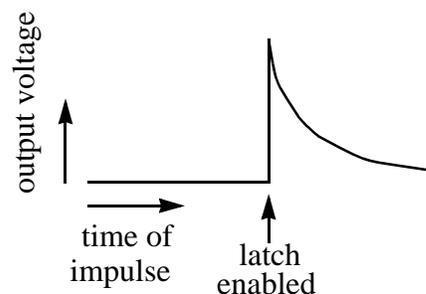


Figure 3.4: Latch Impulse Response

Most regenerative latches use cross-coupled inverters to provide gain, with additional circuitry to couple in the input signal. The latch shown in Figure 3.5 has excellent sampling bandwidth and regeneration rate, and has been used in the StrongArm processors [34] and in high speed receivers [10][15]. The **cross-coupled inverters** (upper left and upper right) are

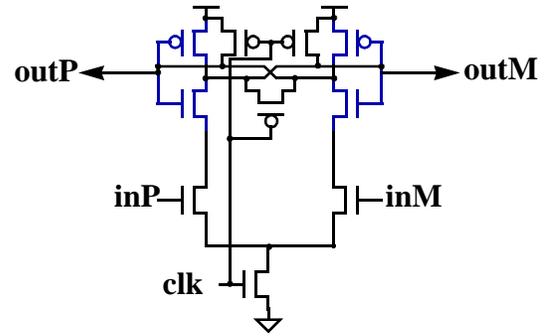


Figure 3.5: StrongArm Latch

loaded by each other, the three central PMOS reset transistors, and the next stage. Thus, the fanout (output to input capacitance) of the inverters is greater than one, and greater than two in any practical implementation. Consequently, a lower limit on the sampling time-constant of the latch is roughly a fanout-of-2 delay or 0.4 FO4. In simulation, a StrongArm latch has a sampling time-constant $\tau = 0.6 \text{ FO4}$, compared to the target of 0.3 FO4¹. The regeneration rate, and thus the sampling time-constant of the StrongArm latch, varies with the input common mode (which changes the current supplied to the cross-coupled inverters), making it unsuitable for accurate, high bandwidth sampling.

Because sampling the input is the most critical function in a flash ADC, a dedicated sampler circuit can improve performance as a front-end to a regenerative latch. The simplest sampler is an NMOS pass gate, which can be modeled as an ideal switch with an on-resistance R , and equal drain and source parasitic capacitances C_p , as shown in the

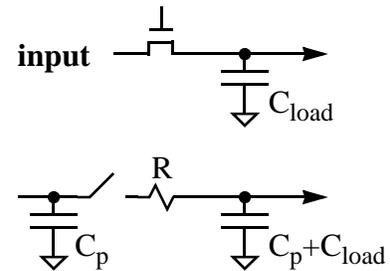


Figure 3.6: Pass Gate Sampler

bottom of Figure 3.6. Thus, the pass gate has a sampling time-constant $\tau = R(C_{load} + C_p)$, neglecting the load on the input for the moment. In order to achieve good sampling bandwidth and linearity, the pass gate must be nearly as large as the load.

1. The regeneration time-constant was measured for a small input signal with common mode $V_{dd} - 0.4V$

Consider the performance of a reasonably large passgate with $C_p = C_{load}$ (note that the load on the input is C_{load} when the pass gate is off, and $3*C_{load}$ when on). This passgate has a very fast sampling time-constant when sampling a small signal near ground under typical conditions. Signals very close to ground are not sampled well, because clock feedthrough as the passgate turns off pushes the sampled signal below ground, and subthreshold conduction attenuates the signal. Therefore, in $0.25\mu\text{m}$ CMOS with a 2.5V supply, the useful signal swing extends down to 0.25V, so that with $V_{th}=0.55\text{V}$, the maximum passgate $V_{gs}-V_{th} = 1.7\text{V}$. In $0.25\mu\text{m}$ CMOS, this results in a $g_{ds}=1200\mu\Omega^{-1}$ for a $1\mu\text{m}$ wide passgate transistor, which drives its own drain capacitance $C_p=4\text{fF}$ and an equal C_{load} , so that the sampling time-constant $\tau=8\text{fF}/1200\mu\Omega^{-1} = 6.7\text{ps} = 0.07\text{FO4}$. This corresponds to results in [38], and would technically allow sampling of a small-swing 150 Gbit/s signal.

However, the sampling time-constant is much larger for reasonable signal swings under worst case conditions. A signal swing of 1V both reduces worst-case V_{gs} and increases V_{th} due to body effect to 0.65V, and worst case process further increases V_{th} to 0.75V. With 10% low supply of 2.25V, worst case $V_{gs}-V_{th}=0.5\text{V}$ with $g_{ds}=300\mu\Omega^{-1}$ (from simulation). Thus, the worst case sampling time-constant $\tau=0.2\text{FO4}$ (27 ps). This time constant is slow enough to be unaffected by the 0.6FO4 fall time of a fanout-of-1 inverter chain driving the gate, although the fall time does increase the sampling time-constant when the input is low¹, from $8\text{fF}/672\mu\Omega^{-1}=0.1\text{FO4}$ to 0.15FO4. The capacitive load of 120 passgates on the 50Ω input (in parallel with the 50Ω termination) provides an additional low pass filter with time constant $25\Omega*120*2C_{load}>0.25\text{FO4}$, assuming $C_{load}>5\text{fF}$ for folded $1.5\mu\text{m}$ wide transistors, and that half the samplers are on (50% duty cycle clocks).

While passgate sampling time-constants are good, passgate samplers suffer from nonlinear variations in sampling time and sampling time-constant with input level, and from charge kickback. The effective sample time of a passgate sample varies with input level, since the passgate turns off when the clock (gate) drops below the input+ V_{th} . Thus,

1. Based on $1/(2\pi*\text{simulated } 3 \text{ dB bandwidth})$ on sampled sine waves with slow process, voltage, temp.

the sample time variation $\Delta_{t,\text{sample}}$ with input level is approximately the signal swing plus body effect divided by the 20/80% clock fall time, as shown in Equation 2.6 and confirmed in simulation:

$$\Delta_{t,\text{sample}} = (0.75 + 0.1V) / (0.8 * 2.5V / 0.6FO4) = 0.25FO4 \quad [2.6]$$

The sample time variation must fit in the timing noise budget, which is $0.23FO4$ for a 3-bit link (see Section 2.1). The variations in sampling time-constant with input level are less of a concern, since the worst-case passgate bandwidth is high. Variations in sampling time-constant on overdrive recovery from full swing voltage steps cause relatively small voltage errors for a symbol time $T=1FO4$:

$$V_{\text{error, rise/fall}} = 750\text{mV} (e^{-T/0.2FO4} - e^{-T/0.15FO4}) = 4\text{mV} @ T=1FO4 \quad [2.7]$$

Overall, the nonlinear variations in sampling time and bandwidth, the charge kickback, and the fairly large input capacitance of a passgate sampler make it difficult to achieve the receiver goals of a $1FO4$ symbol time, a $0.3FO4$ sample time-constant, and 4-bit accuracy. To avoid problems with nonlinearities in samplers, a differential pair can be used to subtract the input from the reference before sampling, driving a differential signal with a fixed common mode to the sampling circuit. The pre-amplifier also reduces the charge kickback from the previous sampled value onto the input and reference ladder.

Although it improves the linearity of the following sampling circuit, a pre-amplifier reduces the sampling bandwidth, with its time constant adding to that of the sampling circuit. Sampling bandwidth is significantly limited even by the high bandwidth open-loop pre-amplifier with grounded PMOS loads shown in Figure 3.7. This amplifier

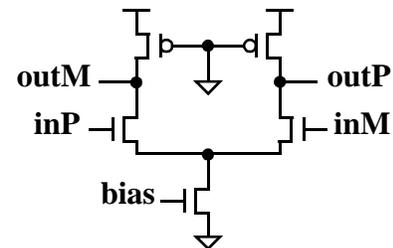


Figure 3.7: Wideband Pre-Amp

is designed for input and output signals that swing from V_{dd} to $V_{dd} - V_{th}$. Simulations of this pre-amplifier show that with a gain of 1, the pre-amplifier time-constant $\tau = 0.3 FO4$, meeting the design goal, and that linearity and kickback are improved over a passgate sampler. However, the overall sampling time-constant is increased by the sampling time

constant of the subsequent sampling circuit, and by the RC time-constant formed by the 25Ω impedance at the input (50Ω in parallel with 50Ω line) and the pre-amplifier input capacitance. The input capacitance is large since a large pre-amplifier is needed to drive the sampling circuit and the next stage load.

A switched differential amplifier that avoids a separate sampling circuit is used as the front-end of the comparators in this work. Because the comparator also needs offset correction, high gain, and an output that is held for a full clock cycle, there are three stages in the comparator, which is described in the next section.

3.2 Comparator Architecture

Figure 3.8 shows the three stages of the comparator: the front end sample/hold amplifier (S/H), a regenerative latch with a DAC for offset calibration, and an unclocked output latch that holds the comparator outputs for a full clock cycle. The comparator uses

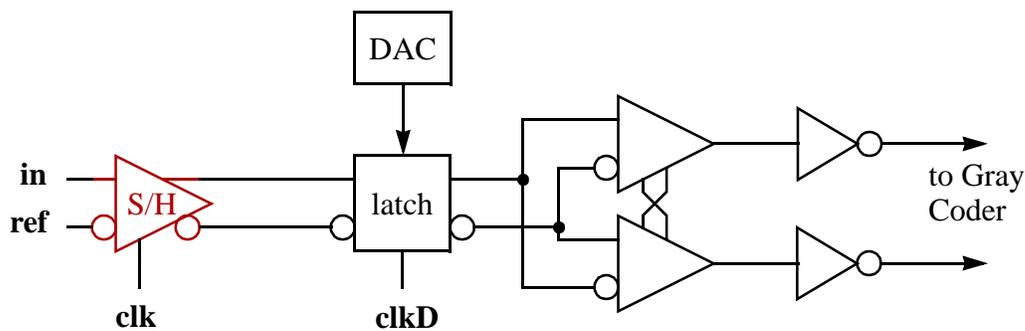


Figure 3.8: Comparator

a digital offset calibration technique in the second stage latch. To correct offsets in the comparator, the desired switching level for each comparator is periodically transmitted into the ADC, and the DAC in each latch is adjusted until the comparator outputs toggle. This approach breaks the tradeoff between bandwidth and offset voltage, allowing small transistors to be used for high bandwidth, and correcting for the large transistor mismatch errors that result. Correcting for offsets in the second stage of the comparator allows the use of a simple, high bandwidth sample/hold amplifier as the first stage.

3.3 Sample/Hold Amplifier

The front-end sample/hold amplifier (shown in Figure 3.9) has a gain of only 1, allowing it to achieve a small sampling time-constant of 0.3 FO4. Because the inputs are applied directly to transistor gates, there is little charge kicked back to the input and reference ladder. The circuit is essentially a switched differential amplifier. It amplifies

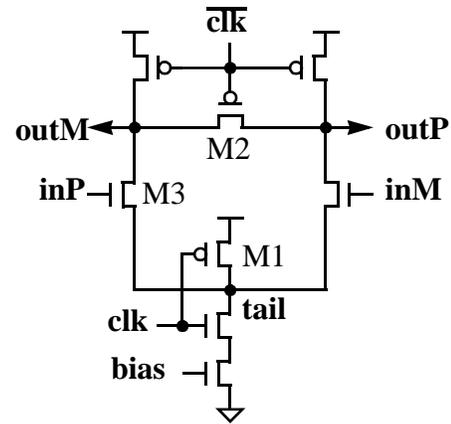


Figure 3.9: Sample/Hold Amplifier

when **clk** is high and $\overline{\text{clk}}$ low, enabling the tail current source and the output loads. It samples the input when **clk** falls and $\overline{\text{clk}}$ rises, disabling the tail current source and turning off the loads, so that the output is held in high impedance. The input (and output) of the comparator swings from V_{dd} to $V_{\text{dd}} - V_{\text{th}}$, so the input pair stays in saturation.

The high common-mode of the receiver signals increases headroom and bandwidth. As process features shrink and supply voltages decrease faster than V_{th} reduces, a simple differential pair has barely enough headroom¹. The high common-mode inputs require the differential pair to have high common-mode outputs to keep the input pair in saturation, precluding the use of diode loads. Grounded gate PMOS resistive loads work well up to V_{dd} , and have low parasitic capacitance, allowing a data path that only drives NMOS gates.

The central PMOS transistors M1 and M2 in Figure 3.9 both reduce the sampling time-constant. M1 actively pulls up the **tail** of the differential amplifier after **clk** falls, rapidly turning off the input transistors so that the sampling time-constant is nearly the same as the steady state time-constant of the amplifier. Although M2 appears to be a reset transistor, it is actually a differential load that reduces the amplifier output impedance, and thus its time constant, with low parasitic capacitance. There is no reset cycle needed for

1. For the 0.25 μm CMOS circuits in this work, with a signal swing of 1V and with $V_{\text{gs}}=1\text{V}$ on the input pair, the bias current source transistor can only have $V_{\text{gs}} - V_{\text{th}}=0.25\text{V}$ to remain in saturation (i.e. **bias**=0.8V) with 10% low supply voltage of 2.25V.

the amplifier, as its time constant is small enough that the previous sampled value decays during the half clock cycle (4 FO4 delays) that the amplifier is enabled.

The function of transistor M2 is illustrated in Figure 3.10. When **clk** is high and the amplifier is enabled, the three upper PMOS transistors act as resistive loads. Assuming they are the same size, the three PMOS transistors have similar impedances¹. The two pullup PMOS loads are represented by resistors **R**. Because the output of the amplifier is differential, there is a **virtual ground** at the center of the differential load transistor M2, so it can be represented by two resistors **R/2**. M2 improves the amplifier bandwidth by presenting half the impedance to each of the outputs for the same parasitic capacitance.

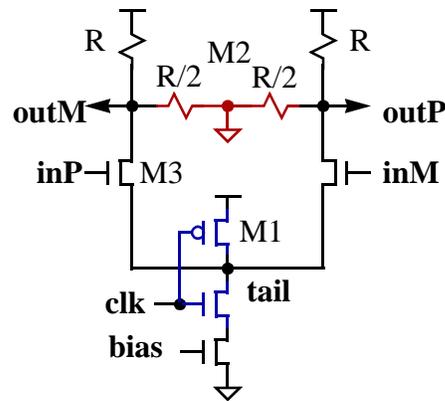


Figure 3.10: Differential Load M2

The sample/hold amplifier is able to sample accurately because it subtracts the reference from the input (for each comparator in the flash ADC), so gain nonlinearities are not an issue, and because it has little sample time variation with input common mode. Accurate sampling is important in a flash ADC, which has different reference voltage inputs to each comparator. Sample time variation is small because during amplification (of small signals), the **tail** node is a V_{gs} below the input common mode. When **clk** falls below the switching point of the “inverter” formed by M1 and the NMOS transistor below it, **tail** is pulled up at a constant rate by M2. The differential pair is disabled when **tail** rises by $V_{gs} - V_{th}$ (about 0.5V) a fixed time after the falling edge of the clock.

Thermal noise is a concern due to the small transistors, low gain and wide bandwidth of the sampler. The RMS noise current of a MOSFET is given by:

$$i_{\text{rms}} = \sqrt{4kT\gamma g_{d0}\Delta_f} \quad [2.8]$$

1. M2 has lower V_{gs} but also lower V_{ds} , so $g_{ds} = k(2(V_{gs} - V_{th}) - V_{ds})$ is similar.

where the noise factor γ has a value of unity for $V_{ds}=0$, and is typically 2-3 for short-channel NMOS devices in saturation [30][45][46][47]. $\Delta_f = \frac{3}{8RC}$ is the noise bandwidth of the RC filter¹ at **outP**, and g_{d0} is the drain-source conductance at zero V_{ds} . The differential noise voltage while the sample/hold amplifier is enabled, (assuming² half of the noise voltage at **outM** due to a noise current into **outM** appears at **outP**) is $v_{rms} = i_{rms} \frac{R}{3}$. The input-referred noise voltage is v_{rms} divided by the gain $G=g_{m3} \frac{R}{3}$ and multiplied by $\sqrt{2}$, since there are two equal, uncorrelated noise sources from the two transistors in the input pair. The noise from the input pair dominates since it is more than twice as large as the noise from the PMOS load transistors and next stage. The total input-referred noise voltage for the sample/hold amplifier in this work, with $g_{d0}=1\text{mA/V}$, $R=5\text{K}$, $C_{outM}=7\text{fF}$, $g_m=0.5\text{mA/V}$, $\gamma=2.5$, and $kT=4.7e^{-21}$, is estimated by:

$$v_{rms,in} = \frac{3\sqrt{2}R}{3g_m R} \sqrt{4kT\gamma g_{d0} \frac{3}{8RC_{outM}}} = \frac{1}{g_m} \sqrt{kT\gamma g_{d0} \frac{3}{RC_{outM}}} = 2\text{mV}_{RMS} \quad [2.9]$$

When the attenuation of noise by the RC filter at the complementary output is considered, the noise from the input pair is 2.5mV_{RMS} . Noise from the loads and 2nd stage latch increase the total input referred noise voltage to 3.1mV_{RMS} , with a corresponding 6σ noise voltage of $19\text{mV}_{pp} = 0.4\text{LSB}_{pp}$.

While the sample/hold amplifier has excellent performance and is simple in concept, it must be carefully biased and clocked to control the output common mode. The bias voltage is set with a replica bias circuit to provide an output swing from V_{dd} to $V_{dd}-V_{th}$ during amplification [15], keeping the differential pair in saturation. At the sample

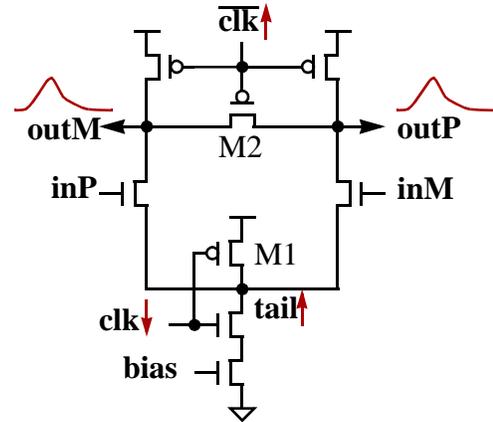


Figure 3.11: Sampling

1. The noise bandwidth of an RC filter is $1/(4RC)$, see [30] page 246.
2. This is conservative since the RC filter at **outP** attenuates high frequency noise. Also, since $4FO4 \gg \tau$ (the sampler time-constant), enabling the latch for only 4 FO4 doesn't affect the noise bandwidth.

time, when **clk** falls and $\overline{\text{clk}}$ rises as illustrated with arrows and waveforms in Figure 3.11, three effects move the output common mode:

1. $\overline{\text{clk}}$ pushes outputs up, coupling through parasitic PMOS gate-drain capacitances
2. $\overline{\text{clk}}$ rises before **tail** rises (due to **clk** falling), so the differential pair pulls outputs down as load impedance increases
3. second stage latch is enabled after 1 inverter delay; coupling pulls outputs down

Figure 3.12 shows simulated waveforms of the outputs of the sample/hold amplifier. During amplification, when **clk** is high, the output common mode settles to $V_{\text{dd}} - V_{\text{th}}$. As $\overline{\text{clk}}$ rises, the outputs are pushed up, and the gain also increases as the output load impedance increases. The outputs drop again as the second stage latch is enabled by **clkEna**, causing internal nodes to fall and couple back to the outputs of the sample/hold amplifier.

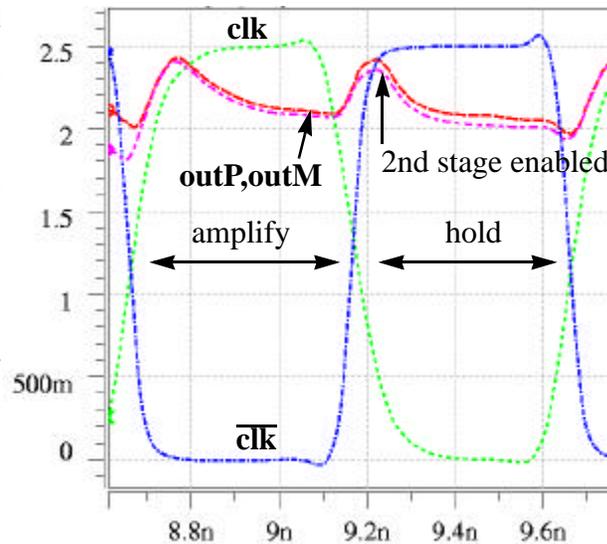


Figure 3.12: Sampling Waveforms

Reduced gain and nonlinear distortion can occur if the outputs drop below $V_{\text{dd}} - V_{\text{th}}$ in hold mode, as the input pair can turn into source followers since **tail** is pulled up to V_{dd} . Fortunately, the V_{th} in this case is a body-affected NMOS V_{th} , which can be 15% larger than for a grounded source NMOS transistor. The outputs must not rise above V_{dd} in the transition from amplify to hold mode, or else subthreshold conduction in the PMOS loads will reduce the amplify/hold gain in a nonlinear (voltage dependent) fashion.

