A CMOS 4-PAM Multi-Gbps Serial Link Transceiver

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF ELECTRICAL ENGINEERING
AND THE COMMITTEE OF GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Ramin Farjad-Rad

August 2000

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.

Thomas H. Lee (Principal Advisor)

_____

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

Mark A. Horowitz

_____

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

Bruce A. Wooley

_____

Approved for University Committee on Graduate Studies:

_____

Abstract

As the demand for higher data-rate communication increases, low-cost, high-speed serial links using copper cables become attractive for distances of 1 to 10 meters. For multi-Gbps applications, the data rate is limited by the channel low-pass characteristics due to cable skin-effect loss and the speed of the process technology.

This design uses a receiver equalizer in combination with a transmitter filter to compensate for the channel low-pass characteristics. The target channel is a 10-meter copper coaxial cable with a -3dB bandwidth of <1GHz. A 4-level pulse amplitude modulation (4-PAM) is used to reduce the symbol rate to half that of the conventional 2-PAM system. This symbol rate reduction decreases the required channel bandwidth for proper data transmission by a factor of two and also allows the designer to halve the on-chip clock frequency. High on-chip frequencies are further lowered by multiplexing and demultiplexing the data directly at the chip pads. This design also takes advantage of a new proportional phase detector for 4-PAM data recovery at multi-Gbps rates. This PLL does not suffer from the stability and bandwidth limitations of traditional bang-bang loops

used in existing Gbps links. A novel frequency acquisition architecture is also designed to enable the receive PLL to lock to the input stream under all process variations.

A chip is implemented in 0.3-um CMOS to demonstrate the feasibility of very high-speed data transmission in standard CMOS technology using the above techniques. A data rate of 8Gbps is successfully generated by the transmitter and detected by the receiver after the 10-m coaxial cable. The chip occupies an area of 2mm x 2mm and consumes 1.1W at 8Gbps with a 3-V supply.

Acknowledgments

It all started in a weekend of November 1988, when my parents woke me up early in the morning to attend Iran's team selection competition for the international physics olympiad. Disliking any type of exam and any early wake-up call on a weekend morning, my initial response to their request was negative, but being a young teenager craving for driving, I could not reject their persistent request after their offer for letting me drive their car to the exam center. That day was followed by an unexpected success for me in the series of those competitions and several memorable months of olympiad preparatory camps with a number of very knowledgeable people, which all set a major turning point in the roadmap of my future. First and for most, I would like to thank my parents and all my teachers and friends in the olympiad camp for their constant encouragement and support, and especially for giving me the opportunity to realize that I can make a better difference in our world as a scientist, rather than my initial dreams of being a talented pilot, a popular athlete, famous physician, or a rich businessman.

I am also deeply indebted to many people who made my Stanford years an enlightening, rewarding, and memorable experience.

First, I would like to thank Professor Thomas Lee for granting me the honor of being a member of his research group and for his continuing support and guidance throughout this research. I have benefited from his vast technical expertise, insight, and his high-standards in research, teaching and publication.

I am grateful to my associate advisor Professor Bruce Wooley for his valuable advice, especially during my first months at Stanford, and his special effort to respond every student's need despite his vast involvements in the E.E. department. Many thanks to Professor Dutton for supporting me during my first year at Stanford and giving me the unique opportunity to search for my preferred field of research. During the Ph.D. qualifying exams, he was also a great source of encouragement and self-confidence for me that I will always remember.

I would like to express my special appreciation to Professor Mark Horowitz for creating a great excitement in me for the field of high-speed digital signaling and helped me with his superb technical supervision from the very first day I started this work.

My colleagues, Professor Ken Yang, Professor Behzad Razavi, Dr. Stefanos Sidiropoulos, and Kevin Yu deserve special recognition for offering their unconditional assistance throughout this project. I would also like to acknowledge Dr. Robert Drost with whom I started my first high-speed link project in SUN laboratories, and SUN microsystems for offering me the financial support to continue my research in this field. I also wish to thank Professor Shoba Krishnan, Lee Sampson, and LSI Logic for fabrication

of the experimental circuits. The financial support received from Powell foundation and LSI Logic is also gratefully acknowledged.