The sample/hold amplifier is sensitive to clock skew between **clk** and $\overline{\text{clk}}$, with 15 ps = 0.12 FO4 of clock skew (due to a 16% increase in clock driver load) causing the peak of the output common mode to change from V_{dd} to $V_{\text{dd}} - 150\text{mV}$. To minimize process dependent skew, a matched 3/2 inverter chain [43] is used to generate **clk** and $\overline{\text{clk}}$ from a single clock for each of the time interleaved ADCs. Simulation shows that the sample/hold amplifier operates best when the **tail** node rises after $\overline{\text{clk}}$ rises. The delay through the three

upper clock inverters in Figure 3.13 is set to be slightly longer than the delay through the two lower clock inverters and the “inverter” driving the **tail**.

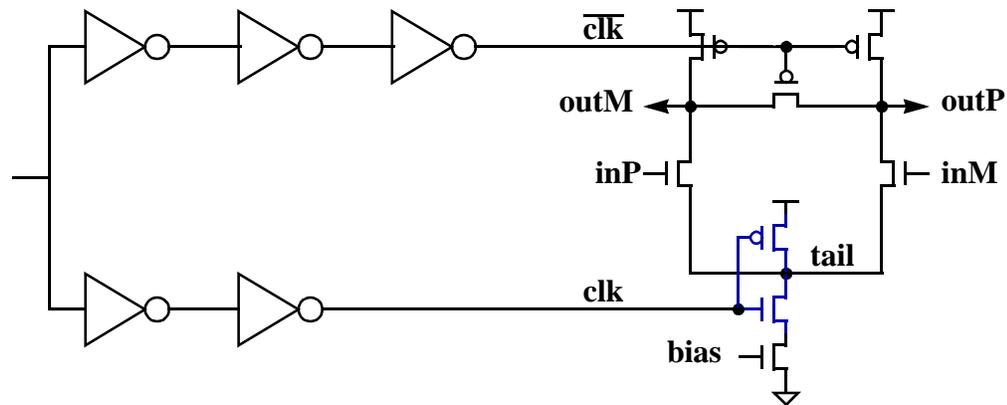


Figure 3.13: Matched Clock Generation

Since the inverter pairs in the matched clock generator can be matched well across process, temperature and voltage variations, the combination of the replica **bias** generator and the matched clock generator produces a differential output with common mode near $V_{dd} - V_{th}/2$. Requirements on the second stage latch are eased with a differential input with a stable common mode, and by performing high bandwidth sampling in the first stage.

3.4 Second Stage Latch

The second stage **latch** regeneratively amplifies the value that is held on the output of the sample/hold amplifier, and compensates for offsets throughout the comparator (shown in Figure 3.8). To correct offset errors, the desired switching point for each comparator is driven into the ADC input, and a 4-bit **DAC** is adjusted until the comparator outputs toggle. Offsets can be calibrated in conjunction with (periodic) link equalizer training, so that the entire ADC can be included in the correction loop and circuits are not burdened with switches and offset storage capacitors. This approach corrects offsets in the simple, high bandwidth sample/hold amplifier from Section 3.3, and allows the use of a regenerative latch as the primary gain element in the comparator.

The regenerative latch is based on the latch used in the StrongArm processors [10], with a current-mode DAC added to correct offsets, as shown in Figure 3.14. This latch was chosen because it resets completely and amplifies well, with a gain of 148 at the short clock period of 8 FO4 used in the

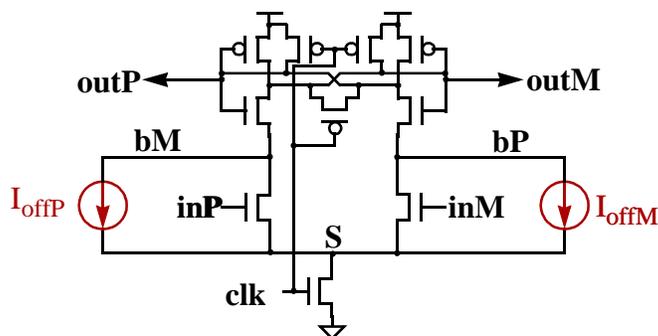


Figure 3.14: Latch with Offset Correction

ADC. When clk is low, the latch resets and its outputs are pulled high. When clk rises, node S is switched low, and the inputs are amplified onto nodes bM and bP , causing one to fall more rapidly. The cross-coupled inverters then regeneratively amplify to full swing. The current-mode DAC is connected to internal nodes of the latch, so it has no effect on the input capacitance, and only slightly slows the regeneration of the latch.

3.4.1 Offset Correction DACs

Each offset correction DAC consists of 2 sets of NMOS current sources (one set is shown in Figure 3.15), connected between node S and nodes bM and bP . Four predrive

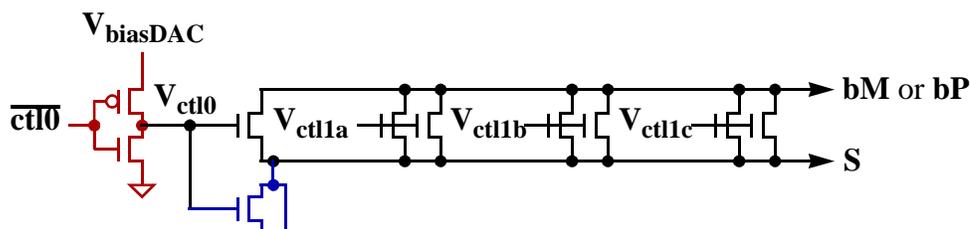


Figure 3.15: Offset Correction DACs

“inverters” either enable or disable each current source by driving its gate to V_{biasDAC} or to Gnd, respectively (only the one for V_{ct10} is shown). The step size of the DAC is controlled by V_{biasDAC} , which is held at a constant voltage by the bias circuit described in Section 3.4.2. The digital inputs of the DACs ($\overline{\text{ct10}}$ through $\overline{\text{ct13}}$) are periodically adjusted during offset calibration to minimize the input offset voltage of each comparator. The least significant bit directly drives $\overline{\text{ct10}}$, while $\overline{\text{ct11}}$ and $\overline{\text{ct12}}$ are thermometer encoded into V_{ct11a} , V_{ct11b} , and V_{ct11c} which drive three pairs of identical current source transistors.

The most significant control bit $\overline{\text{ctl3}}$ determines whether the 3-bit DAC connected to bM or to bP is enabled, so that each comparator has a total of 4 bits of offset correction. The [dummy transistor](#) at the bottom of Figure 3.15 is included in the layout for symmetry, with its drain and source connected to S . This connection makes the capacitive coupling from the out and S nodes (which drop by about the same voltage when the latch is enabled) similar on V_{ctl0} and the V_{ctl1x} control lines.

The two offset correction DACs in each latch generate digitally controlled currents that cancel offsets throughout the ADC. Since the first stage has a small and constant gain, its output offset can be corrected in the second stage. The high performance latch presented in this section is used as the primary gain element in the comparator, without requiring closed loop offset correction that would reduce performance. Even reference ladder errors, clock coupling, and offsets in the subsequent digital flip-flops can be corrected, since the final output of the ADC is used to determine the sign of the residual offset and close the correction loop.

The latch has adequate supply noise rejection because it is a differential circuit, and because the input gain doesn't vary much with changes in input common mode, since the input transistors have large V_{gs} and are velocity saturated. Offsets in the latch are modeled as normally distributed variations in V_{th} and effective transistors widths with standard deviations that vary inversely with the square root of the area (width*length) of the gate [32]. Simulation shows that the input-referred offset voltage of the comparator (due to transistor mismatches) changes by less than 0.5% with a 100mV change in V_{dd} , and by 0.2%/°C. Thus, $\text{V}_{\text{biasDAC}}$ should be biased so that the input-referred offset correction voltage is also insensitive to supply and temperature.

3.4.2 Offset DAC Bias Generator

The bias generator in Figure 3.16 outputs $V_{\text{biasDAC}} = 2V_{\text{th}}$ by driving a small current from three weak PMOS transistors into two NMOS diode-connected transistors. An external current (not shown) can be added to increase V_{biasDAC} in case offsets are larger than estimated with simulation. This results in an input-referred offset voltage that varies by 18% for a 100mV change in V_{dd} , and by 0.3%/°C. The supply

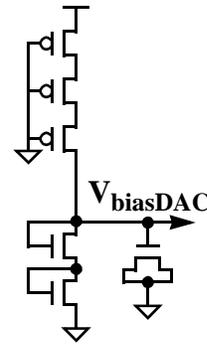


Figure 3.16: Offset DAC Bias

sensitivity is due to variations in input common mode that cause variations in input gain in the second stage latch. A bias generator that allows V_{biasDAC} to track the input common mode would improve the supply sensitivity of the offset-corrected receiver, but the circuit described here is adequate for 4-bit accuracy (with periodic offset correction). The remaining challenge in the comparator is to sample and hold the latch output at its peak (just before reset), with low hysteresis.

3.4.3 Latch Output Sampling

The outputs of the second stage latch are high during reset, and then are regeneratively amplified until sampled by the next stage, as shown in Figure 3.17. The next stage should sample the latch outputs just before it is reset to achieve a maximum gain in the latch. This is particularly important with the short clock period of 8 FO4, which enables the latch for nominally 4 FO4.

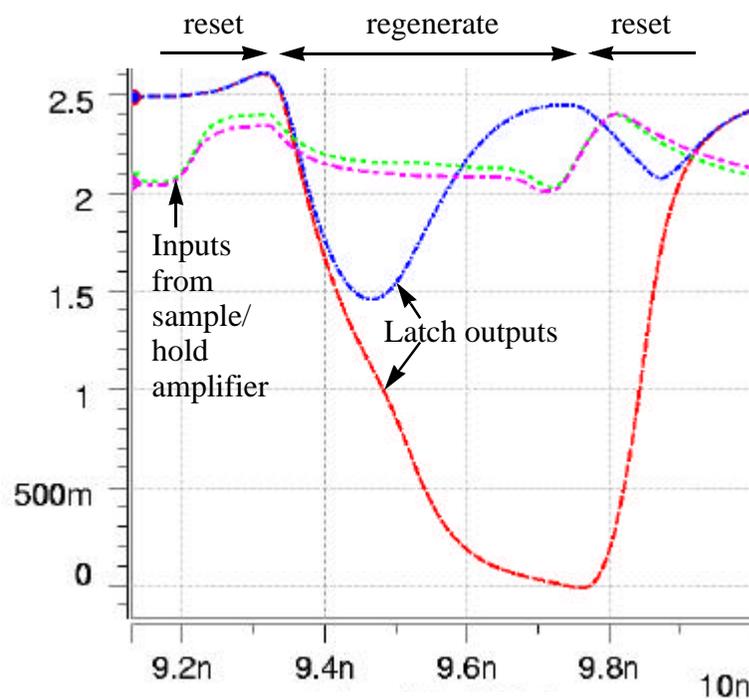
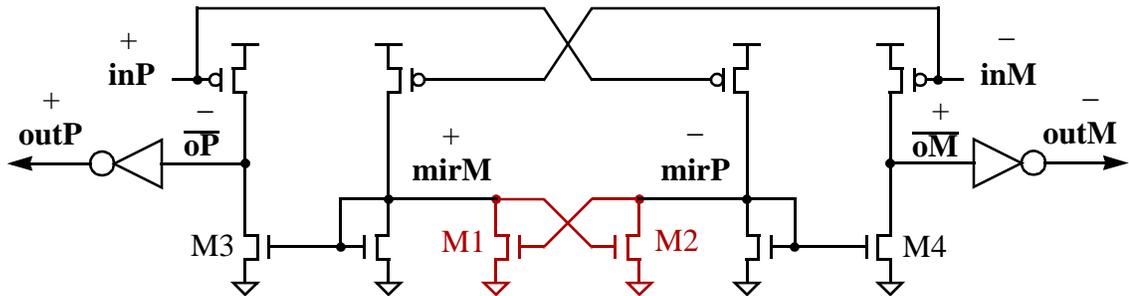


Figure 3.17: Output Latch Waveforms

The latch has an excellent regeneration time-constant $\tau=0.6FO4$, but 1 FO4 is needed to enable the latch and couple the input in, so the latch gain is $e^{3/0.6}=148$. Due to its finite sampling time-constant, a clocked sampler would effectively sample the latch outputs about 1 FO4 delay before the latch is reset, so the gain of the latch would be $e^{2/0.6}=28$. Because it is good practice to design circuits with some operating frequency margin to tolerate clock skew and duty cycle errors, the gain of the latch for a 7 FO4 clock period is considered, and is $e^{1/0.6}=5.3$. This small gain would allow the hysteresis and offset of the next stage(s) to dominate the input referred hysteresis and offset of the comparator, so a better output latch is used.

3.5 Output Latch

An unlocked output **latch** captures the output of the second stage latch without losing part of the cycle to clock skew. The unlocked latch is based on simple op-amps with PMOS inputs that hold their value when the second stage latch outputs go high during reset. Complementary CMOS drivers are used to drive the long wires out of the comparator array to the Gray coder so there is no net charge transfer out of the comparator, and thus no local supply bounce.



PMOS 2 x min. size; M3,M4 1.5 x min. size; other NMOS min. size

Figure 3.18: Output Latch Schematic

The basic operation of the output latch can be understood with **inP** high and **inM** low, as denoted with “+” and “-” in the schematic in Figure 3.18. The upper PMOS transistors pull **oM** and **mirM** high, and **mirM** high pulls **oP** low and **mirP** low. Then **inM** and **inP** go high when the second stage latch is reset, turning off the PMOS input transistors and holding the values at **oM** and **oP**. The output latch amplifies the peak

output voltage of the second stage latch (peak gain 148) by 31 for an overall comparator gain of 4600 with a clock cycle of 8 FO4 delays.

The cross-coupled NMOS transistors **M1** and **M2** in Figure 3.18 are needed to avoid degraded output levels. A careful examination of the second stage latch outputs in Figure 3.17 shows that they equalize before they are fully reset high, so that all the PMOS inputs in Figure 3.18 will be partially on during reset. **M1** and **M2** prevent **mirP** from rising unless **inP** is lower than **inM**. Without the cross-coupled NMOS transistors, **mirP** would rise above threshold and pull $\overline{\text{om}}$ down from V_{dd} . **M1** and **M2** don't add much hysteresis because **mirM** and **mirP** drop below V_{th} during reset of the second stage latch (when **inP** and **inM** are reset to V_{dd}). The simulated input-referred hysteresis of the comparator is 1.5mV, with the majority (1.3mV) due to the output latch, because the internal nodes $\overline{\text{om}}$ and $\overline{\text{op}}$ must be driven full-swing to change the state of the latch.

3.6 High Speed Receiver Logic

The outputs from the comparator output latches are full swing CMOS signals, so the subsequent retiming and encoding logic (see Figure 3.19) is high speed but otherwise straightforward. The complementary data from the comparators is received with simple inverters, and the true data passed to the following logic. Dummy inverters match the load on the inverted data, which is discarded. Since the analog comparator circuits presented in

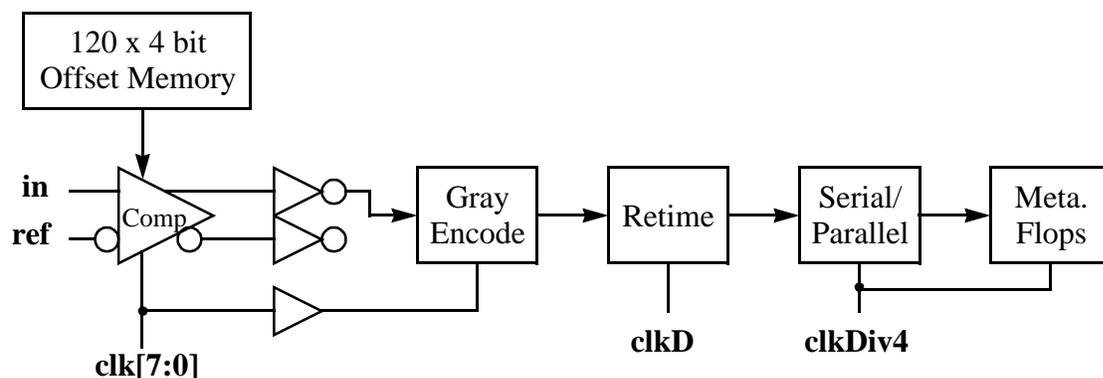


Figure 3.19: High Speed ADC Circuitry

Sections 3.3 through 3.5 have a gain of 4600, assuming a uniformly distributed input, the outputs have a chance of about 1 in 4600 of being metastable, so further metastability

resolution is needed. Because Gray encoding is tolerant of a single metastable input, the data is encoded first to reduce the number of signals. The clock for the Gray encode logic is buffered to reduce the load and thus reduce the jitter on the comparator clocks, and to match some of the delay from **clk** to the comparator outputs. The Gray code outputs are retimed to **clkD**, converted from serial to parallel and relocked to the lower rate **clkDiv4**. Metastability resolution is completed in flip flops at the lower clock rate, where regenerative amplification is more efficient with fewer transfers between latches (transfers reduce the time available for regeneration). The outputs are passed to synthesized logic and to a high-speed test output port.

3.7 Reference Voltage Generation

A resistor ladder is used to generate the reference voltages V_{ref15} , V_{ref14} , ... V_{ref1} for the comparators in the flash ADC. The ladder is tightly coupled to the signal return of the ADC so that noise appears as common mode to the differential comparators. Metal 4 is used for the ladder, shielded with Metal 3. Each ladder step resistor $R=66\Omega$ for a total ladder resistance of 1K. Poly resistors were not used to keep the ADC circuits portable to any digital process (absolute resistor accuracy is not needed). The small mismatches between resistors are corrected by the offset-correction circuitry in the comparators.

An active bias circuit is used to generate the current for the resistor ladder. The active bias circuit allows a high impedance reference voltage input that is filtered to reduce noise, and adjusts the current to compensate for variations in the metal ladder resistance with temperature. A control loop based on a simple op-amp keeps V_{ref} equal to V_{ref0} , and is compensated by the NMOS compensation capacitor at the high impedance output of the op-amp, V_{biasRef} .

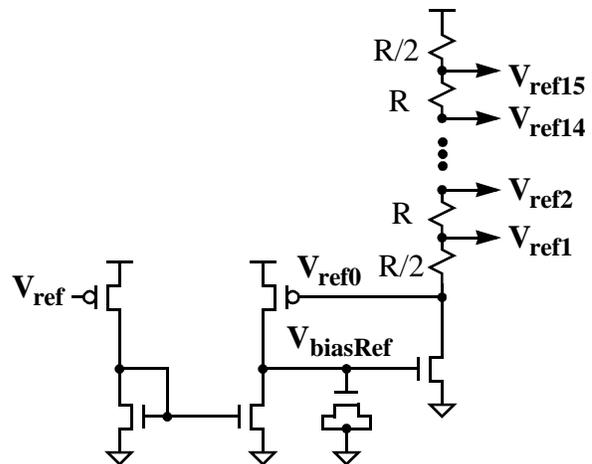


Figure 3.20: Reference Bias Circuit

3.8 Simulated Comparator Performance

The performance of the comparator is summarized in Table 3.1. Simulations were performed in the slow process corner (weak transistors, low voltage, high temperature). The sampling time-constant is measured using a voltage source input, and thus does not include loading effects of the 120 comparators on the 50Ω input.

Description	Value
Sampling time-constant	0.3FO4 (38 ps)
Input Capacitance	$3 C_{\min}^a$ (6 fF)
Sample Period	1 FO4 (125 ps)
Input Range	V_{dd} to $V_{dd}-V_{th}$ (2.5 to 1.75V)
Sample Time Variation Over Input Range	0.1FO4 (12 ps)
Input Referred Hysteresis	1.5mV (0.03 LSB)
Gain at 8 FO4 clock period (1 ns)	5000
Charge coupled back to input	$0.2fCoulomb = 1mV * 200fF$
Offset Before Correction (6 sigma)	121 mV (2.4 LSB)
Offset After Correction	12 mV (0.25 LSB)
Power per comparator, including clocks	4 mW

Table 3.1: Simulated Comparator Performance

a. C_{\min} is the gate capacitance of a minimum sized transistor

Offsets are measured in simulation by adjusting a DC voltage input to the comparator so that its output (after clocking) is metastable. This technique allows the offset of a regenerative latch to be measured, but does not allow small signal analysis. However, the main interest is in operation near the switching point of the comparator, where the comparator can be analyzed as a linear system from input to output at a given clock rate. The input offsets due to each transistor in the comparator are measured and statistically combined (root-sum-squared), with the sample/hold differential pair accounting for most of the input-referred offset. The offsets were estimated in simulation

with a V_{th} mismatch voltage in series with the gates with a standard deviation of $V = \frac{A_{Vt}}{\sqrt{WL}}$ and a width variation with standard deviation $\frac{\Delta W}{W} = \frac{A_{beta}}{\sqrt{WL}}$ where $A_{Vt}=10V\mu m$ and $A_{beta}=0.01\mu m$. Offsets should increase mildly with scaling to finer linewidth CMOS processes [31][32][33], perhaps requiring another bit of digital correction. Offsets are only reduced by a factor of 10 with 4 bit offset correction because the offset correction DAC has mismatch errors that reduce its effective resolution.

3.9 Comparator Array Layout

Two goals drive the layout of the comparator array: low input capacitance and low noise coupling. An isolated supply is used for the comparator array, which contains all of the analog circuitry in the ADC. To keep the array small and the wires short, data is driven out of the array immediately after amplification to full swing. Clocks are distributed from right to left, to partially cancel the sample time variation with reference level that

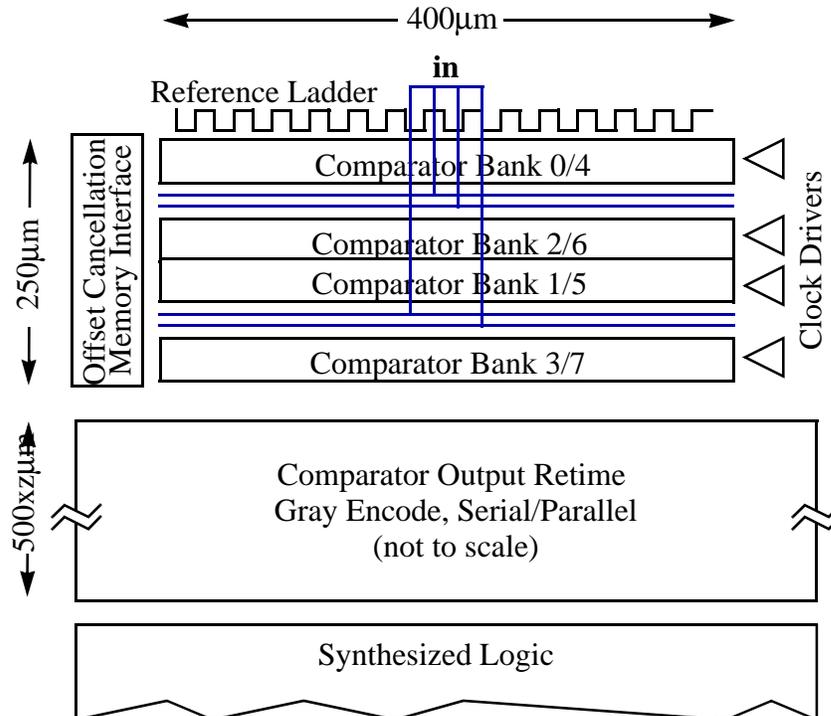
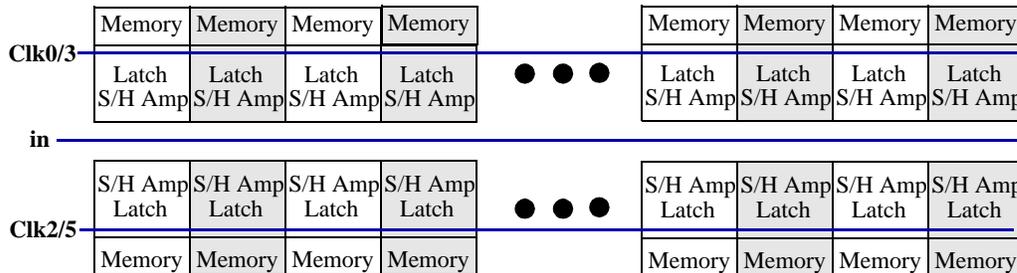


Figure 3.21: Comparator Array Layout

increases from left to right. The metal 4 reference ladder runs serpentine between the comparators and the input pads. Comparators with complementary clocks are interleaved (as shown in Figure 3.22), sharing the same input, reference and supply wires so that charge kickback and supply noise are cancelled to first order.



Tile 2 of these vertically, total size each 3000x400 lambda

1 inch=200 Lambda

Figure 3.22: Interleaved Comparator Bank Layout

The offset cancellation memory cells are interspersed with the comparators to reduce the number of wires run into the array. The offset memory, decode logic and DAC current sources approximately double the area of the comparator array. The bit lines for the memory run out to the left of the array, and word lines are input from the synthesized logic at the bottom of the comparator array (see Figure 3.21). Along with the received data that flows from top to bottom through the various stages, these wires use up nearly all available wiring channels, but the layout is still limited by local interconnect and transistors. The dense layout minimizes wire capacitances, so that most nodes are dominated by transistor parasitics and the performance of the ADC is primarily determined by the circuit design.

Both the input and the reference ladder are shielded in U-shaped troughs (see Figure 3.23) to tightly couple them to the signal return for good noise rejection. The weak links in this approach are the bond wires, which are susceptible to inductive coupling, and allow the signal return and the reference ladder to move independently in response to noise and supply bounce.

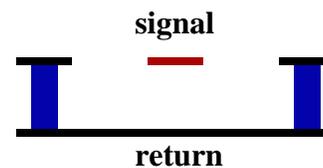


Figure 3.23: Shield Troughs

3.10 Summary

This chapter presents the design of a receiver based on a 4-bit time-interleaved flash ADC. The primary goals of high sampling rate and bandwidth are met with a high bandwidth, low-gain sample/hold amplifier. With many comparators loading the analog input due to the time-interleaved flash architecture, small transistors are used to minimize input capacitance. The resultant large offsets are reduced with digital calibration in the second stage latch. The small, low-gain sample/hold amplifier also results in thermal noise of 0.4 LSB, showing that noise analysis is increasingly important in high performance CMOS design.

Subtraction of the input from the reference before sampling improves linearity. By using switchable PMOS loads, the sample/hold amplifier combines reference subtraction and input sampling to achieve high sampling bandwidth. The performance of the sample/hold amplifier is preserved by performing offset correction in the second stage latch. To ensure gain that is adequate to prevent the offsets and hysteresis of the subsequent CMOS latches from dominating performance, an output latch is used that efficiently captures the output from the second stage latch.

The comparator performance demonstrates that CMOS circuits can achieve high bandwidth with good linearity. However, with the capacitive load of the full comparator array, the input sampling time-constant is well above the goal of $0.3 \cdot FO_4$. The next chapter presents the remainder of the transceiver, shows that the transmitter also suffers from capacitive loading, and explains how inductors are used to distribute parasitic receiver and transmitter capacitances to improve bandwidth.

Chapter 4

Transceiver Chip Design

This chapter first presents the architecture of the full transceiver, highlighting issues critical for high speed operation. It provides an overview of the transmitter and clock generation circuitry (designed by colleagues), and shows how to extend transceiver bandwidth with an inductive network. The logic and memories that complete the transceiver are presented, followed by a discussion of the communication techniques that are supported, setting up the experimental results in Chapter 5.

4.1 Transceiver Architecture

The transceiver architecture minimizes hardware complexity by pushing most signal processing into software. In addition to simplifying the chip, this approach allows a range of modulation, equalization and coding techniques to be explored. Because it is not possible to transfer the full bandwidth of the test link to a PC for software processing, on-chip memories are used to store raw link data, as shown in Figure 4.1. The transmitter converts 8-bit digital data from the transmit memory to an analog waveform, sending periodic sequences up to 1024 symbols long. The receiver either stores or compares the received data to the contents of the 1024 symbol receive memory. The PLLs generate clocks for the transceiver from a reference clock. Raw transmit data sequences are generated by a PC and downloaded into on-chip memory (dashed blocks are implemented off-chip in software).

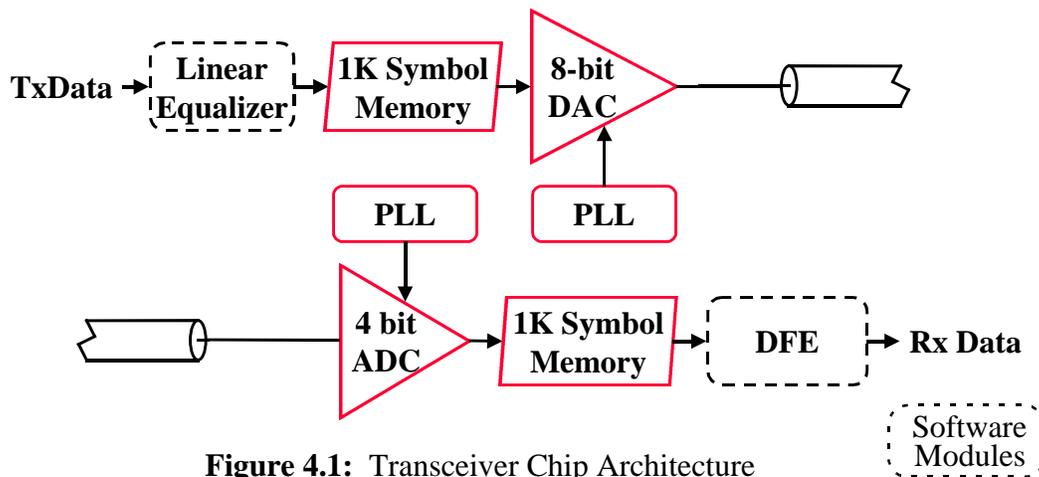


Figure 4.1: Transceiver Chip Architecture

Decision Feedback Equalization (DFE) can be performed off-chip as a post-process on 1024-symbol snapshots of the received data. The transceiver has a symbol time goal of one FO4 delay, which corresponds to 8Gsamples/sec in the 0.25 μ m CMOS process used to fabricate the prototype chips (slow process, voltage, temperature). The ADC and DAC each need Nyquist bandwidth of 4GHz (half the sample rate) to prevent the circuits from limiting the link performance. The receiver uses 2 bits of resolution to support 4-level signalling, and 2 more bits to allow a DFE as a post-process. The transmitter also needs 2 bits of resolution for 4-level signalling, 4 more bits for equalization and interference correction, and 2 additional bits to explore the limits of high-speed DAC resolution.

Although the circuits are differential, singled-ended input and output signals are used to allow transmission over inexpensive coaxial cables. In this way, noise that couples equally to the reference ladder and the input appears as common mode to the differential comparators.

4.2 Transmit DAC Design

The 8-bit output stage was designed by Siamak Modjtahedi and Ken Yang, extending Ken's binary output stage [3]. Time interleaving is used to achieve a high sample rate, as shown in Figure 4.2. The transmitters consist of 8 time-interleaved DACs running with a clock period of 8 FO4s. Each of the 8 interleaved DACs generates a narrow

current pulse for one symbol time, enabled by the overlap of two 50% duty-cycle clocks (see Section 2.3). The output pulses of the 8 interleaved DACs are summed onto the 25Ω impedance formed by the 50Ω on-chip termination in parallel with the 50Ω transmission line.

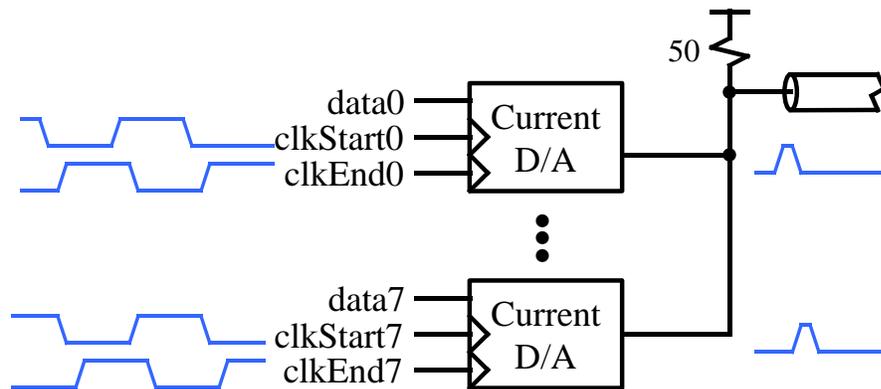


Figure 4.2: Time Interleaved Transmitter

To achieve multiple bits of resolution, each interleaved DAC consists of many small current sources, as shown in Figure 4.3. The least significant bits are implemented with 5 binary weighted current sources, and the upper 3 bits are implemented with 7 identical current sources to reduce nonlinearity due to transistor mismatches. The decrease in DAC output resistance with output current as more current sources are switched on causes integral non-linearity in the DAC, resulting in slightly smaller step sizes for low output levels. The DACs are designed with extra resolution to digitally correct for this nonlinearity.

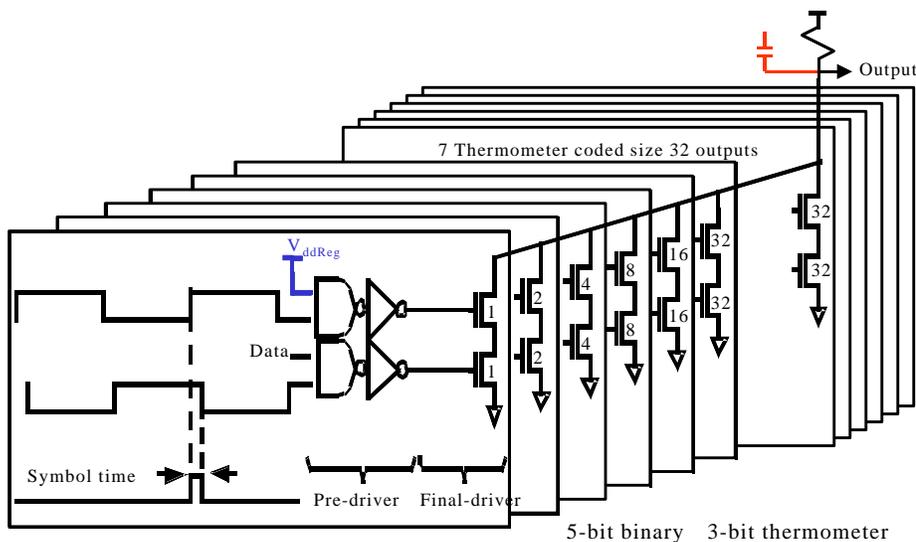


Figure 4.3: Interleaved DAC Design

Each of the current sources consists of 2 NMOS transistors in series with the bottom source connected to ground. The output level is controlled with a simple linear regulator on the V_{ddReg} supply of the predrivers. This maximizes the gate overdrive voltage, allowing smaller transistors to be used. The fanout of the predrivers is kept low to provide a fast output time constant of 35 ps, simulated for a single DAC with minimal wire capacitance. However, because of the 8 way interleaving, the RC pole at the DAC output limits the bandwidth of the transmitter.

The **parasitic capacitance** shown at the DAC **output** in Figure 4.3 is estimated from layout as 0.9pF of drain capacitance from output transistors, 0.5pF of pad capacitance, and 3.0pF from the long, wide wires that run to the output pads. The wire capacitance is large because 10 μ m-wide wiring is needed to carry the 40mA output current, despite the thick top layer 5 metal with copper content and high (4mA/um) electromigration current density limit. With the benefit of hindsight, the wiring capacitance could be reduced by a factor of more than 2 by moving transistors closer to the pads and tapering down wires. However, the output RC is 4.4pF * 25 Ω = 110ps, and limits the output time constant of the transmitter to 114ps=0.9FO4 (cascaded time constants add in root-sum-squared fashion [30]). This is well above the design goal of 0.3FO4, and results in a mismatched source termination.

4.3 Inductors to Distribute Parasitic Capacitance

To increase the bandwidth of the transmitter with its large output capacitance, an old, good idea is used that for years has been incorporated into distributed amplifiers for

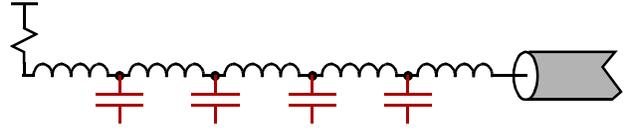


Figure 4.4: Lumped 50 Ohm Transmission Line

oscilloscopes [42]. A lumped 50Ω transmission line is constructed as shown in Figure 4.4 by dividing up the large DAC parasitic output capacitance, and inserting inductors so that $\sqrt{\frac{L}{C}} = 50\Omega$.

Pairs of transmitters with complementary clocks are wired to separate pads, with optional bond wire inductors between the pads, as shown in Figure 4.5. Each point on the LC line still sees a 25Ω impedance, 50Ω to the right, and 50Ω to the left. Thus, the current pulses from each transmitter develop into voltage pulses that propagate in both directions. Since the line is terminated at both ends, the pulses don't reflect (much).

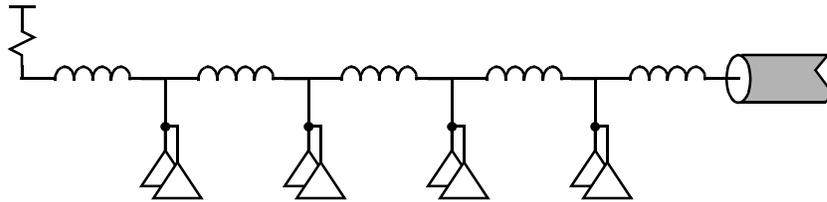


Figure 4.5: DAC Output Inductors

Each transmitter outputs a pulse that appears at a different time on the line, so transmitters on the left must transmit early to compensate for the delay down the LC line. Because each DAC has its own clock, the LC delay along the lumped transmission line is compensated by adjusting the phase of the clocks. The same ability to adjust clock phases is needed to cancel static phase errors (see Section 2.2: “Clock Generation” on page 10).

However, interference between pulses can occur, because DAC pulses propagate in both directions on the LC line. An earlier pulse from a transmitter on the right can propagate left and add to a pulse from a transmitter on the left, resulting in a total voltage that takes the NMOS output transistors out of saturation. While little interference is seen in the current chip because a reduced output swing of 750mV is used, future “distributed”

transmitters may need to address this issue. One possible approach is to code the data for transmitters that are separated by a roundtrip LC delay equal to their output time spacing, so that the sum of the transmitter outputs is less than full swing.

A simulation of the LC circuit in Figure 4.4 with a model of the package and connector evaluates the sensitivity to inductance value. Figure 4.6a shows simulated DAC output bandwidth (not including output circuit bandwidth) with inductance values near $L_o=50\Omega^2*4.4\text{pF}/4=2.7\text{nH}$, and at $1.3L_o$ and $0.7L_o$. The results above 3 GHz show an average improvement of 6 dB with inductors (resulting in double the output voltage), and are insensitive to variations in inductance of +/- 30%. The peaking at higher frequencies is reduced by low-pass filters formed by circuit and wire parasitics that are not included in this simulation.

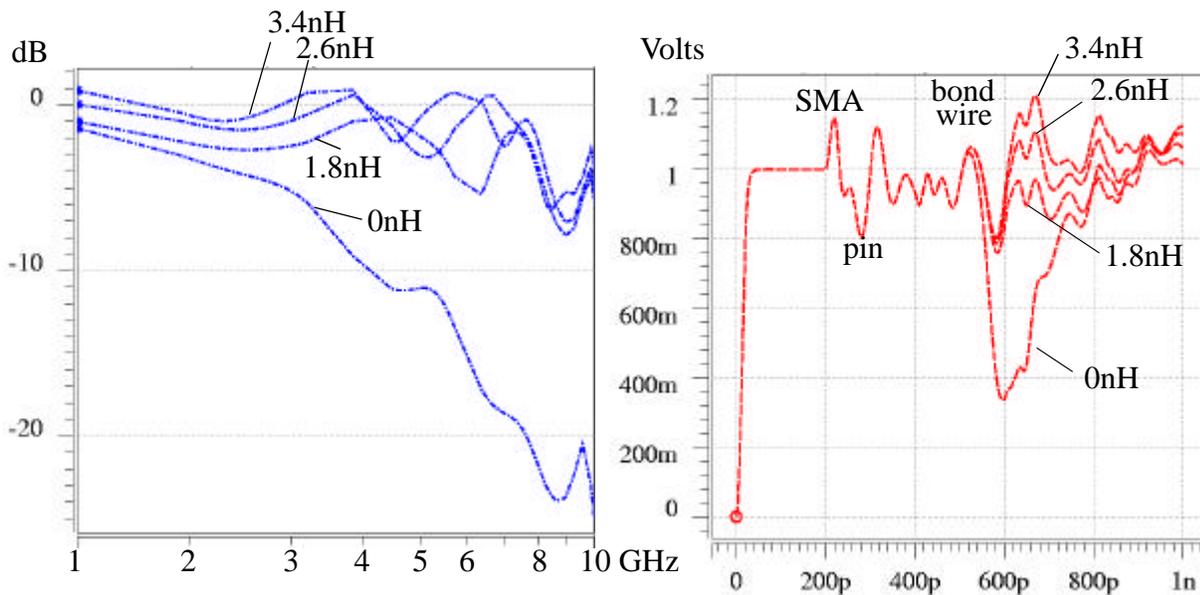


Figure 4.6: a) Simulated DAC Output Bandwidth b) Simulated TDR of DAC Output

To help experimentally verify inductance values, Figure 4.6b shows simulated Time Domain Reflectometer (TDR) traces on the LC network with the same values. The upper trace shows the step from the TDR, with a positive reflection indicating a high impedance at the SMA connector, then a negative reflection from the package pin capacitance. The signal then propagates along a transmission line that models the controlled impedance trace in the package, and then has a peak from the inductive bond wire. Without inductors, the 4.4pF input capacitance causes a large dip, while the LC network results in a series of smaller dips and peaks.

A similar but less severe input bandwidth problem exists in the receiver. The input load from the 120 comparators in the 8 time-interleaved ADCs adds up to 720 fF, and along with wire and pad capacitance, forms an RC pole with a time constant of $25\Omega * 1.8\text{pF} = 45\text{ps} \cong 0.36\text{FO}_4$. When combined with the comparator circuit sampling time constant of 0.33FO_4 , the design goal of $\tau = 0.3\text{FO}_4$ is missed by 60%, despite the small devices used in the comparators. The input load could be lightened by reducing the number of comparators, which would require reducing the sample rate or precision of the ADC, but the use of inductors allows delay to be traded for high bandwidth instead.