I am grateful to Ann Guerra for her kind assistance over the years and for being the best secretary of all times. So much of what I have learned in the past few years has been through interaction with many peers. All the members of SMIrC group, Horowitz students, and many other students in CIS have been helpful with their comments and friendship. In particular, I would like to thank Tamara Ahrens, Birdy Amrutur, Rafael Betancourt, Ken Chang, Bill Ellersick, Katayoun Falakshahi, Ali Hajimiri, Mar Hershenson, Mehrdad Heshami, Kambiz Kaviani, John Maneatis, S. S. Mohan, Hirad Samavati, Hamid Rategh, Derek Shaeffer, Arvin Shahani, Gu Wei, Evelina Yeung, Masoud Zargari.

There were many friends who made my years at Stanford the most wonderful moments of my life. I am exceptionally thankful to my dearest friends Tabassom, Philip, Nilofar, and Laleh for their unprecedented friendship and unconditional support at all times, and for filling me with hope and serenity in many hardships. I also owe most of my pleasant memories during these years to my friends Ali, Alireza, Amir, Amirali, Anita, Afshin, Arash, Ardavan, Azita, Babak, Eela, Fareed, Farnad, Farshid, Fati, Homan, Kamran, Katy, Keyvan, Kian, Mondana, Rasool, Reza, Salar, Samira, Salomeh, Serena, Shideh, Shahram, Siamak, Susan, Taraneh, Tiba, and many more.

I am deeply grateful for the endless and constant love and affection of my family members: Yasaman, Nilofar, and Afsar during so many years of separation. And last, but definitely not the least are my parents that I could not find the words to thank them.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Motivation

The exponential growth in speed and integration levels of digital integrated circuits (ICs) has increased the demand for high IC-to-IC communication bandwidths to maximize overall system performance. Traditionally, system designers have addressed this demand by increasing the number of high-speed signals, leading to an increase in the cost and complexity of the system. In order to maintain a balanced system, however, the per-pin interconnection bandwidth must scale with the speed and integration level of the IC.

There are two dominant approaches to high-speed signaling: shared buses and point-to-point links. In the conventional shared bus model, many links are integrated within a single system to increase the total signaling bandwidth. The large number of such links makes the overall overhead (area, power, latency) the key constraint of the design. These constraints dictate not only a shared bus approach, but also a simple link circuitry, which both limit the maximum bandwidth per channel. Shared buses have typically been

used for short-range interconnections within a single system and rates up to 100-200Mbps, such as in multi-processor systems ([79], [85], [86]), processor-to-memory interfaces ([88], [87]), and network switches ([89]). In contrast, the goal in high-speed point-to-point links is to maximize the communication bandwidth and distance on a single cable, where number of wires are restricted and high bandwidth is required. Therefore, the link circuitry needs to be more complex to meet the required performance. While this complexity can increase the aggregate system latency, this increase is of secondary concern, as the overall latency is usually dominated by the channel delay. Point-to-point links offer an excellent solution for applications that require multi-gigabit per second (multi-Gbps) rates over distances of several meters such as computer-to-computer or computer-to-peripheral interconnection [references], or several kilometers such as SONET and ATM.

High-speed serial signaling is also growing rapidly in the chip-to-chip interconnection domain [89], where designers of high-throughput chips are restricted by the limited number of package pins and printed circuit board (PCB) traces. The critical importance of this problem has resulted in some emerging standards such as Intel's and Compaq's *Infiniband* that advocates point-to-point high-speed signaling for chip to chip interconnection. These new growing needs make an integrated design of such links quite attractive, as it can increase the total bandwidth per channel and therefore the total chip throughput for a fixed package pin count.