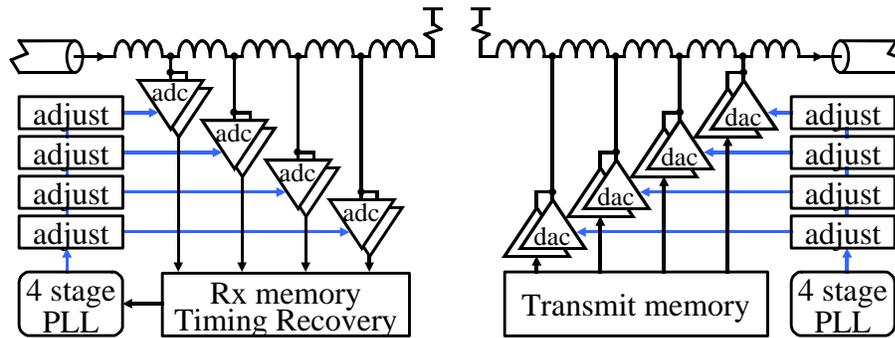


Figure 4.7: Transceiver Block Diagram with Inductors

Thus, as in the transmitter, pairs of time-interleaved ADC inputs are wired to separate pads as shown on the left in Figure 4.7. This allows bond wire inductors to optionally be inserted between the input capacitances to evaluate the ADC performance with and without the inductors (the package to die bond wire forms the leftmost inductor). The **clocks** for each time-interleaved ADC are phase adjusted to compensate for the LC delay along the lumped transmission line.

Figure 4.8a shows simulation results of a model of the ADC input network with an average 4 dB reduction in loss above 3 GHz that is insensitive to $\pm 30\%$ variations in inductance. The dips and peaks in the frequency response disappear when the package model is removed from the simulation, indicating that they are caused by secondary reflections between the package and the ADC input capacitance. The simulated TDR results in Figure 4.6b can be compared to measured TDR results to experimentally verify ADC input inductance values.

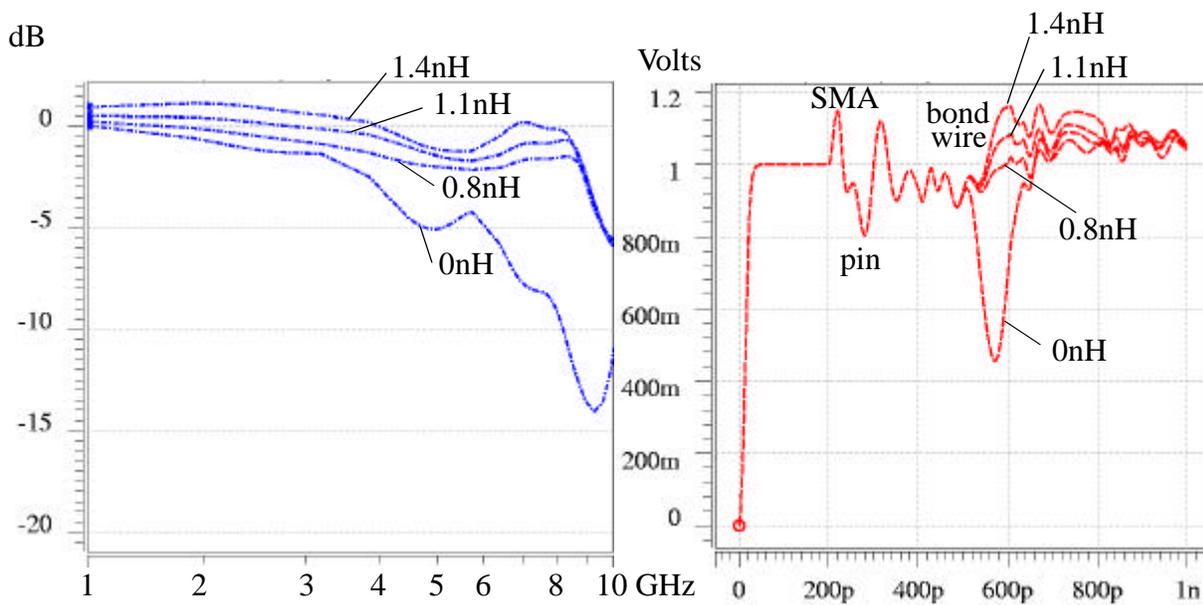


Figure 4.8: a) Simulated ADC Input Bandwidth b) Simulated TDR of ADC input

Bond wire inductors are used for their low parasitics¹ and for flexibility in testing, and can be controlled with a tolerance of about 20%. While bond wire inductors are not practical for volume production, similar inductors can be constructed in flip-chip packages, with the signal going up one solder ball from the die to a spiral trace in the package and back down another solder ball to the die. Inductance is roughly proportional to length, with the inductance of a loop of diameter D and cross-sectional diameter d (a typical bond wire has $d=18\mu\text{m}$) given by [40]:

$$L/\text{length} = 0.57 * \log_{10}(8D/d-6.4) \text{ nH/mm, for } D > 5d \quad [2.10]$$

In fact, the rule of thumb $L/\text{length} = 1\text{nH/mm}$ is surprisingly accurate, within 20% for a loop with $20d > D > 4d$, or for a wire a distance H above a plane with $100d > H > 14d$. Inductance is also insensitive to variations in loop shape, with the inductance of a square loop 16% larger, and an octagon 4% larger than the inductance of a circle of the same perimeter [37]. Coupling between the inductive loops is harder to calculate, but the overall inductance values can be verified with a TDR.

1. On-chip spiral inductors are not used, because their series resistance causes significant attenuation.

4.4 Clock Generation

The demands of providing stable and accurate multi-phase clocks are met with an enclosed-dual-loop clock generator, shown in Figure 4.9. The **inner loop** phase-locks a 4-stage VCO to a one-quarter rate reference clock, **refClk**, with high loop bandwidth to minimize random and supply noise. The **outer loop** (only used in the receiver) aligns the ADC clocks to the receive signal, with lower bandwidth digital filters to reduce phase noise. The clock generator is based on the dual-loop architecture and circuits described in Section 2.2, with an additional Feedback Phase Adjuster to improve phase-locking of the multi-phase clock outputs to the receive signal. By adjusting the Feedback Phase Adjuster, the multi-phase clock outputs of the VCO are shifted in phase relative to **refClk** to track the received signal. Thus, the Static Phase Adjuster are only responsible for correcting static timing errors, and only need enough range to correct for clock wiring and transistor mismatches. Changing the feedback phase adjuster shifts all of the multi-phase ADC clocks evenly, maintaining the corrected phase spacing.

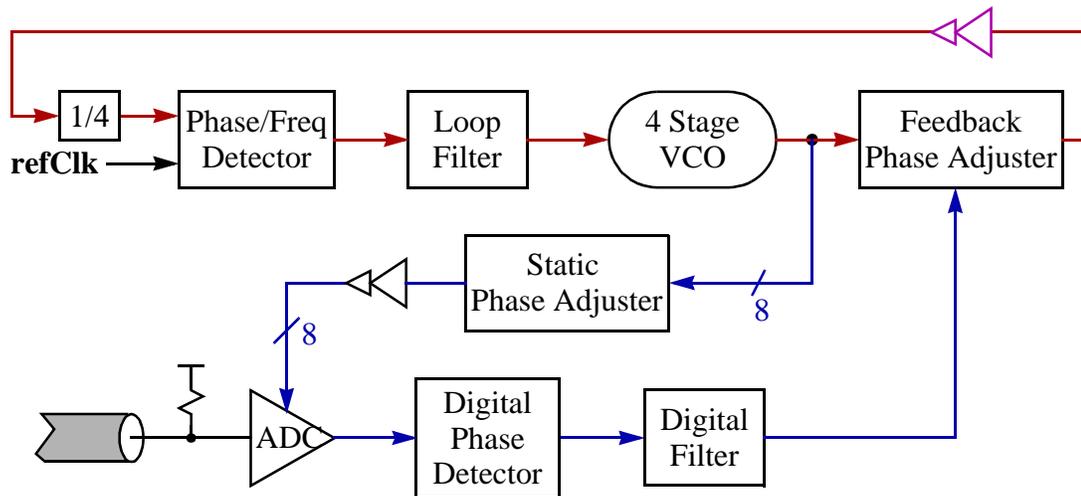


Figure 4.9: Enclosed Dual Loop Clock Generation

The digital outer loop waits for the inner loop to settle after each phase adjustment before measuring the new phase of the received waveform and making a new adjustment. This minimizes the dithering due to the digital “bang-bang” control, so that the loop phase alternates between two states in the absence of noise. The outer loop tracks noise within the slew rate capabilities determined by its adjustment rate, and uses digital filtering to

eliminate high frequency noise. In addition to being limited by the inner loop settling time, the adjustment rate of the outer loop is limited by logic delays, since it is implemented in synthesized Verilog logic, running at the 250 MHz **clkDiv4**. The logic delays total about 10 **clkDiv4** cycles for receiver data encoding, parallelization, metastability resolution, phase detection and filtering. Adding a few inner loop time constants (about 3 **clkDiv4** cycles each) and some margin, the outer loop makes one adjustment every 32 **clkDiv4** periods for an 8 MHz adjustment rate. Thus, small amplitude phase noise below about 8 MHz can be tracked by the loop. Larger amplitude noise has a larger time constant, since multiple adjustments are needed to track the noise. Thus, the outer loop is suited to tracking low frequency $1/f^2$ noise, where the noise amplitude is large at low frequencies. A frequency offset of about 100 parts per million can also be tracked.

The relative phase of the clocks is corrected by the Static Phase Adjust block within approximately 1/20th of a symbol time. When the outer loop changes the control for the Feedback Phase Adjust, all of the clock phases smoothly slew to the new setting with the same relative phase. Extra feedback phase adjustment precision is obtained with a state machine that dithers between two feedback phase settings. Since the dithering occurs at a rate above the inner loop bandwidth, the loop settles at an intermediate phase.

A previous chip demonstrated clocks stable enough to receive data without a high bandwidth timing recovery loop by including **replica clock buffers** (see Figure 4.9 on page 51) in the feedback path of the inner loop [6]. The replica clock buffers match the supply-noise induced phase errors in the clock buffers that follow the Static Phase Adjust block. These phase errors are compensated by feedback around the inner loop, within its bandwidth. The cost of this stability is increased delay in the feedback path of the inner loop, which reduces the phase margin. A significant experimental benefit is that clocks can be stable enough to run clock recovery experiments with phase control loops in off-chip software, by injecting phase noise at frequencies low enough to be tracked by these loops.

4.4.1 Digital Phase Detector

The receive digital phase detector is the only element in the timing recovery loop that is affected by multi-level signalling or complex equalization techniques. The Digital Phase Detector in Figure 4.9 on page 51 has several modes of operation that work with samples of the receive signal taken either at symbol transition times, or at symbol data sample times. Transition timing samples are taken when the received waveform has high slope, so voltage noise maps to small phase noise. Data timing samples have the advantage of using the same clock to sample data and timing, avoiding a source of static phase error.

In its simplest, traditional mode of operation, the phase detector samples at symbol transition times, and an extra sampler is needed in addition to the data samplers. To provide this extra sampler, the 7th of the eight time-interleaved ADCs has its clock shifted by half a bit time, so it samples at symbol transition times, and the 7th data symbol is not sampled. The transmit data is coded so that the 7th symbol is the complement of the 6th symbol, so the received waveform crosses the middle of the voltage swing halfway between the 6th and 7th symbol sample times. Thus, a single comparator at the transition between the 6th and 7th symbol can sense whether the sampling clock is early or late, using the 6th symbol to determine if a rising or falling edge was transmitted, and assuming the 7th symbol is the complement of the 6th symbol. This algorithm only reduces the data rate by 1/8th, and provides a timing sample every **clk** cycle (every 8 symbols).

In its other, data sampling modes of operation, the phase detector uses extra resolution in the ADC to measure the phase of the received waveform. Two algorithms are supported to determine whether the sampling clock is early or late, which span the range of possible and published [55][56] approaches. To illustrate the design space, consider the transition types in multi-level transmitter-equalized signaling¹ shown in Figure 4.10.

1. Transmitter equalization results in well-formed data eyes at the receiver (see [10][59][60] or [61]).

The two transitions around a data symbol are labelled with “+” for a rising edge, “-” for a falling edge, and “0” for no edge. The first case, labeled “++”, yields reliable timing information that indicates the sample clock is early if the received signal is lower than the ideal voltage. The second case, labeled “+0”, indicates possible early timing with a lower than ideal received signal, but late timing can result in either a high or a low. The third and fifth cases, labelled “+-” and “00”, do not yield timing information. The fourth case, labelled “0+”, indicates possible late timing with a higher than ideal received signal, but early timing can result in either a high or a low. The bottom four cases are symmetric to the top four.

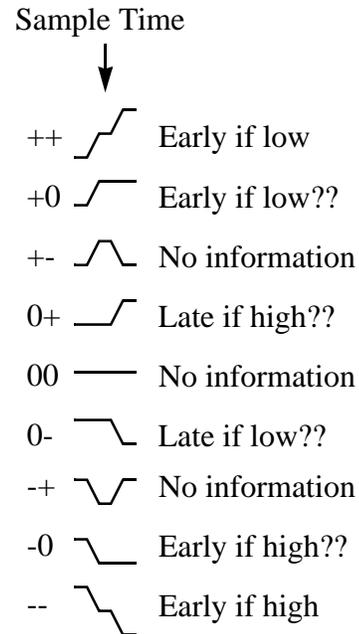


Figure 4.10: Transition Types

The first data sampling algorithm uses only the most reliable timing information, from “++” and “--” transitions, and thus is named the “Rising-Rising” algorithm. A single comparator at the ideal received voltage is used to determine if the sample clock is early or late. Timing information is obtained only when the adjacent data samples indicate a “++” or “--” transition. With 4-PAM signaling, these transitions occur at 1/8th of randomly distributed symbols. Because the phase detection logic runs at the **clkDiv4** rate, 32 symbols are processed at a time. To simplify the logic, only four timing samples are taken from the 32 symbols, so that a useful transition occurs on average in half of the **clkDiv4** cycles.

In contrast, the second data sampling algorithm uses every possible timing sample, expanding on a stochastic binary algorithm in [56]. The algorithm uses questionable information (for example, the “Early if low??” information from a “+0” transition), counting on statistical averaging to filter out random noise if the clock is late. With equal distribution of transition types, and adequate filtering of noise, this algorithm locks to the

desired phase. However, simulation shows that the required filtering reduces the bandwidth of the loop significantly, so that the first data sampling algorithm outperforms this approach despite using information from fewer transitions. Again, to simplify the logic, only four timing samples are taken every **clkDiv4** cycle.

Phase detectors for signals that are equalized after ADC sampling (using a DFE implemented in software as described in Section 4.6.4) can also be implemented. The slope of the waveform can be determined from the adjacent received data symbols, since the channel response is known at the receiver. These more complex algorithms can be implemented in software, and tested with very low frequency phase noise injected into the system that is below the software loop bandwidth, as long as data can be received reliably without a phase recovery loop running in hardware, as demonstrated in [6].

Two levels of filtering are performed on the early/late data from the phase detector. First, the four early/late bits are checked for unanimous early or late, and discarded if discord exists (this may discard useful information in some cases, but discord generally indicates that the current phase is near optimal). Second, unanimous early and late votes are counted¹ over multiple **clkDiv4** cycles. The feedback phase is adjusted up if the count reaches a (programmable) threshold, and down if the count reaches the negative of the threshold. Once an adjustment is made, a wait counter is loaded, and subsequent early or late votes are discarded until the counter expires to indicate that the inner loop has settled and that receive data with the new clock phases has been sampled by the ADC, coded, parallelized, resolved, phase detected and filtered (see Figure 4.9 on page 51).

4.5 Synthesized Logic and Memories

To complete the transceiver, non-critical functions are implemented with Verilog code that is synthesized into CMOS logic blocks. The logic and embedded memories allow a PC to generate raw transmit data sequences and download them into on-chip memory, and to upload and process raw receive data sequences from onchip memory. The

1. The counter has two modes: pure up/down counting until the (programmable) positive or negative threshold is reached, or reset-on-reverse that zeros the count when the early/late indication reverses.

PC also controls, configures, and monitors the transceiver via a serial interface. To simplify the place and route, the synthesized logic is partitioned into three blocks, clocked by the receive **rxclkDiv4**, the transmit **txclkDiv4**, and the serial PC interface **scclk**, respectively.

The receive logic either stores or compares the ADC receive data to the contents of an onchip 1024 symbol memory. Because the received data is parallelized into 128-bit words to support the 32 Gbit/s data rate, the receive logic block is large. The logic also interfaces to the offset correction memory in the ADC comparator array, and provides control and status

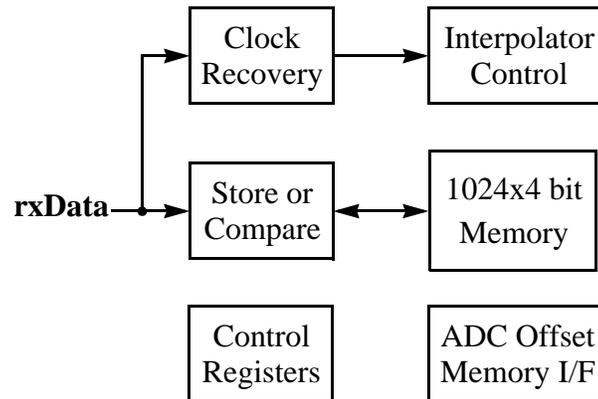


Figure 4.11: Synthesized Receive Logic

registers to configure and monitor the receiver. The digitally controlled clock interpolators (see Section 2.2: “Clock Generation” on page 10) are also controlled by the synthesized logic, supporting a range of clock recovery algorithms described in Section 4.4.1.

The transmit logic is simpler, primarily providing an interface to the 1024-symbol transmit memory that continuously transmits a periodic data pattern. Still, a large die area is consumed to support the 256-bit wide data path. Transmitter control and status registers are also implemented, along with static control registers for the transmit phase interpolators.

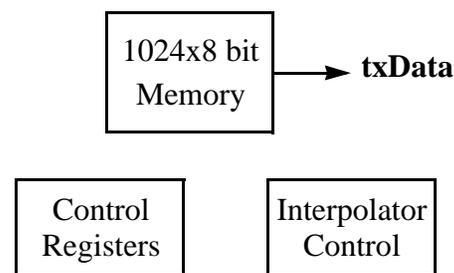


Figure 4.12: Transmit Logic

4.6 Transceiver Capabilities

The emphasis on off-line data processing allows the transceiver to be used as a tool to research a range of communication techniques. The on-chip bit error rate test capability allows comparisons to be made at error rates below 10^{-12} . The testing is performed by comparing one or more received bits against the contents of the 1024 symbol memory, allowing 2-PAM, 4-PAM, 8-PAM or 16-PAM modulation. Because the data is checked against memory data instead of a fixed sequence, a variety of coding schemes from straight NRZ to 8B10B or partial response (e.g. 1+D) signalling can be verified.

The application of several communication techniques is discussed in the final sections of this chapter as a prelude to the experimental results in Chapter 5. First, the correction of interference from a variety of sources is discussed. Then, an analysis of multi-level modulation is presented, followed by a discussion of linear equalization. Finally, decision feedback equalization is evaluated, although it can only be performed as a post-process in software.

4.6.1 Interference Correction

Because extra resolution can be designed into the transmitter without reducing performance, it is natural to use this resolution to correct for interference throughout the transceiver. Interference correction is particularly important at high bit rates, where signal levels are smaller and interference is larger¹.

Reflections from impedance discontinuities cause linear interference that changes the frequency response of the channel. These changes can be corrected with the same equalizer that corrects for losses in the channel. In addition, additive noise from clocks and other known signals can be corrected with extra resolution in the transmitter. Finally, the nonlinearities in the transmitter can be corrected to first order.

In this work, interference correction is performed in software for flexibility. Software generates the (possibly coded) desired data sequence, equalizes, corrects for interference, predistorts to compensate for transmitter nonlinearity, and stores the data in

1. For example, capacitive coupling and reflections are more of a problem with faster edge rates.

the onchip transmit memory, as shown in Figure 4.13. Because the performance of the time interleaved transmitters varies, different equalizer, interference correction and predistortion weights are used for each transmitter. In a hardware implementation, these functions would be implemented as lookup tables, running at the transmit **clk** or the **clkDiv4** rat. Individual lookup tables for each transmitter do not increase complexity as multiple tables are needed to handle the high data rate. An onchip state machine continuously reads the 8-bit values from memory and transmits them.

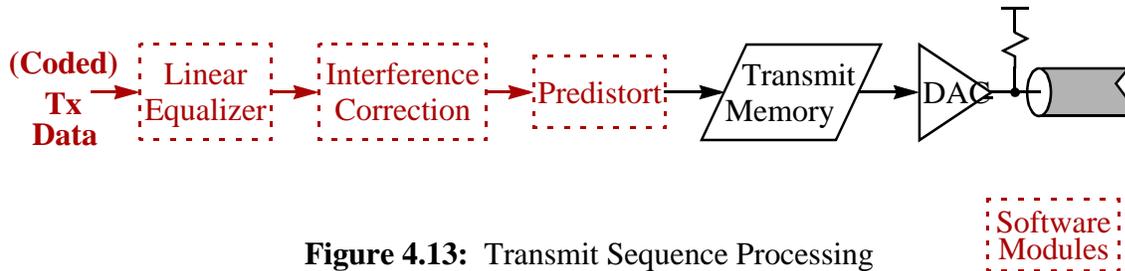


Figure 4.13: Transmit Sequence Processing

4.6.2 Multi-level Modulation

Figure 4.14 compares eye diagrams of equalized 2-PAM and 4-PAM signaling at 10 Gbit/sec on an 18m RG59 cable (simulated using measured cable step response data with no PCB board or IC package). Although the 4-PAM modulation runs at half the symbol rate, the more complex transitions result in a phase margin similar to that of 2-PAM modulation. The 2-PAM signaling has higher bandwidth, so it suffers more attenuation in the cable. The 4-PAM signaling is less attenuated, but 3 eyes are formed, so that the voltage margin (eye height) is similar to that of the more attenuated, single 2-PAM eye.



Figure 4.14: 2-PAM and 4-PAM Modulation on 18m RG59 cable (6 tap equalizers)

The tradeoff between 2-PAM and 4-PAM depends on how the attenuation varies with frequency, as is shown in Figure 4.15, which compares the attenuation of 2-PAM and 4-PAM eye heights in two idealized wires. The [upper right curves](#) show the attenuation of

the eyes on a wire with only skin-effect limited conductor loss, where the loss in dB increases with the square root of the signalling frequency. The **lower left curves** show the attenuation of the eyes on a wire with only dielectric loss, where the losses in dB increase directly with frequency. The solid curves are for 2-PAM, while the dashed curves are for a 4-PAM system. The 2-PAM and 4-PAM curves for each of the two wires are identical, but the 4-PAM curves are shifted down by 9.5 dB (a factor of 3 to account for the 3 eyes in 4-PAM modulation), and shifted to the right by a factor of 2 in frequency (because the 4-PAM *symbol* rate is half that of 2-PAM with equal *bit* rate).

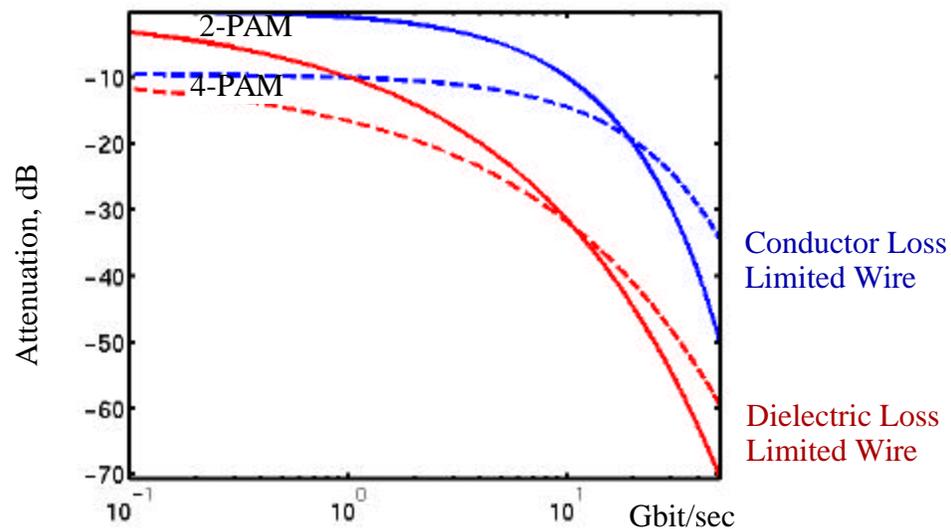


Figure 4.15: 2-PAM and 4-PAM Wire Attenuation vs. Frequency

The crossover points of the curves in this simple analysis indicate the frequencies at which 2-PAM and 4-PAM eye heights are equal. The crossover points occur well beyond the 3 dB bandwidth of the wires, at 19 dB for the conductor loss limited wire, and at 32 dB for the dielectric loss limited wire. Real wires fall between these two extremes. While this analysis makes 4-PAM look unattractive for most practical wires, impairments due to circuit and package parasitics, crosstalk and thermal noise are worse at the higher frequencies needed for 2-PAM (receiver offsets remain roughly unchanged). These impairments make multi-level modulation more attractive than this analysis shows. In the chip in this work, 4-PAM signalling has a larger eye height than 2-PAM at 8 Gbit/s even without any wire losses, due to the 119 ps overall time constant of the transceiver.

4.6.3 Linear Equalization

For either 2-PAM or 4-PAM systems, equalization is needed to correct for the wire and parasitic losses in high-speed CMOS links. These equalizers correct for the frequency-dependent attenuation from the wires and parasitics. A number of equalizer architectures are possible, with the linear equalizer the simplest and the most common [7][10][61]. Since most frequency dependent losses are low pass, most equalizers are high pass filters, to provide an overall channel that is flat, as shown in Figure 4.16. This extends the usable bandwidth of wires and circuits, and allows higher link data rates.

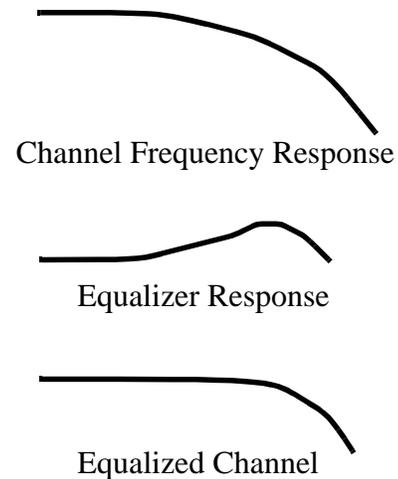


Figure 4.16: Frequency Response of Channel and Equalizer

Equalizers are often based on Finite Impulse Response (FIR) filters, which can form arbitrary linear functions, are easy to implement in CMOS, and can be automatically adjusted. A symbol-spaced FIR filter adds versions of a signal delayed by symbol times, with adjustable weights (gains) on each delayed version, as shown in Figure 4.17. A 5-tap

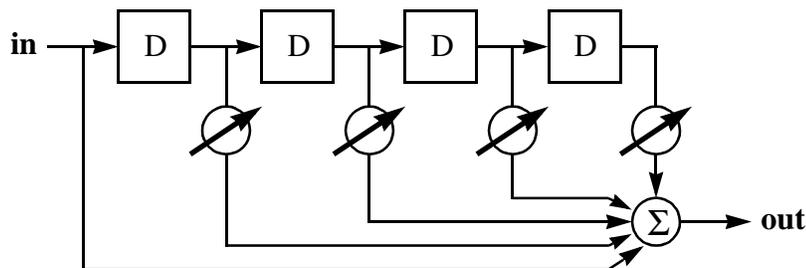


Figure 4.17: 5-tap FIR filter

FIR filter has 4 degrees of freedom (the weights), and can generally¹ correct the interference from four adjacent symbols at each symbol sample time. Because linear

1. The pulse response (ISI) is inverted in the equalizer, with each term in the inverse to first order negating the corresponding term in the pulse response. The inverse also must correct for the second order pulse response from the first order correction, and so on. For a typical pulse response that drops rapidly toward zero, the second order corrections are much smaller than the first order, and the significant terms in the inverse correspond to the significant terms in the pulse response.

elements can be transposed, the equalizer can be put either before the wire at the transmitter, or after the wire in the receiver, with similar results.

Linear receive equalization is simple from a system standpoint, since the required feedback from the receiver to adjust the equalizer (dashed line in Figure 4.18) is easily available. Unfortunately, it requires signals to be sampled, held for multiple symbol times, and added. Analog sample and hold circuits and adders are difficult to implement at high speeds, because they need time to reset and settle to final values. A digital implementation requires a more complex receiver with an ADC with extra bits of resolution to keep the quantization errors small.

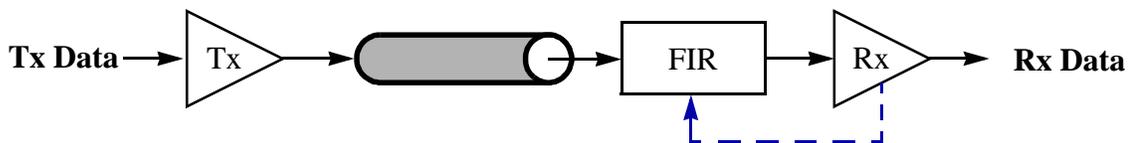


Figure 4.18: Receive Equalization

Implementing a linear equalizer circuit in the transmitter is easier since the data being filtered starts as binary data, making the needed multiplications trivial. There are two common implementations: DACs with adjustable weights and summed analog outputs [7], or digital addition of adjustable values convolved with the data sequence [10]. The digital implementation (used in this work) also allows nonlinear equalization and cancellation of other interference: a 5-tap linear equalizer is implemented by convolving 5 digital data symbols with adjustable digital weights. To avoid the need for high speed multipliers and adders to perform convolution, table lookups can be substituted. Only the final result needs full resolution to keep quantization errors small. This result is transmitted by a DAC. Transmit equalization does require feedback from the receiver (dashed line) to adjust the equalizer at initialization, and periodically if high accuracy is needed.

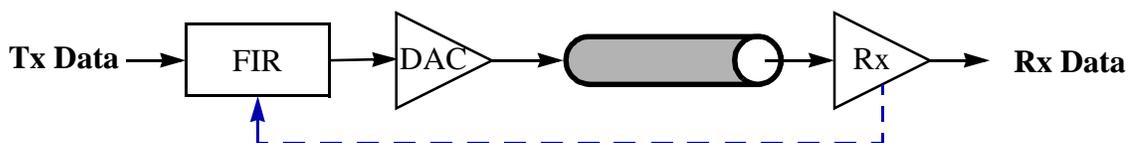


Figure 4.19: Transmit Equalization

Linear equalizers reduce the Signal-to-Noise Ratio (SNR) by either attenuating the signal or amplifying noise. Transmit equalizers attenuate low frequency signals to match the high frequency wire and parasitic losses, while receive equalizers amplify high frequency noise. In either case, the reduction in SNR is the same for noise that is coupled after the transmitter and before the receiver.

4.6.4 Decision Feedback Equalization

To avoid degrading the SNR, a Decision Feedback Equalizer (DFE) uses received data that has been fully resolved to digital symbols, and thus contains no noise, as shown in Figure 4.20. Since the digitizer in the receiver (Rx) is a nonlinear element, a DFE is not a linear equalizer, but it equalizes in a similar manner, by adding fractions of previous symbols to the analog received signal. The transmitter outputs simple, full swing data, with the largest possible swing each bit-time at the receiver. Feedback from the receiver is again used to adjust the DFE (dashed arrow).

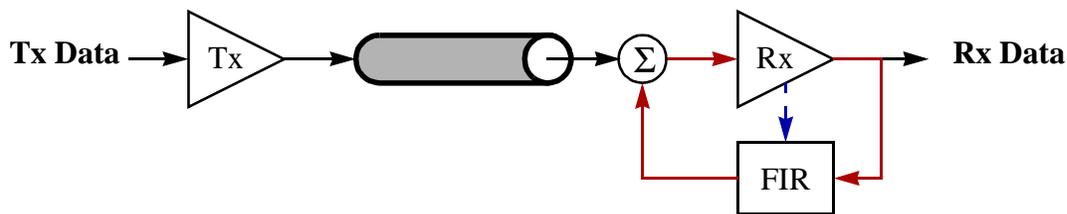


Figure 4.20: Decision Feedback Equalizer

While a DFE improves performance over a linear equalizer in channels with severe ISI [59][60], the performance gains are typically no more than 3-6dB, and DFE implementation is difficult, particularly in high speed links. The major impediment to the use of a DFE in a high speed link is the need to fully resolve and filter the previous symbol before the next symbol is received, as shown by the **loop** in Figure 4.20. This loop cannot run at the target symbol period of 1 FO4 gate delay, because it takes many FO4 delays to resolve a signal to full swing with a low probability of metastability (see Section 3.6: “High Speed Receiver Logic” on page 36). Most DFE receivers use an ADC to capture the analog received signal in real time, and then run the DFE loop as a digital post-process.

Also, since the received signal is no longer flat at the sample time, a DFE may result in lower phase noise tolerance, especially if the jitter period is shorter than the DFE transfer function. A DFE can only correct for causal interference, so a linear equalizer is still needed as a pre-filter.

4.7 Summary

This chapter presents the circuitry that completes the transceiver: the transmitter, inductive network, clock generator, logic and memories. The chapter concludes with a discussion of the communication techniques that the transceiver supports.

An overview of the time-interleaved transmitter shows that high bandwidth is achieved with a grounded-source output stage driven by low-fanout predrivers with a regulated supply voltage to control the output swing. 8 bits of resolution are achieved by individually controlling small current sources, without sacrificing bandwidth. The extra resolution is used to correct for transmitter nonlinearities and interference, and to reduce quantization errors in equalization.

The output wiring prevents the transmitters from meeting the design goals, so inductors are used to distribute the parasitic output capacitance of the transmitter. A similar inductive network in the receiver allows the full ADC to approach the sampling bandwidth of a single comparator by distributing the parasitic input capacitance.

An enclosed-dual-loop PLL generates stable, equally spaced multi-phase clocks. The inner loop has high bandwidth to reduce random noise generated in its VCO. The outer digital loop has lower bandwidth because it waits after each phase adjustment for the inner loop to settle and the new clock phases to propagate through the ADC and timing recovery logic, and because it filters phase information before making each adjustment. The phase detector has several modes to allow the performance of transition and data sampling to be compared, and to support a range of modulation and equalization approaches.

The synthesized logic and memories allow configuration and monitoring of the transceiver, as well as offline generation and processing of raw link data to support experimentation with communication techniques. The ability to compare received data

against the receive memory contents allows these techniques to be analyzed at the low bit error rates (below 10^{-12}) that are typically required.

The capabilities of the transceiver chip allow experimentation with several communication techniques. The extra resolution in the transmitter corrects for offsets, nonlinearities and interference throughout the transceiver, reducing design cost and risk by tolerating lower quality circuits, layout, packaging and wires. Multi-level PAM shows promise for increasing data rates on long, lossy wires. Linear equalization is effective at compensating for frequency dependent losses, while Decision Feedback Equalization performs the compensation with less reduction in SNR. To quantify the performance of the transceiver and of the communication techniques, Chapter 5 presents experimental results from prototype chips.

Chapter 5

Experimental Results

This chapter presents results collected from experiments performed with the 0.25 μm CMOS transceiver shown in Figure 5.1. The transceiver consists of independent receiver and transmitter sections, each with their own PLL, with the receiver making up the bottom half of the chip.

The time-interleaved receiver comparator array (Rcvr) and analog output circuitry (Transmitter) are quite small. Note the ADC input and termination pads in the lower left of Figure 5.1. These pads can either be connected as shown with short, flat bond wires with no appreciable inductance, or with small inductor coils to distribute the ADC input capacitance. Similar inductive coils are used to distribute the DAC output capacitance.

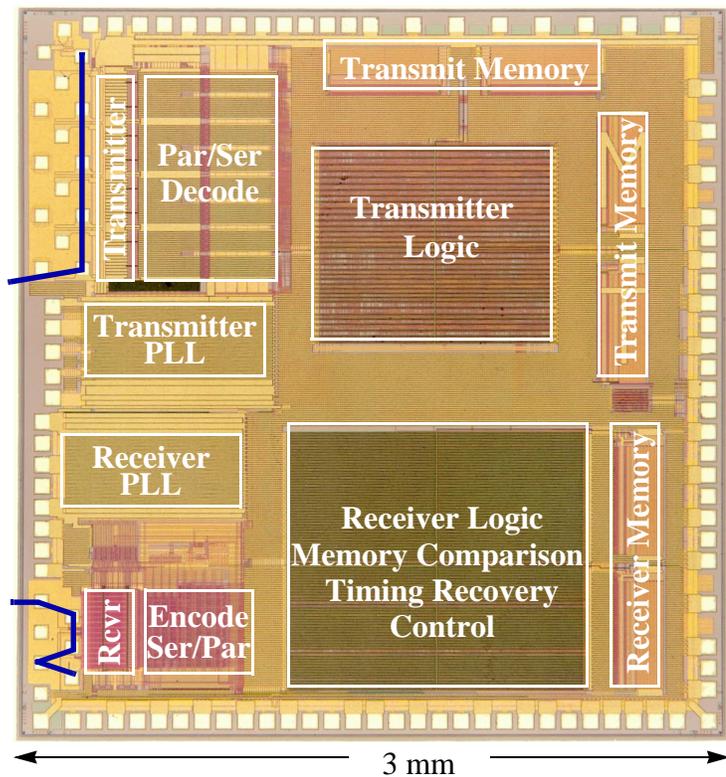


Figure 5.1: Transceiver Die Photo

This chapter begins with TDR results to verify inductance values so that the performance of the ADC and DAC over frequency can be correctly interpreted. Timing noise results are analyzed next because they also affect the remainder of the measurements. Results from the receiver are then presented, starting with the correction of static voltage and timing errors. Performance over frequency is characterized, with and without inductors. SNDR plots demonstrate the raw performance of the ADC, and LSNDR plots indicate performance in a synchronized link.

Results from the transmitter also begin with the correction of static voltage and timing errors. Frequency response data is presented with and without output inductors. SNDR and LSNDR plots demonstrate performance as an asynchronous DAC and as a synchronized transmitter. Equalization results are shown first on the transmitter alone because they are easy to understand and measure. Equalization results from the full transceiver conclude the experimental results.

5.1 Inductors to Distribute Parasitic Capacitances

Figure 5.2 is a closeup of the ADC input pads of a chip bonded with inductive coils. A diagram of a 2-turn coil with about 1nH of inductance is shown to the right of the die photo for reference. There are four inductor coils connecting the four ADC input pads and the termination resistor pad (forming the circuit on the left in Figure 4.7 on page 49). This is a beautiful piece of bonding work by Pauline Prather, who works in our lab, but is obviously not mass manufacturable. Bond wire inductors were used instead of on-chip spiral inductors to keep series resistance low, with the goal to gather data to help design mass manufacturable inductors

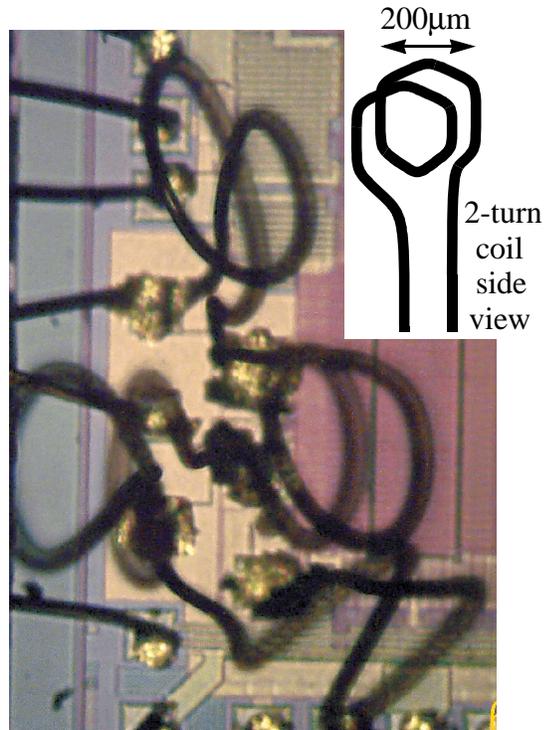


Figure 5.2: Bond Wire Inductors

into a flip chip package. Flip-chip packaging and inductors will also reduce coupling of other signals to the ADC input, which was 50% larger with the bond wire inductors than with just a single input bond wire. The TDR results below show that the bond wire inductance values are controlled within the +/- 30% range that effectively distributes the ADC input capacitance (see Section 4.3).

The upper TDR trace in Figure 5.3 shows a large dip at the ADC input due to the large input capacitance without inductors. The bottom TDR trace in Figure 5.3 shows the ADC input of a chip bonded with 2-turn inductors. The same reflections are seen from the SMA connector, board and package, but the large dip from the input capacitance is eliminated. This matches the

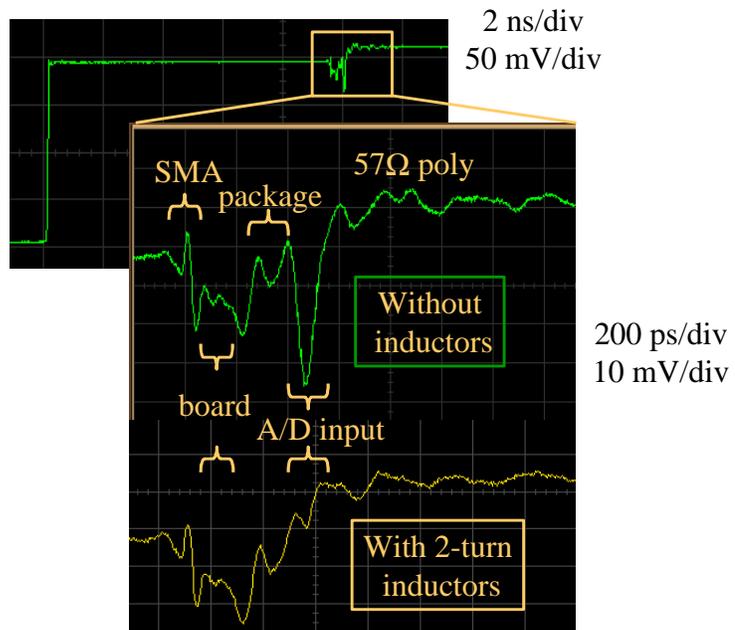


Figure 5.3: ADC Input TDR Results

trace in Figure 4.8 on page 50 for $L=1.1\text{ nH}$, and indicates that the ADC input capacitance has been effectively distributed. Thus, the ADC input bandwidth should be limited only by the sample/hold amplifier bandwidth, which is about 6 GHz.

Figure 5.4 is a TDR trace from the transmitter used in this work, with the large dip showing that the DAC output inductors are not well matched to the output capacitance of the transmitter. Because of an error in estimating the DAC output

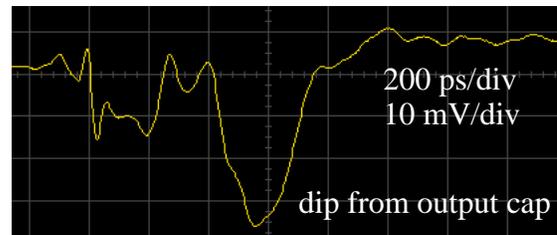


Figure 5.4: TDR of DAC Output

capacitance, inductors of half the ideal value were bonded.