Traditionally, high-speed links in the Gb/s range have been implemented in GaAs or bipolar technologies. The primary advantage provided by those technologies is faster intrinsic device speed (higher f_T), but they suffer from the main drawback of limited level

of integration. On the other hand CMOS, despite its slower device speed, is becoming the target technology for high-speed integrated systems due to its widespread availability and higher integration levels compared to the faster technologies. This availability makes high-speed links built in CMOS very attractive for large-volume applications that require such links. Furthermore with higher integration, links can be built as a macro-block in a single-chip system, allowing high throughput chips with low pin counts together with significant cost savings. Finally, the speed of CMOS technology is improving faster than the speed of other technologies, because of the extensive investment and momentum in CMOS technology developments.

There are also two distinct communication channels for serial links: optical fibers and copper cables. Optical fibers provide a large communication bandwidth over very long distances, but the fiber, necessary optical components and terminal electronics makes this approach costly and area-inefficient [90]. The optical links are the only solution for the gigabit/s communication over very long distances. Copper cables, on the other hand, are a much cheaper solution, but suffer from a very limited data bandwidth that decreases with length. Hence, copper links, which is the focus of this dissertation, are mostly used for short distance applications, such as system-to-system interconnections in server rooms.

This dissertation explores various communication techniques and circuit architectures for a low-cost serial link with an objective to maximize the operating bandwidth and transmission distance in a modest CMOS technology, over copper cables. This work demonstrates the possibility of building a 10-Gbps serial link in CMOS technology over 10 meters of copper cables.

1.2 Organization

Since choosing a suitable communication architectures is the first essential step in the design of a high-speed serial link, Chapter 2 discusses the trade-offs between different modulation, equalization and detection methods considering the specified transmission medium (copper co-axial cable) and the circuit limitations of CMOS process technologies at GHz speeds. The discussions lead to the selection of a communication scheme which uses a 4-level pulse amplitude modulation (4-PAM) and linear transmit and receive equalization.

The circuit architecture of the transmitter and receiver front-end that realizes the proposed system is presented in Chapter 3. Circuit techniques necessary to achieve the required performance with medium complexity are presented. The compensating transmit and receive FIR filters are designed in an analog fashion to reduce circuit complexity and extend speed capabilities of the process technology. Chapter 4 discusses the multi-phase clock generation and data recovery circuit. Since clock and data recovery affects the system bit-error rate (BER), which is a main performance metric, a new linear high-bandwidth data phase locked loop (PLL) is designed that keeps the clock jitter and offset low. This PLL is also capable of extracting the clock phase information from a 4-PAM data stream. Furthermore, the PLL's on-chip voltage controlled oscillators (VCOs) experience large frequency variations due to the large process variations. As the data-recovery PLL does not correct for the loop frequency error, a novel frequency acquisition technique is invented to adjust the VCO frequency to the data frequency at start-up.

Chapter 5 is devoted to the measurement results from two prototype links fabricated in 0.4-μm and 0.3-μm CMOS process technologies. The test setup and the precautions necessary for the high-frequency characterization of the transceiver are explained in this chapter. The conclusion of this dissertation and the future work in this field to expand the bandwidth of serial links further is the subject of Chapter 6.

Chapter 2

Limitations and Architecture of the Signaling System

Very high data rates in binary signaling have been demonstrated in CMOS technology (e.g. [10] and [11]), where *point-to-point incident-wave* signaling makes it possible to achieved bit-times (actually, symbol periods) less than 250ps over electrical interconnections. The basic speed limitations of these links are set by the maximum operating frequencies intrinsic to the transistors and the transmission medium bandwidth. The on-chip signaling rate is expected to scale up with the transistor speed, and thus should follow a Moore's law curve. However the interconnection bandwidth, which improves at a much slower pace, is becoming the dominant bottleneck for the signaling speed. In order to overcome these limitations, this work investigates two areas: 1) dealing with the impact of the bandwidth limitation of the medium (specifically, copper coaxial cable), and 2) transmission of more complex symbols to increase the number of bits per symbol.