5.2 Timing Noise

Before ADC and DAC performance can be verified over frequency, however, timing noise must be characterized. The timing noise is twice that measured on previous chips [2][7][12][53], due to an error in the layout technology file that causes the PLL filter capacitance to be less than half the desired value. Simulations show this results in larger phase noise and an underdamped PLL. This is confirmed by experimental measurements that show a peak in the magnitude of output jitter in response to noise at 42MHz, even though the PLL was designed for a loop bandwidth of 12 MHz.

To measure the noise response of the PLL directly, sinusoidal “noise” is added to the 300 MHz reference clock and the effect on the ADC PLL is measured. The resultant p-p timing noise on a test clock output is plotted in Figure 5.5 versus the input “noise” frequency. With no input “noise”, jitter is 40 ps_{pp}. Low frequencies are tracked by the PLL with a gain of 1 so output

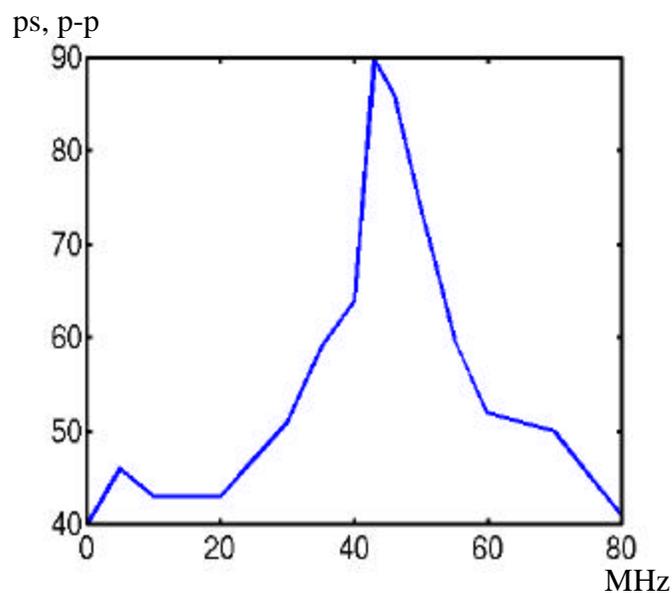


Figure 5.5: PLL Phase Noise vs. Noise Frequency

phase noise only slightly increases up to 20 MHz. Then, noise is amplified by the PLL, with a peak of 90 ps of jitter near 40 MHz.

Although Figure 5.5 shows results with a reference clock frequency of 300 MHz, the lowest timing noise is seen with a reference clock frequency of 256 MHz, where 30 ps_{pp} is measured. This corresponds to $0.24T_{\text{symbol}}$ or approximately the target for a 3-bit transceiver (see Section 2.1). Thus, the remainder of the results in this chapter use a reference clock rate of 256 MHz, for a symbol rate of 8.2 GHz (given the 4:1 PLL multiplication and 8:1 time-interleaving).

5.3 ADC Measurements

After measuring timing noise, the first step in characterizing the ADC is to correct static¹ voltage and timing errors. Figure 5.6 shows a histogram of the uncorrected input offset errors of the 120 comparators in one ADC. The input offset voltage is 27mV_{RMS} and 131mV_{pp} over the 120 comparators in one chip, which is large due to the small transistors used (for low input capacitance) and due to the low gain

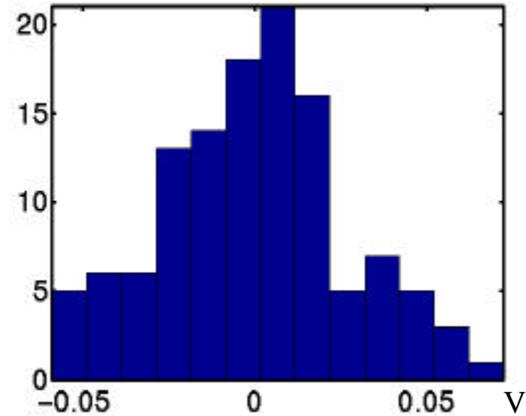


Figure 5.6: ADC Input Offset Histogram

of the sample/hold amplifier. The offset voltage is somewhat larger than the simulated 20mV_{RMS} of the transistor mismatch errors (see Section 3.8 on page 38). Coupling to the high speed ADC **clk** phases² is $0.82\text{ LSB}_{\text{pp}}$ (41 mV_{pp}) and has been removed from Figure 5.6, reducing the RMS and the p-p input offset voltage by about 10%.

The digital inputs of the 4-bit DACs in each comparator are adjusted to correct for input offset errors throughout the ADC. The calibration loop includes the entire comparator and reference ladder by driving voltages into the ADC and making corrections based on the digital output of the ADC. However, the 4-bit offset correction DACs are themselves nonideal, as can be seen from the scatter plot in

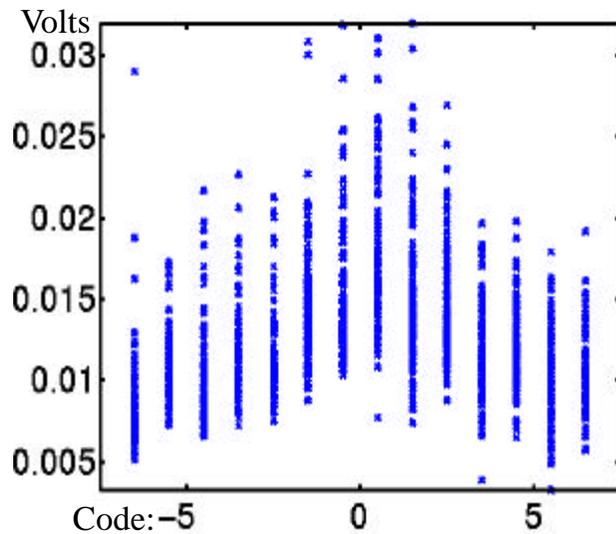


Figure 5.7: Offset Correction DAC Step Sizes

1. Voltage errors are time averaged to filter out random noise, which is discussed below.

2. Coupling errors are measured as the mean offset of the 15 comparators connected to each **clk** phase.

Figure 5.7 of the DAC step sizes vs. the DAC control values. To generate this plot, the switching point of each of the 120 comparators is measured for each of the 15 valid DAC control values, and subtracted from the switching points for adjacent control values. Random mismatches in the DAC current source transistors cause the 15-20mV_{pp} spread in the step size at each control value. There are also two types of systematic variations in step sizes: steps are larger for smaller control values, and the least significant bit is slightly larger due to smaller clock coupling to the least significant control voltage (see Section 3.4.1 on page 32). The DAC step sizes¹ average 13mV and range from 3mV to 30mV, so peak-to-peak offset voltage errors are reduced to about 30mV_{pp} (0.6LSB_{pp}).

After calibration, the averaged Integral Non-Linearity (INL) errors are 0.6LSB_{pp} (see Figure 5.8). The errors are primarily due to quantization errors from the offset correction DACs. INL errors on rising and falling slopes are nearly identical, with an average comparator hysteresis of 2.2mV (0.04LSB). No code bit errors are seen over 192,000 samples.

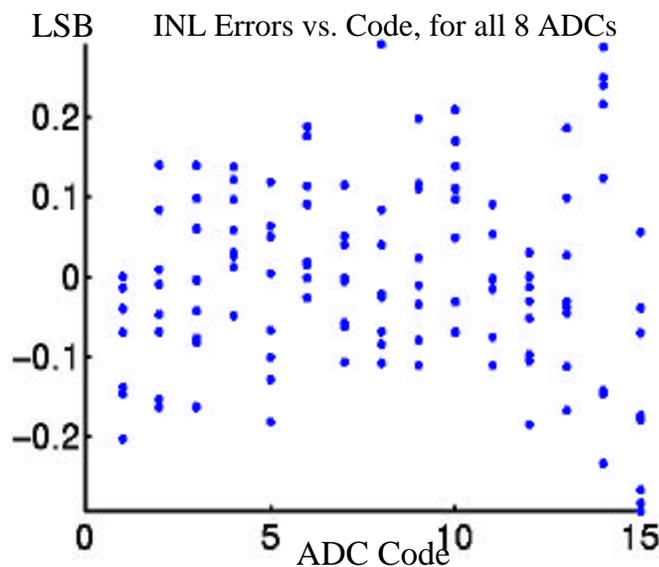


Figure 5.8: ADC INL After Calibration

Note that offset calibration corrects for coupling to the high speed ADC **clk** phases, since this coupling is the same every cycle. However, the receiver is not designed to correct for coupling to **clkDiv4**, and inductive bondwire coupling from **clkDiv4** to the ADC input is measured at 0.9LSB_{pp} (by averaging the offset errors against the four phases of **clkDiv4**). The transmitter can calibrate out **clkDiv4** coupling, so this coupling is removed from Figure 5.8 and Figure 5.9 to allow other noise sources to be investigated.

1. Some offset control values resulted in overlapping codes or switching beyond the linear input range, so the step sizes weren't measured.

The remaining noise is shown in the probability density distribution of the non-averaged INL errors in Figure 5.9, normalized to the distribution of ideal quantization noise. The total $1.6\text{LSB}_{\text{pp}}$ INL error data (dots) is compared to a simulated model (line). To generate the measured dots in Figure 5.9, 750 DC voltages from V_{dd} to $V_{\text{dd}}-750\text{mV}$ are driven into the ADC, and the digital

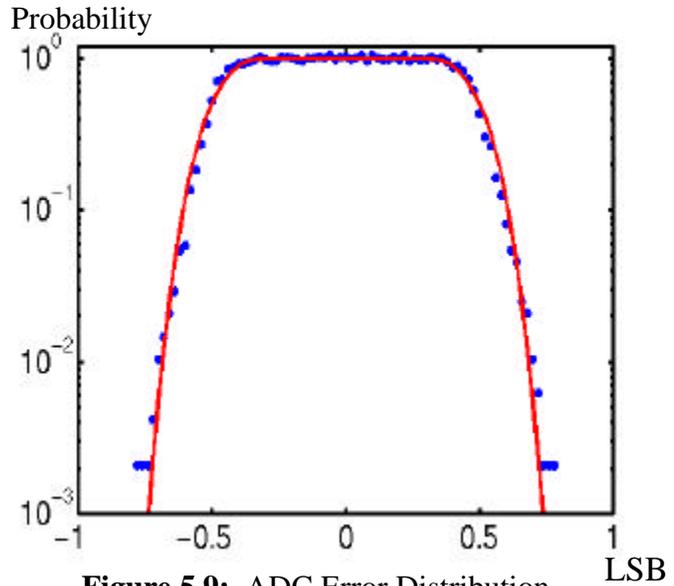


Figure 5.9: ADC Error Distribution

output captured for 128 samples at each voltage. The average switching points of each comparator are computed and subtracted from the data, so that the comparator switching points are at $\pm 0.5\text{LSB}$. The probability of the remaining errors are plotted vs. distance from the nearest comparator. The data was fit to a model error distribution (line) consisting of a 1LSB_{pp} rectangular distribution for uniform quantization noise convolved with normally distributed noise with $\sigma=0.08\text{LSB}$ (vs. $0.07\text{LSB}_{\text{pp}}$ estimated in Section 3.3), and with $0.2\text{LSB}_{\text{pp}}$ of uniform noise¹. The total non-averaged errors from all sources are $0.39\text{LSB}_{\text{rms}}$ for a Signal-to-Noise-and-Distortion Ratio (SNDR) of 22.7dB which corresponds to 3.5 effective bits at low frequency.

1. Variation in the 0.7LSB of **clkDiv4** coupling with 250ps rise/fall times, due to 30ps_{pp} phase noise is estimated: $0.7\text{LSB} \cdot 30\text{ps} / 250\text{ps} = 0.1\text{LSB}_{\text{pp}}$. Estimated errors in measuring offsets: $0.1\text{LSB}_{\text{pp}}$.

Now that the low frequency response of the receiver has been analyzed, performance over frequency can be investigated. The ADC frequency response in Figure 5.10 shows a bandwidth of about 6 GHz with inductors, with several dB less loss above 4 GHz than without inductors. This plot was generated by driving asynchronous sine waves into the ADC, capturing the digital output of the ADC from

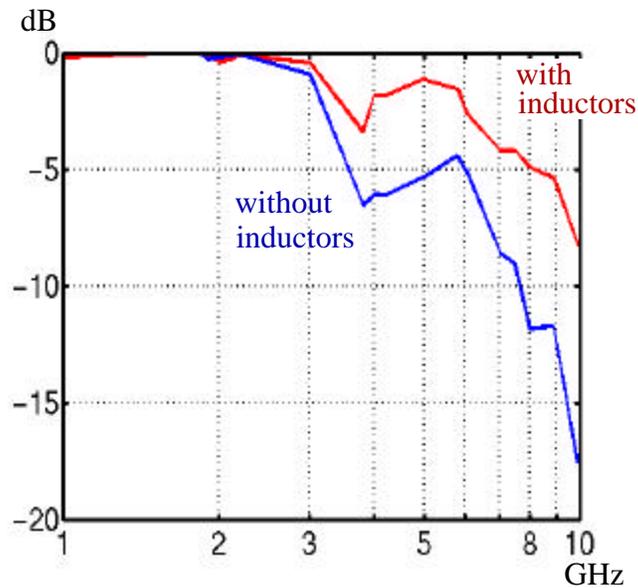


Figure 5.10: ADC Frequency Response

the on-chip memory, fitting ideal sine waves to that data, and plotting the amplitudes of the ideal sine waves. The frequency response without inductors is better than expected at some frequencies due to reflections between the ADC input and packaging (see discussion of Figure 4.8 on page 50), which causes the dips and peaks in the frequency response of the input network. This is actually demonstrating a similar principle, that the bond wire inductance can help tune out the input capacitance.

While the frequency response data in Figure 5.10 indicates that inductors increase the input bandwidth of the ADC, a broader measure of performance must be considered. Figure 5.11 is a plot of the SNDR of the ADC, showing a resolution bandwidth of only about 3 GHz. This is due to phase noise, which causes distortion that increases with frequency and is proportional to

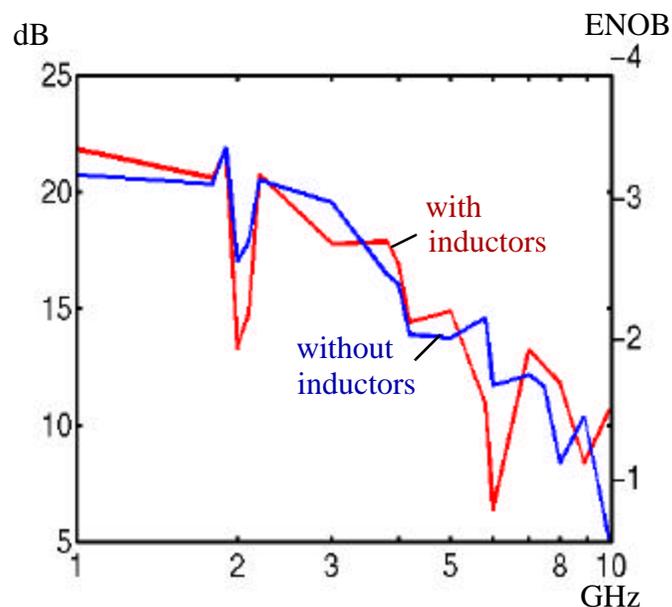


Figure 5.11: ADC SNDR

the sampled signal. Inductors do not improve the SNDR, because the larger received

signal with inductors is accompanied by an increase in the distortion caused by phase noise.

The measured phase noise of the test chip is double that of prior art (see “Timing Noise” on page 68) due to a layout technology file error. This error also heightens sensitivity to coupling from the single-ended ADC input to the PLL reference clock. The coupled ADC input is downsampled by the PLL, exciting its large peak in phase noise sensitivity, so that large dips in SNDR are seen at input frequencies near clock harmonics¹.

Three of the four possible sources of noise that increase with frequency and are proportional to the input signal are deterministic, and are modeled and fit to the data to see if they explain the increase in distortion with frequency. First, the ADC sample time varies with input level, but only by 15ps/V, and when fit and removed as a source of error, does not significantly change the SNDR plot. Second, static phase errors are reduced from 47 ps to 10 ps p-p with calibration, and when fit and removed, also leave the SNDR primarily unchanged. Third, no significant phase noise correlated to the subsampled input frequency is measured except at the dips in SNDR in Figure 5.11. This leaves random phase noise as the culprit, measured at 30ps_{pp} on a test clock output, which results in distortion that reduces the SNDR of a 4-bit ADC by 6dB at 4 GHz (see Figure 2.4 on page 9). A PLL with less phase noise would reduce the drop in SNDR: similar designs have achieved 2%/cycle of jitter p-p, which would cause a drop in SNDR of 1dB at 4 GHz.

1. Many points near 2 GHz are plotted to show the dip shape, but most other peaks do not coincide with the infrequencies that are measured in Figure 5.11.

As discussed in Section 2.1, SNDR and ENOB are not the best measures of ADC performance in a link. When the link is synchronized, phase noise is centered at the lowest slope of the signal, and thus maps to relatively little voltage noise. Figure 5.12 plots the p-p Link SNDR (LSNDR) and Link ENOB (LENOB) defined in Section 2.1. The large phase

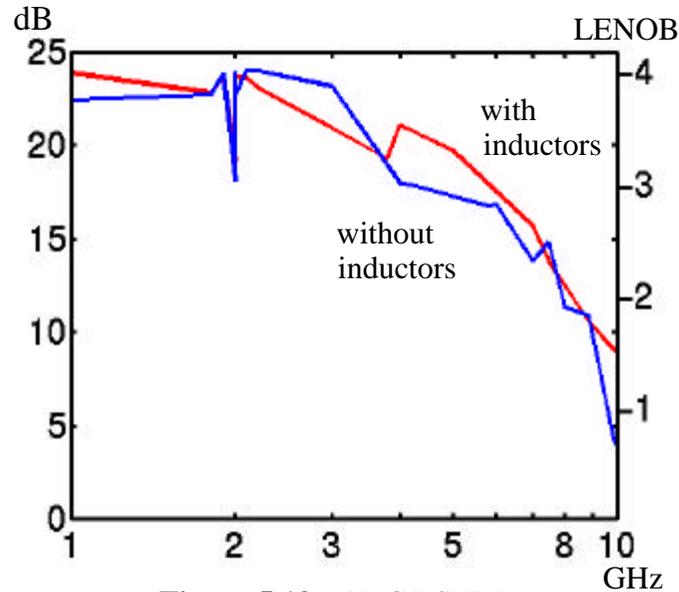


Figure 5.12: ADC LSNDR

noise still limits performance, so little improvement is seen with inductors. But note that the measured 30ps_{pp} of phase noise has little effect on the LSNDR of a synchronized 4GHz sine wave. Because PAM signalling is approximated by a sine wave at half the sample rate, a 4 GHz bandwidth can carry 8 GSymbols/sec. Thus, the LENOB indicates the number of bits/symbol that can be received in an equalized link at a symbol rate equal to twice the signal bandwidth, (neglecting equalizer losses). This represents an upper limit on receiver performance if the transmitter has the same signal swing as the receiver, as it does in this work.

Table 5.1 lists key performance parameters of the ADC. While the ADC draws a little over 1 Watt, dominated by the oversize PLL and phase adjusters, this power consumption could be cut in half through better sizing of clock buffers. Offset correction using the DACs in each comparator reduces offset and clock coupling errors from 2.9 to 0.49 LSBpp. The phase adjusters reduce static timing errors from 47 to 10 ps_{pp} . The calibration is stable over environmental conditions, with results [5] showing only a 2dB reduction in SNDR due to variations in temperature from about 0 to 70°C and to variations in supply of +/-10%. Losses and nonidealities in the transmitter also reduce performance, so the transmitter is characterized next.

Supply	1.1 Watt from 2.5V at 10 GSa/s
Technology	National 0.25μm CMOS
Analog ADC Circuit Area ADC and logic, memory, PLL, pads	0.3 x 0.2 mm 1.7 x 3.5 mm
Sample Rate	10 GSample/sec
Input Bandwidth	3.2 GHz (without inductors) 6 GHz (with inductors)
Effective Number of Bits (ENOB)	3.5@DC, 2.2@4 GHz
Link Effective Number of Bits (LENOB)	3.8@DC, 3.2@4 GHz
Reduction in SNDR over Temp. and Supply	2 dB
Voltage Offset Errors (including high speed clk coupling)	144 mV_{pp} before calibration 30mV_{pp} after calibration
Static Phase Errors	47 ps_{pp} before calibration 10 ps_{pp} after calibration
Phase Jitter	30 ps_{pp}
Sample Time Variation vs. DC level	11 ps_{pp}
Input Range	V_{dd} to V_{dd}-0.75 Volt

Table 5.1: ADC Performance Summary

5.4 DAC Measurements

The operation of the DAC is easier to visualize than the ADC, because the output waveform can be displayed directly on an oscilloscope. The ease of visualization is fortunate, because the calibration of the interleaved DAC is complex. Results from cancellation of static phase and voltage errors are first presented, followed by results from the equalization of frequency dependent effects.

The correction of static phase errors is the first step in the calibration of the transmitter. Each of the 8 time-interleaved DACs is controlled by two adjustable clocks - one to start the pulse, and the second to end the pulse 1/8th of a clock cycle later (see Section 4.2 on page 44). Because the internal clock transition time is typically more than a bit time long, the start clock and end clock adjustments affect each other (their transition times overlap slightly).

Therefore, they are adjusted in concert by a software algorithm that repeatedly moves each pulse toward its ideal position, adjusting the pulse width at each new position. Figure 5.13 shows pulses from two transmitters before phase calibration (the scattering of points on the sampling scope is due to timing and voltage noise). Because the DACs use NMOS current sources, the pulses are negative going, corresponding to a 10111011 data pattern. The second pulse is smaller, and a bit later than it should be.

Figure 5.14 shows the output pulses after calibration. The pulses are approximately the same width and height, and the second pulse has been moved closer to its ideal position. Now that the pulses are correctly positioned, the lower bits of the transmitter are used to compensate for variations in clock coupling.

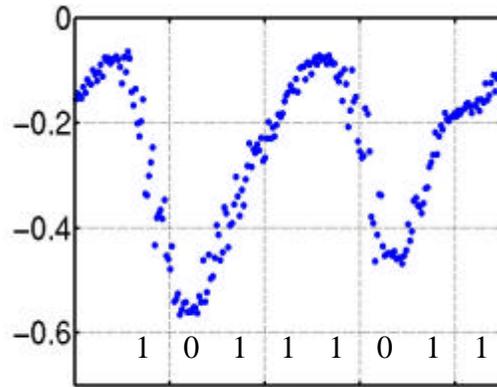


Figure 5.13: Uncalibrated Pulses

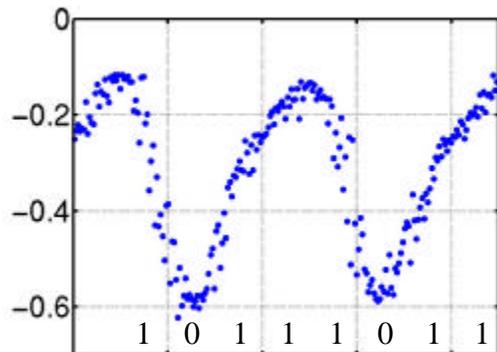


Figure 5.14: Calibrated Pulses

Coupling from the DAC clocks to the output causes periodic voltage noise that can be cancelled by transmitting “anti-noise”. Figure 5.15 shows the coupling noise from the transmitter with all outputs off. The clock coupling is primarily through the upper transistors in the DAC current sources. Although the coupling to the 8 clock phases that drive the upper transistors nominally cancels, static variations in the clock

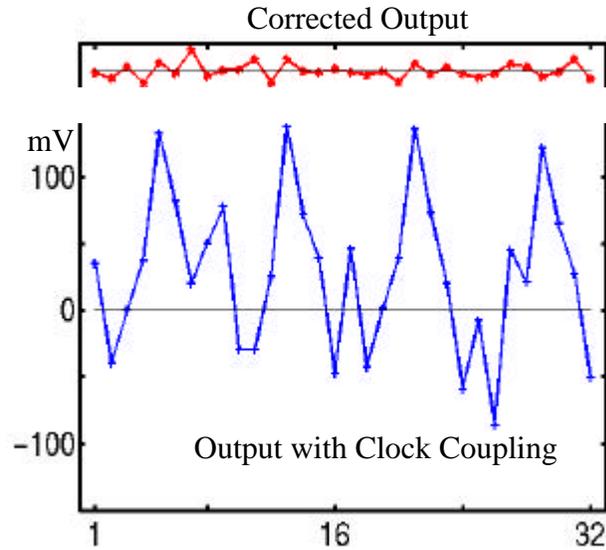


Figure 5.15: Transceiver Clock Coupling

phases and in the coupling parasitics, and coupling to **tlclkDiv4**, result in periodic voltage noise on the DAC output which can be cancelled. The clock coupling in the DAC is cancelled by adding small values to the transmitted data to negate the noise at each symbol time. This demonstrates the extent to which DAC performance can be improved with calibration, albeit at the cost of a reduction in usable output swing.

After cancelling periodic noise, the voltage errors due to transistor mismatches and output nonlinearity are also reduced from $3.7\text{LSB}_{\text{pp}}$ to 1LSB_{pp} after calibration. The envelope plots in Figure 5.16 show the worst case INL and Differential Non-Linearity (DNL) errors measured on averaged voltage pulses from the 8 time-interleaved DACs, at each output code (this measurement

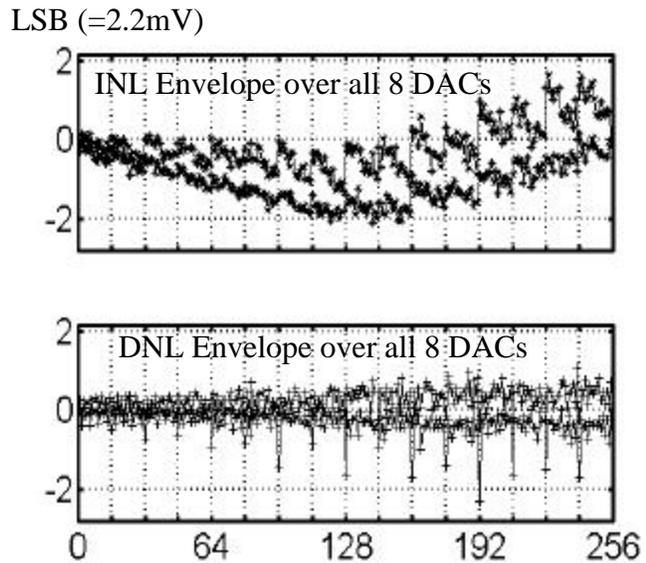


Figure 5.16: DAC INL and DNL vs. Code

does not include output noise). The DNL and INL errors are plotted in Least Significant Bits (LSBs), with a mean LSB of 2.2 mV, versus the output code from 0 to 255, for an

output voltage swing of 561 mV (this data is measured on single pulses, which do not reach the full 750 mV swing). Three effects can be seen: random transistor mismatch errors, nonlinearity due to variation in the current source output impedance with output voltage (seen in the INL plot only), and a systematic mismatch every 16th code. Fortunately, the large systematic DNL errors every 16th code are negative, and thus cause overlapping codes instead of large gaps¹. Therefore, they do not increase errors after calibration, which are determined by the distance to the closest possible output voltage.

Non-averaged measurements of DAC output pulses reveal the noise shown in Figure 5.17 (clock coupling and INL errors are removed). This noise is roughly proportional to output current, with $50\text{LSB}_{\text{pp}}=110\text{mV}_{\text{pp}}$ of noise at full output. The measured transmitter phase noise of $36.5\text{ps}_{\text{pp}}$ would cause $36.5\text{ps}/125\text{ps}/2*750\text{mV}=110\text{mV}_{\text{pp}}$ on an ideal triangle wave. Since the output waveform with inductors is approximately triangular,

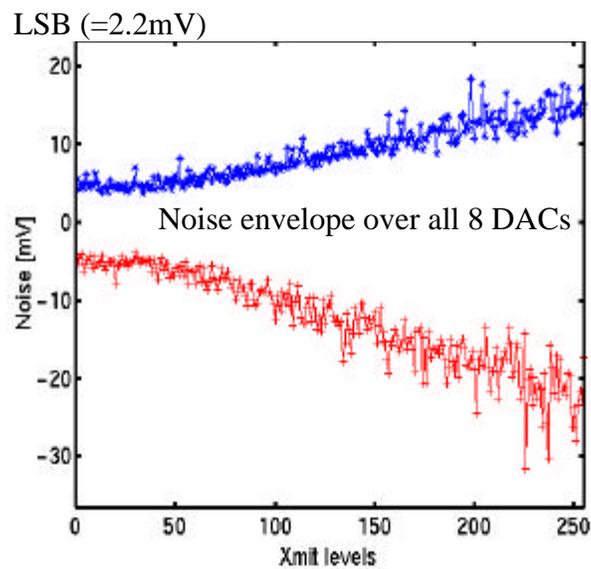


Figure 5.17: DAC Noise vs. Code

phase noise is the dominant source of DAC voltage noise. Noise on low frequency output waveforms (instead of pulses) is smaller, varying from 15 to 25mV_{pp} with output code.

The large DAC noise limits performance both as a link transmitter and as a DAC. The average DAC noise of $33\text{LSB}_{\text{pp}}*2.2\text{mV}=71\text{mV}_{\text{pp}}$ results in an estimated LSNDNR of $20\log(750/71)=20.5\text{dB}$ and LENOB of 3.4 bits at DC. The 3 GHz output bandwidth (see Figure 5.20) leads to a loss of about 4 dB at 4 GHz, and an estimated LENOB of 2.7 bits. To compute the low frequency SNDR, the RMS noise is estimated to be 1/6 of the asynchronous p-p noise (assuming a normal distribution, 10^5 samples), which varies from

1. In fact, intentionally designing upper binary bits of a DAC to be slightly smaller than the lower bits (designing a radix < 2 DAC) can decrease the effect of transistor mismatches in calibrated binary DACs, and avoid the need for thermometer encoding of upper bits.

29mV_{pp} to 110mV_{pp}. With the simplifying assumption that the signal is high half the time and low half the time (i.e. for a square wave output), $V_{\text{noise,RMS}} = \frac{1}{6} \sqrt{\frac{29^2 + 110^2}{2}} = 13.4\text{mV}_{\text{RMS}}$. Adding 9mV_{pp}=3.7mV_{RMS} of uniform quantization and clock coupling and DNL errors after calibration, the low frequency SNDR is estimated as $20\log(750\text{mV}/17.1\text{mV})=32.8\text{dB}$, which corresponds to 5.2 effective bits. The 36.5ps_{pp} of phase noise and the 4dB of loss at 4GHz limit the SNDR to 2 effective bits at 4 GHz (see Figure 2.4 on page 9).

Table 5.2 lists DAC performance data. The DAC dissipates about twice as much power as the ADC, because it has twice as many clocks, and this power could also be cut in half by better sizing of the clock buffers.

Supply	1.7 Watts from 2.5V at 8GSa/s
Technology	National 0.25μm CMOS, 2.5V Supply
Analog DAC Circuit Area DAC and logic, memory, PLL, pads	0.9 x 0.1 mm 1.8 x 3.5 mm
Sample Rate	8 GSample/sec
Output Bandwidth	1.5 GHz (without inductors) 3 GHz (with inductors)
Effective Number of Bits (ENOB)	5.2@DC; 2@4 GHz (estimated from data)
Link Effective Number of Bits (LENOB)	3.4@DC; 2.7@4 GHz (estimated from data)
Phase Jitter	36.5 ps_{pp} (σ = 4.8ps)
Output Range	V_{dd} to V_{dd}-750mV

Table 5.2: DAC Performance Summary

5.5 Equalized Transmitter Results

To demonstrate link operation, equalization and interference correction results are presented (without inductors). Since transmit waveforms are easier to capture and display than waveforms inside the receiver, the techniques are first developed on the transmitter alone. Measurements from the full transceiver (with inductors) conclude the experimental results.

Figure 5.18 shows an oscilloscope trace of a single pulse from one interleaved transmitter at 9.5 GSymbols/sec. Phase noise results in thicker vertical traces than horizontal traces. The pulse is spread out by the parasitic transmitter capacitance so that there is a significant pulse response in the symbol time before the peak, and in the two symbol times after the peak. The pulse response \mathbf{p} is approximately:

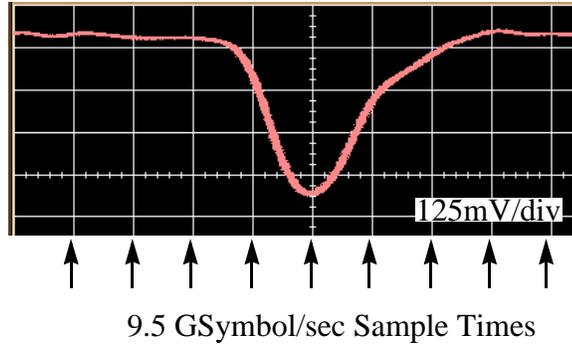


Figure 5.18: Raw DAC Pulse

$$\mathbf{p} = 0.3 + 1.0D + 0.4D^2 + 0.2D^3 \quad [2.11]$$

where D represents a one symbol delay (equivalent to Z^{-1}). The 0.3 represents a measurable response one symbol time before the main symbol peak represented by $1.0D$. To equalize the pulse response, negative values are added to the transmitted waveform before (to cancel the 0.3) and after (to cancel $0.4D^2 + 0.2D^3$) each symbol. This eliminates pulse responses in adjacent symbol times (inter-symbol interference).

Figure 5.19 shows the single transmitter pulse after linear equalization. Since the transmitter has a limited output voltage swing, the equalized voltage swing is reduced to allow high frequency signal components (transitions) to be driven

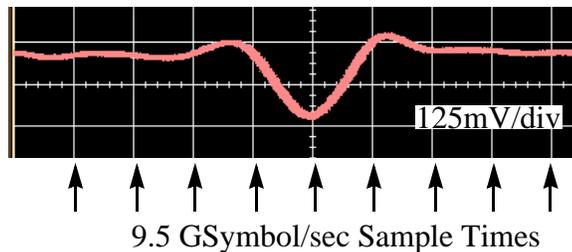


Figure 5.19: Equalized DAC Pulse

harder than low frequency components. Also, because only one value per symbol is transmitted, the intermediate signal values cannot be controlled, and signal components above the Nyquist frequency are not equalized, causing the ripples between sample times. Note, however, that the values at the symbol times (denoted by arrows at the bottom of Figure 5.19) are nearly zero with a single peak. The factor of 2.1 reduction in pulse height (6.5 dB) from Figure 5.18 to Figure 5.19 is due to loss in the linear transmit equalizer. More advanced equalization techniques such as Decision Feedback Equalization (DFE) and Tomlinson precoding have less loss.

Figure 5.20 shows the measured pulse responses (on the left) at each symbol time of the eight interleaved DACs. The raw pulse response data is used to compute the equalized values to transmit. A closed form expression for the ideal equalized values developed for time-interleaved transceivers is described in [13], and is based on the theoretical approach derived in [62]. A zero-forcing algorithm is used, which reduces the inter-symbol interference to zero. The FFT of the sampled pulse responses on the right in Figure 5.20 shows equalization in the frequency domain. By attenuating low frequency signal components, the transmit equalizer produces a received signal that is evenly attenuated over frequency.

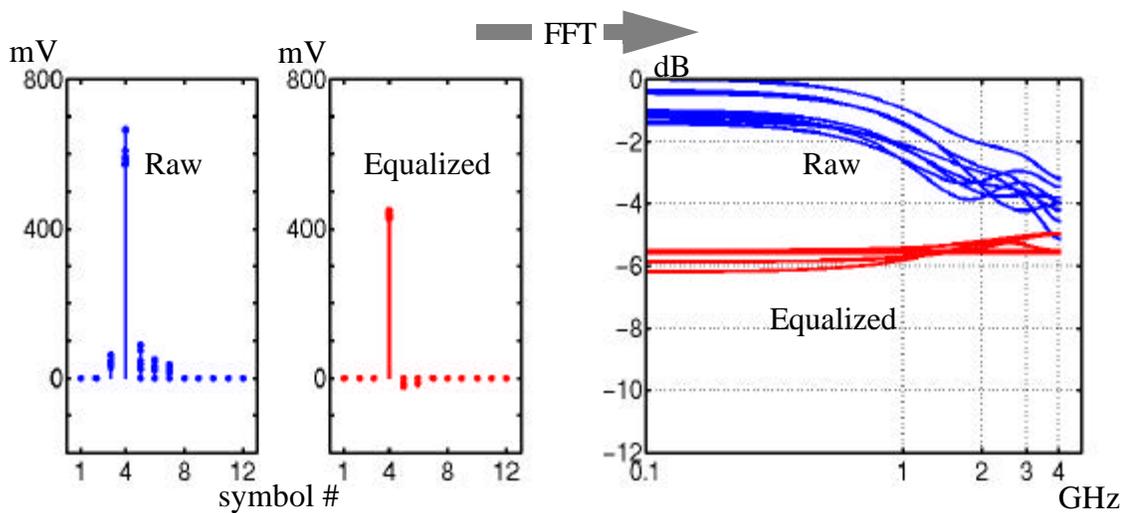


Figure 5.20: 8 GSa/s Transmitter Pulse and Frequency Response

5.6 Full Transceiver Results

The equalization techniques developed on the transmitter are applied to a full link from transmitter over a short cable to the receiver. Only a short cable is used, as the transceiver bandwidth of 3 GHz presents a good challenge to 8 Gsymbol/s data transmission. Figure 5.21 shows the pulse and frequency response of the transceiver before and after equalization. The pulses have varying heights due to residual pulse width variations between the time-interleaved DACs, and have varying inter-symbol interference due to differing positions on the lumped LC transmission lines. Thus, each transmitter uses a unique equalization lookup table (see Section 4.6.1 on page 57). After equalization and gain correction, the pulse heights are nearly the same, and the interference in adjacent symbol times is significantly smaller.

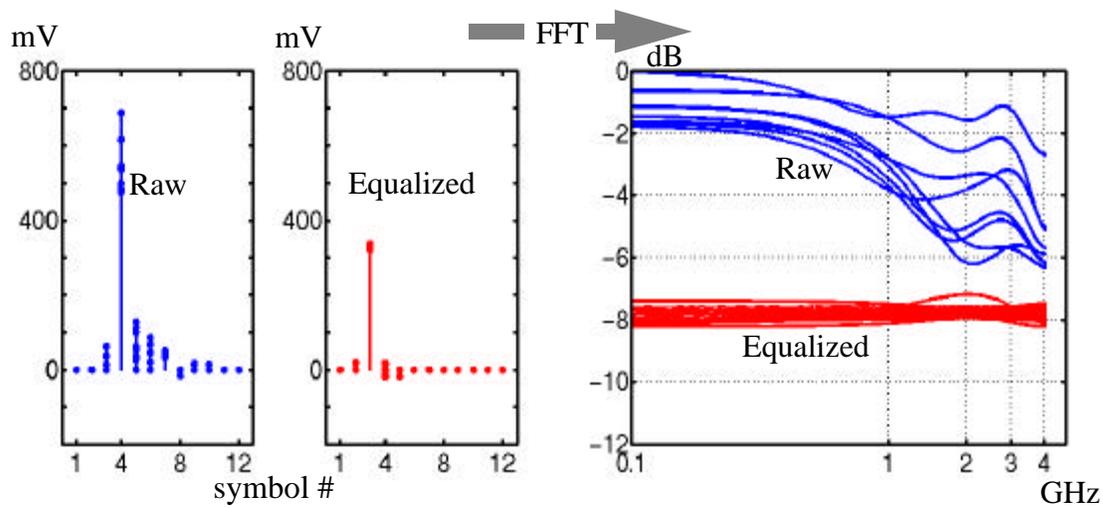


Figure 5.21: 8 GSa/s Transceiver Pulse and Frequency Response

To avoid quantization errors in equalization, the pulse responses are measured by adding a varying DC voltage to the signal to find the average switching point of the most significant bit comparator in each time-interleaved A/D converter. Therefore, this data does not include noise or ADC quantization errors, although it does include parasitic filtering and clock coupling from both the ADC and the DAC.

A similar technique is used to measure the voltage and phase margins of the link. The schmo plots shown in Figure 5.22 are generated by varying a DC voltage added to the signal and varying the phase of the receiver clocks, and plotting a dot where bit errors

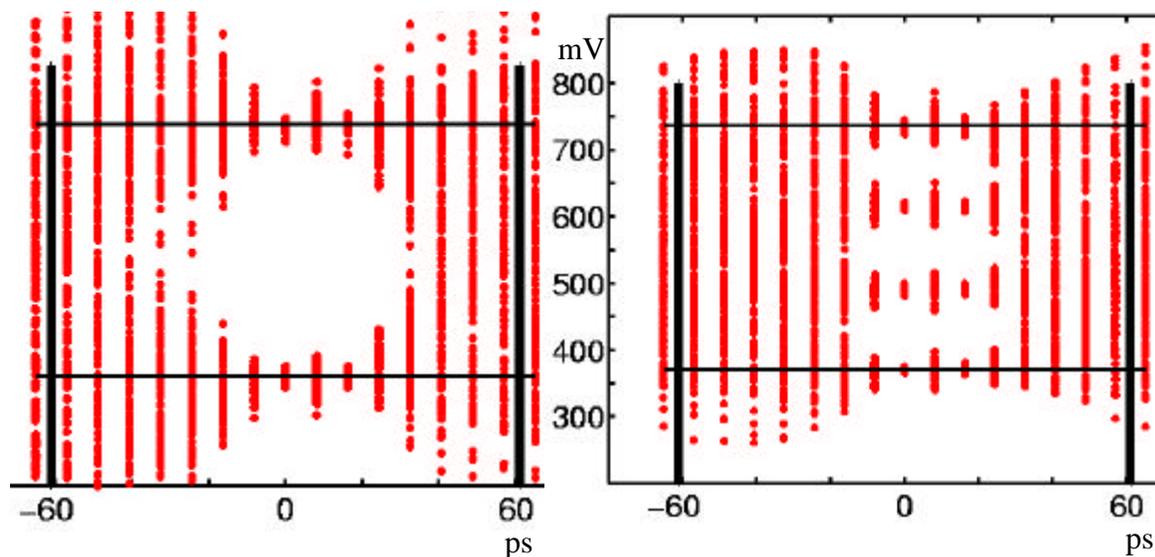


Figure 5.22: Equalized Binary and 4-Level Transceiver Schmo Plots

are recorded. The eye openings are less than the 750 mV signal swing and 125 ps symbol time, but can still support 4-level signalling. The received signal swing is reduced by the 3 GHz bandwidth, and by clock coupling correction. Performance is limited by the combined transceiver phase noise of 47ps_{pp} , but 300mV_{pp} binary and 100mV_{pp} 4-level eye heights are seen, with 45ps_{pp} eye widths. These schmoo plots correspond to a bit-error-rate of about 10^{-4} .