This chapter begins in Section 2.1 by examining the limitations of conventional CMOS signaling and showing how they are overcome by the point-to-point low-swing incident-wave signaling used in this work. The characteristics of the electrical interconnections, the copper cables, are illustrated in Section 2.2. Section 2.4 proposes two equalization techniques to better utilize the available bandwidth of wires at multi-Gbps rates. This work also takes advantage of 4-level pulse amplitude modulation (4-PAM) to reduce the bandwidth of the signal, while keeping the data rate constant, in order to relax the bandwidth limitations of both process and channel. The modulation scheme and simple coding used to improve system performance, are the subject of Section 2.5. In order to alleviate the on-chip frequency requirements to permit signaling at high symbol rates, this design multiplexes the on-chip low-speed parallel data onto a single high-speed output stream and demultiplexes the high-speed input stream back to low-speed parallel data at the chip pads. Section 2.3 briefly describes this parallelizing architecture, followed by detailed implementation discussions in Chapter 3

2.1 Signaling Techniques

This section describes two different approaches to high-speed signaling in digital systems. The first method is traditional high-swing signaling, such as TTL or CMOS, which has been used in most computer systems over the past several years, especially for chip-to-chip communication. This conventional method has been inherently limited to maximum frequencies of ~100MHz and these frequencies have not scaled with improving process technology. As the speeds in the modern digital systems increase, this conventional method of signaling is therefore becoming a major bandwidth bottleneck,

and system architects are looking into techniques to avoid such shortcomings. The second scheme discussed in this section, point-to-point incident-wave signaling, does not suffer from the inherent limitations of the conventional method, and thus its data rates can scale with the process technology. Because of this scaling, this new signaling technique is emerging in high-speed systems, e.g., network switches and backplanes.

2.1.1 Traditional large-swing signaling

Traditional signaling systems are limited to data rates of about 100Mb/s per wire or less and dissipate large amounts of energy per bit transmitted. More important, traditional signaling rates have not scaled with improvements in process technology. Because of this limitations, many modern microprocessors operate their external buses at a small fraction of the internal clock rate.

In traditional CMOS systems, a CMOS inverter is used both as the driver and receiver. The transmission medium, typically a cable or PCB trace, has a characteristic impedance of about 50Ω to 100Ω. The driver has an output impedance of ~300Ω and the line is unterminated at the receiver. The two power supply voltages (Vdd and Gnd) are used to represent logic 1 and 0 respectively.

These signaling systems are slow because the high impedance driver is unable to switch the line voltage completely on the incident wave [12]. Instead, the driver must charge up the line as the signal propagates over several round trips, as shown in Figure 2.1.

The figure shows the voltage at the far end of the line as a function of time. The incident wave only switches the line to 15% of the final value, which is doubled (30%) after the first reflection at the far end. As seen in this example, a minimum of 11 traversals

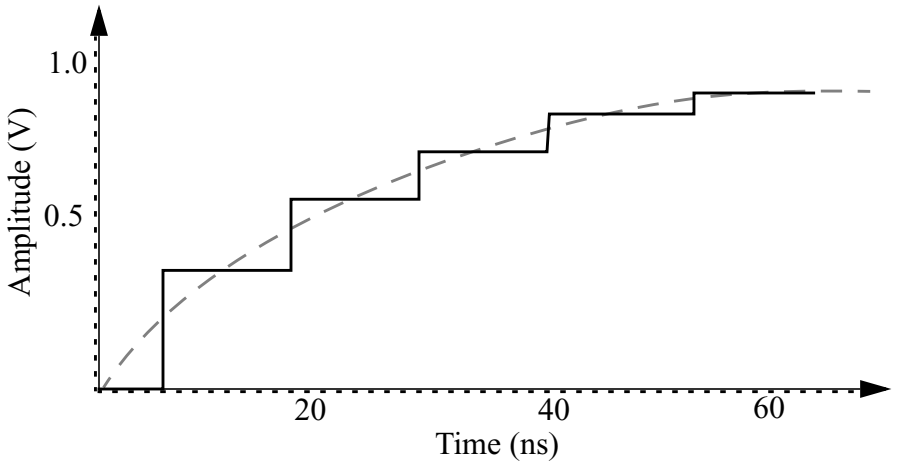

Figure 2.1: A CMOS inverter applies a logic level of 1 to a 4-ns line: Waveform at the far end of the line

are required before the line settles to within 10% of its final value. This effect also causes the maximum allowable length of the interconnections to decrease linearly as the data rates increase, even in a lossless wire.