Figure 5.23 is generated using the on-chip sequence compare circuitry, with each point corresponding to 10^{10} samples. The larger number of samples increases the peak-to-peak random phase and voltage noise, from 5.5σ at 10^{-4} to 9σ

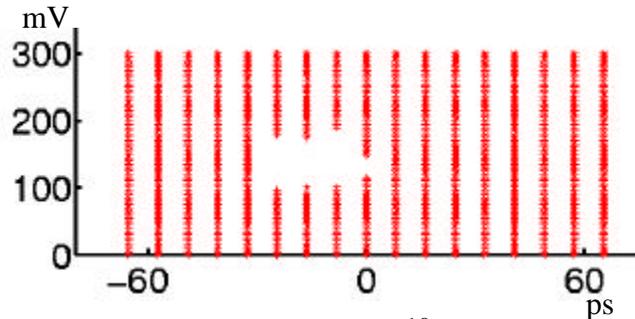


Figure 5.23: Transceiver 10^{-10} BER Schmoo

at 10^{-10} . Binary operation is still verified, although with a smaller eye opening of 75mV by 30ps. Thus, despite large phase and voltage errors and significant clock coupling and bandwidth limitations, Figure 5.22 and Figure 5.23 show that an accurate 8 GSymbol/s multi-level transceiver can be built in $0.25\ \mu\text{m}$ CMOS.

Table 5.3 summarizes the performance of the full transceiver.

Technology	National $0.25\ \mu\text{m}$ CMOS, 2.5V Supply
Die Size	3.5 x 3.5 mm
Package	52-pin LCC
Sample Rate	8 GSample/sec
Bit Error Rate	$<10^{-10}$ at 8 Gbit/sec (2-PAM) $<10^{-4}$ at 16 Gbit/sec (4-PAM)
Transceiver Bandwidth	1.4 GHz (without inductors) 3 GHz (with inductors)
Signal Swing	V_{dd} to $V_{\text{dd}}-0.75$ Volt into $50\ \Omega$
Power Dissipation	2.7 Watts total @8 GSymbol/s
On-chip Memory	1024 x 4 bits and 1024 x 8 bits

Table 5.3: Transceiver Performance Summary

5.7 Summary

This chapter presents experimental results that demonstrate the performance possible in a calibrated CMOS transceiver. The value of the inductors used to distribute parasitic capacitances are shown to be well-matched for the ADC and too low for the DAC. Timing noise is measured at 35 ps with undesirable sensitivity to noise well beyond the intended loop bandwidth, complicating link performance analysis at high frequency.

Static voltage and phase offsets were cancelled in the receiver, but performance is limited by random voltage noise. ADC input bandwidth was shown to approach the 6 GHz bandwidth of the ADC circuits alone with inductors that distribute the parasitic input capacitance. Thermal noise was shown to be a concern at the bandwidth and accuracy goals of the ADC.

Transmitter results showed bandwidth of 1.5 GHz that only improved to 3 GHz with inductors because of the incorrect inductance values. Equalization of the transmitter resulted in large losses to compensate for the low bandwidth without inductors and for the variations in output bandwidth with the poorly matched inductors. The fortuitous systematic undersizing of the 4th binary bit of the DAC suggests that calibrated DACs can be designed with a radix less than 2, avoiding thermometer-encoded upper bits that increase DAC cell size and wire capacitance. The DAC resolution is limited by random noise at low frequencies, and phase noise at high frequencies.

Results from the full transceiver demonstrate that the effects of a wide array of parasitic losses, noise and interference can be reduced in a calibrated transceiver. An equalized eye diagram for the full transceiver shows that reliable communication can be established despite significant interference and parasitics. Multi-level signaling is possible, especially with reductions in equalizer loss and improvements in timing and thermal noise.

Chapter 6

Conclusion

This work shows that high bandwidth, accurate communications circuits can be implemented with CMOS transistors, using small transistors to minimize parasitic capacitances, and calibration to correct static errors. The range of correctable, static errors in the transceiver is large, and includes several unanticipated error sources, showing that calibration can reduce both errors and risk. CMOS ADCs and DACs are demonstrated at binary transceiver rates, enabling digital communications techniques on high speed links to increase data capacity.

In the receiver, a high bandwidth sampler is demonstrated, with bandwidth comparable to a pass-gate sampler, but with less sample-time distortion and kickback, less variation in sampling bandwidth with input level, and an input range extending up to V_{dd} . The sampling amplifier requires well-matched complementary clocks, but more advanced bias circuitry, or a redesign for single clock operation may allow the high performance of tri-state sample/hold amplifiers to be applied more widely. The StrongArm latch is shown to be an effective sample/hold circuit, especially when preceded by a differential amplifier. The output latch efficiently captures the output waveform of the StrongArm latch, holding it while the latch is reset without needing a separate clock. Offset calibration reduces the static input offset errors of the receiver from $144mV_{pp}$ to $25mV_{pp}$. The small, low gain sample/hold amplifier also results in thermal noise of 0.4 LSB, showing that thermal noise is increasingly important in high performance CMOS circuits.

The transmitter DAC design proved to be more difficult than simply dividing a binary transmitter into segments. Clock coupling, transistor matching and output impedance variations emerged as significant challenges. Fortunately, calibration can correct these flaws. Still, transistor matching errors limit the resolution to 5 effective bits after calibration. Bandwidth is limited by output wire capacitance, which can be improved significantly with careful layout.

Clock generation limits the performance of the transceiver. The PLL bandwidth does not meet simulated specifications and measured timing errors are twice as large as those measured in other similar designs. Performance is only measured at the best clock frequencies, and is expected to increase with an improved PLL. However, the performance of the transmitter and receiver match simulated results in bandwidth and static offsets. Since the circuits work at a symbol time of 1 FO4 delay with $47\text{ps}_{\text{p-p}}$ of jitter, and since that jitter has been more than halved in previous clock generation circuits, symbol times below 1 FO4 delay are possible.

Calibration techniques improve accuracy while maintaining speed: extra DAC resolution corrects for interference and simple nonlinearities, and ADC offsets are corrected with a DAC inside each comparator. Calibration reduces static errors in both the receiver and transmitter. The static errors that are corrected include transistor mismatch errors, clock coupling, and both random and systematic layout asymmetries. Calibrated equalization allows frequency dependent effects including wire losses, reflections, and circuit bandwidth limitations to be corrected. Timing accuracy is improved with clock interpolators to correct static phase errors. In all, these corrections allow high performance to be achieved, and lower risk by tolerating imperfections in layout and fabrication. This was crucial, because we learned more than we ever expected to on this project, but were able to use calibration techniques to correct most of the mistakes that we learned from.

High bandwidth is achieved by using inductors to trade delay for bandwidth, and by designing with small transistors for their low parasitics and low power consumption. The use of inductors increases the bandwidth of the receiver more than the transmitter because the transmitter inductors were incorrectly sized. Still, the inductors significantly increase the transceiver bandwidth. The use of inductors breaks the tradeoff between

parasitic capacitance and bandwidth in interleaved data converters, allowing more interleaved transceivers or larger devices. Since the resulting LC delay can be compensated by phase shifting the clocks in an interleaved data converter, heavily interleaved comparators or large device sizes are possible. However, inductive coupling may increase unless inductor layout is well engineered.

Experimental results on communications techniques were clouded by excessive phase noise and output bandwidth limitations, but results from linear equalization of the transceiver (with a short cable) did show that the significant frequency dependent losses due to reflections and circuit and package parasitics can be corrected. Significant reduction in timing errors, noise and interference were demonstrated in a calibrated transceiver, showing that multi-level signaling at binary transceiver rates is possible, especially with improvements in SNR through increased DAC bandwidth and lower timing errors and noise.

While more advanced communications techniques may result in higher SNR, and in particular may prove to be more tolerant of phase noise, implementation complexity limits their applicability to long links where wires are expensive. Similarly, multi-level modulation has promise for increasing data rates on long wires, but its complexity will likely delay widespread application until simpler techniques such as the use of low loss dielectrics, larger conductors, and wider busses have been pushed further. Results that could demonstrate the performance of a wire-loss-dominated link are not possible with the current transceiver chip due to increased phase noise at low frequencies. However, with less jitter, interleaved data converters enable sophisticated communications techniques to further increase CMOS link data rates.

A number of aggressive approaches are taken in this work, including multi-level modulation, single-ended signaling, high symbol rates, wideband receiver circuits, high resolution data converters, an aggressive technology file, a low fanout PLL capable of high clock rates, and the use of inductors to distribute parasitic capacitances. These aggressive approaches were taken to understand the tradeoffs and performance limits of high frequency CMOS circuits, but did result in some difficulties. Hopefully, the experience will guide future efforts toward low-risk and high performance designs.

Future research includes both circuits advances and link architecture work. The receiver can be improved with a more robust sampler biasing approach, by reducing residual offsets, or by adding resolution with a multi-step ADC architecture (provided that phase noise can be reduced). Transmitter clock coupling can be reduced, and the generation of high rate clocks with very low jitter remains open to improvement despite the volume of prior work. The high speed ADCs and DACs in the transceiver in this work can be used to experimentally verify advanced link algorithms, supporting research into link interference correction, application of Decision Feedback equalizations, and clock recovery algorithms.

Bibliography

Serial Links

- [1] S. Sidiropoulos, et al., "A 700-Mb/s/pin CMOS signaling interface using current integrating receivers," *IEEE Journal of Solid-State Circuits*, May 1997, vol.32, no.5, pp. 681-90
- [2] C.K. Yang, et al., "A 0.8- μm CMOS 2.5 Gb/s oversampling receiver and transmitter for serial links," *IEEE Journal of Solid-State Circuits*, Dec. 1996, vol.31, no.12, pp. 2015-23
- [3] C.K. Yang et al. "A 0.6- μm CMOS 4-Gbps Transceiver with Data Recovery using Oversampling", *Proceedings of 1997 IEEE Symposium on VLSI Circuits*, p. 71-72
- [4] W. Ellersick et al., "A Serial-Link Transceiver Based on 8-GSa/s A/D and D/A Converters in 0.25- μm CMOS", *2001 IEEE International Solid-State Circuits Conference, Digest of Technical Papers*.
- [5] W. Ellersick et al, "GAD: A 12-GS/s CMOS 4-bit A/D Converter for an Equalized Multi-Level Link", *Proceedings of 1999 IEEE Symposium on VLSI Circuits*, p. 49-52.
- [6] K. Chang, W. Ellersick, et al. "A 2Gb/s/pin CMOS Asymmetric Serial Link," *Proceedings of 1998 IEEE Symposium on VLSI Circuits*.

- [7] R. Farjad-Rad et al. "A 0.4- μm CMOS 10-Gb/s 4-PAM Pre-Emphasis Serial Link Transmitter", *IEEE Journal of Solid-State Circuits*, p.580-5, May 1999.
- [8] A.X. Widmer et al. "Single-chip 4*500-MBd CMOS transceiver", *IEEE Journal of Solid-State Circuits*, vol.31, no.12 p. 2004-14, Dec. 1996
- [9] A. Fiedler et al. "A 1.0625 Gbps transceiver with 2x-oversampling and transmit signal pre-emphasis", *Proceeding of 1997 IEEE International Solids-State Circuits Conference*, p. 238-9.
- [10] W.J. Dally et al. "Transmitter equalization for 4-Gbps signalling" *IEEE Micro*, Jan.-Feb. 1997. vol.17, no.1, p. 48-56
- [11] C.K.K. Yang, "*Design of High-speed Serial Links in CMOS*", PhD Dissertation, Stanford University, Stanford, CA, 1998.
- [12] K.Y.K. Chang et al., "A 50 Gb/s 32*32 CMOS crossbar chip using asymmetric serial links", *Proceedings of 1999 Symposium on VLSI Circuits*, p. 19-22.
- [13] C.K. Yang, et al., "A Serial-Link Transceiver Based on 8-GSa/s A/D and D/A Converters in 0.25- μm CMOS," *IEEE Journal of Solid-State Circuits*, Nov. 2001.
- [14] J. Zerbe, et al., "A 2Gb/s/pin 4-PAM Parallel Bus Interfaces with Transmit Crosstalk Cancellation Equalization and Integrating Receivers," *IEEE ISSCC Dig. of Tech. Papers*, Feb. 2001, San Francisco, pp. 66-7.
- [15] K. Donnelly et al., "A 660 MB/s Interface Megacell Portable Circuit in 0.3 μm -0.7 μm CMOS ASIC," *IEEE Journal of Solid-State Circuits*, vol. 31, no. 12, Dec. 1996.
- [16] T.H. Hu and P.R. Gray, "A Monolithic 480 MB/s Parallel AGC/Decision/Clock-Recovery Circuit in 1.2- μm CMOS," *IEEE Journal of Solid-State Circuits*, vol. 28, no. 12, Dec. 1993.

A/D Converters

- [17] A. Yukawa, et al., "A CMOS 8-bit high speed A/D converter IC," *1988 Proceedings of the Tenth European Solid-State Circuits Conference*, pp. 193-6.

- [18] W.C. Black Jr., D.A. Hodges, "Time interleaved converter arrays", *IEEE Journal of Solid-State Circuits*, Dec. 1980, vol. SC-15, no.6, pp. 1022-9
- [19] C.S.G. Conroy, et al., "An 8-b 85-MS/s parallel pipeline A/D converter in 1 μ m CMOS", *IEEE Journal of Solid-State Circuits*, April 1993, vol.28, no.4, p. 447-54.
- [20] A.M. Abo, P.R. Gray, "A 1.5-V, 10-bit, 14.3-MS/s CMOS Pipeline Analog-to-Digital Converter", *IEEE Journal of Solid-State Circuits*, May 1999, vol.34, no.5, pp. 599-606.
- [21] K. Poulton et al. "A 6-b, 4 GSa/s GaAs HBT ADC", *IEEE Journal of Solid-State Circuits*, vol. 30, no. 10, Oct. 1995.
- [22] B. Rasavi, "Design of Sample-and-Hold Amplifiers for High-Speed Low-Voltage A/D Converters", *IEEE 1997 Custom Integrated Circuits Conference*, p.59-65.
- [23] C. Mangelsdorf "A 400-MHz Input Flash Converter with Error Correction", *IEEE Journal of Solid-State Circuits*, vol. 25, no. 1, Feb. 1990.
- [24] R. Van de Plassche, *Integrated Analog-to-Digital and Digital-to-Analog Converters*, Kluwer Academic Publishers, 1994.
- [25] K. Azadet et al., "A Mismatch Free CMOS Dynamic Voltage Comparator", 1995 IEEE International Symposium on Circuits and Systems, p. 2116-19, vol. 3.
- [26] B. Razavi, *Principles of Data Conversion System Design*, IEEE Press, 1995.
- [27] S.-H. Lee et al. "Digital-domain calibration of multistep analog-to-digital converters", *IEEE Journal of Solid-State Circuits*, vol.27, no.12, Dec. 1992, p. 1679-88
- [28] B.-S. Song and M.F. Tompsett, "A 12-bit 1-Msample/s Capacitor Error-Averaging Pipelined A/D converter," *IEEE Journal of Solid-State Circuits*, vol. 23, no. 6, Dec. 1988.
- [29] R. van de Grift et al., "An 8-bit Video ADC Incorporating Folding and Interpolation Techniques", *IEEE Journal of Solid-State Circuits*, vol. 22, no. 6, Dec. 1987.

Circuits

- [30] T.H. Lee, *The Design of CMOS Radio Frequency Integrated Circuits*, Cambridge University Press, 1998.
- [31] M. Pelgrom et al, "Transistor Matching in Analog CMOS Applications", *IEDM 1998 Technical Digest*, p. 915-8.
- [32] M. Pelgrom et al, "Matching Properties of MOS Transistors", *IEEE Journal of Solid-State Circuits*, vol. 24, No. 5, Oct. 1989.
- [33] T. Mizuno et al, "Experimental Study of Threshold Voltage Fluctuation Due to Statistical Variation of Channel Dopant Number in MOSFET's". *IEEE Transactions on Electron Devices*, vol. 41, no. 11, Nov. 1994.
- [34] J. Montanaro et. al., "A 160-MHz, 32-b, 0.5-W CMOS RISC microprocessor". *IEEE Journal of Solid-State Circuits*, vol.31, no.11 p.1703-14, Nov. 1996
- [35] M. Shoji et al. "Elimination of process-dependent clock skew in CMOS VLSI". *IEEE Journal of Solid-State Circuits*, vol. SC-21, no. 5, p. 875-80 Oct. 1986.
- [36] B. Kleveland et al., "Monolithic CMOS distributed amplifier and oscillator," *1999 Proceedings of IEEE International Solid-State Circuits Conference*, pp. 70-1.
- [37] F. Grover, *Inductance Calculations*, Dover Publications, 1973.
- [38] H.O. Johansson and C. Svensson, "Time Resolution of NMOS Sampling Switches Used on Low-Swing Signals", *IEEE Journal of Solid-State Circuits*, Feb. 1998.
- [39] P. Gray and R. Meyer, *Analysis and Design of Analog Integrated Circuits, 3rd Edition*, John Wiley and Sons, 1993.
- [40] H.P. Westman et al., *Reference Data for Radio Engineers*, ITT, 1969.
- [41] C. Walker, *Capacitance, Inductance and Crosstalk Analysis*, Artech House, 1990.
- [42] E. Ginzton, W. Hewlett, et al, "Distributed Amplification," *Proceedings of IRE*, August 1948, pp. 956-69.
- [43] P.V. Argade, "Sizing an Inverter with a Precise Delay: Generation of Complementary Signals with Minimal Skew and Pulsewidth Distortion in

- CMOS,” *IEEE Transactions on Computer Aided Design of Integrated Circuits and Systems*, Jan. 1989, vol.8, no.1, pp. 33-44
- [44] I. Sutherland and D. Harris, “*Logical Effort: Design Fast CMOS Circuits*”, Morgan Kaufmann, 1999.
- [45] G. Knoblinger et al., “Thermal Channel Noise of Quarter and Sub-Quarter Micron NMOS FET’s”, *Proceedings of ICMTS*, March 2000, Monterey, CA.
- [46] A.J. Scholten et al, “Accurate Thermal Noise Model for Deep-Submicron CMOS”, *International Electron Devices Meeting, Technical Digest*, Dec. 1999.
- [47] J.-S. Goo et al, “An Accurate and Efficient High Frequency Noise Simulation Technique for Deep Submicron MOSFETs”, *IEEE Transactions on Electron Devices*, Dec. 2000.

Clock Generation

- [48] J.G. Maneatis, “Low-Jitter Process-Independent DLL and PLL Based on Self-Biased Techniques”, *IEEE Journal of Solid State Circuits*, Nov. 1996. vol.31, no.11, p. 1723-32.
- [49] S. Sidiropoulos, M.A. Horowitz, “A Semidigital Dual Delay-Locked Loop,” *IEEE Journal of Solid-State Circuits*, Nov. 1997, vol.32, no.11, pp. 1683-92
- [50] J. Poulton et. al. “A Tracking Clock Recovery Receiver for 4Gb/s Signalling”, *Hot Interconnects V Symposium Record*, August 1997, Stanford, 157-169
- [51] A. Hajimiri, *Jitter and Phase Noise in Electrical Oscillators*, PhD Dissertation, Stanford University, Stanford, CA, 1998.
- [52] J. Savoj, B. Razavi, “A 10Gb/s CMOS Clock and Data Recovery Circuit with Frequency Detection,” *IEEE ISSCC Dig. of Tech. Papers*, Feb. 2001, San Francisco, pp. 78-9
- [53] D. Weinlader et al., “An Eight Channel 36-GS/s CMOS Timing Analyzer,” *IEEE ISSCC Dig. of Tech. Papers*, Feb. 2000, San Francisco, pp. 170-1.

- [54] B. Razavi et al. "Design of high-speed, low-power frequency dividers and phase-locked loops in deep submicron CMOS" *IEEE Journal of Solid State Circuits*, Feb. 1995. vol.30, no.2, p. 101-9.
- [55] P. Roo et al., "A CMOS Analog Timing Recovery Architecture for PRML Detectors", *IEEE Journal of Solid-State Circuits*, vol. 35, no. 1, Jan. 2000.
- [56] K.H. Mueller and M. Muller, "Timing Recovery in Digital Synchronous Data Receivers", *IEEE Transactions on Communications*, p.516-31, May 1976.
- [57] T. Weigandt et al, "Analysis of Timing Jitter in CMOS Ring Oscillators", *Proceedings 1994 ISCAS*, p. 27-30, May 1994.
- [58] J. McNeill, "Jitter in Ring Oscillators", *IEEE Journal of Solid-State Circuits*, Vol. 32, No. 6, June 1997.

Communications

- [59] J. Proakis and M. Salehi, *Communication Systems Engineering*, Prentice Hall 1994
- [60] E. Lee and D. Messerschmitt, *Digital Communication, Second Edition*, Kluwer 1994.
- [61] W.J. Dally and J.W. Poulton, *Digital Systems Engineering*, Cambridge University Press, 1998.
- [62] J. Cioffi, "*EE 379A: Digital Data Transmission*", Stanford University, 1996.