The traditional CMOS system is power hungry at 100-200Mbps rates, dissipating 1nJ or more to transmit each bit[1], because its line is not terminated or clamped to some intermediate voltage and it has to use rail-to-rail signal swings for transmission. Also due to its parallel single-ended signaling approach, the system typically derives transmit and receive voltage reference from noisy power supplies[2]. To overcome the large noise of the reference, a large input signal swing is required by the receiver.

An additional problem is that traditional systems typically use the interconnections as a multi-drop, shared bus. As the signaling speeds on the lines begin to increase, these taps off of the line start to appear as discontinuities. Signal transitions are reflected from

1. Assuming a total line capacitance of ~100pF and a typical CMOS voltage of 3.3V, the energy dissipated to charge up the line for each bit change is $E=CV^2= 1nJ$.
2. The voltage reference for traditional CMOS signaling is the inverter's threshold voltage that is a function of PMOS to NMOS strength ratio and supply voltage.

these discontinuities and interfere with other transmitted symbols in the line. This effect can frustrate attempts to signal at rates over 100Mbps.

2.1.2 Point-to-point, low-swing, incident wave signaling

A signaling system that overcomes the limitations of traditional signaling is shown in Figure 2.2. A current-source transmitter drives the line with currents that typically range

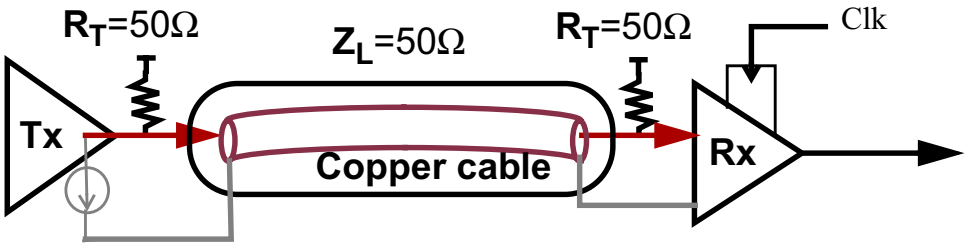

Figure 2.2: A point-point, low-swing, incident-wave system

from a couple of milliamperes to tens of milliamperes, resulting in a voltage swing range of 100mV to about 1V. The line is terminated at both ends in its characteristic impedance. The receiver termination absorbs the incident wave, preventing any reflections. The source termination makes the system more tolerant of crosstalk and impedance discontinuities by absorbing any reflected wave. A high-gain clocked regenerative receive amplifier can possess both low offset (~30-60mV) and high gain. The input sensitivity is determined in part by the regeneration time of the amplifier and increases exponentially with this time.

Taking advantage of the improved receive detection with low offset and high sensitivity, and also coupling the transmit and receive reference tightly to the signal, this

system can operate reliably using very small voltage swings. Therefore this signaling method also offers a considerable reduction in power dissipation of the system.

The system described can also operate at data rates independent of the line length. A new symbol can be driven onto the line before the previous symbol arrives at the receiver. This results in a system whose data rate, to a first approximation, scales linearly with the device speed.

2.2 Limits of Electrical Signaling

The data rate of a signaling system is limited by both the electronics used to generate and receive the signal and the medium over which the signal propagates. The following two parts examine more closely these two limitations.

2.2.1 Electronic limitations

As illustrated in Figure 2.3, limits on signaling rate arise from rise time, aperture

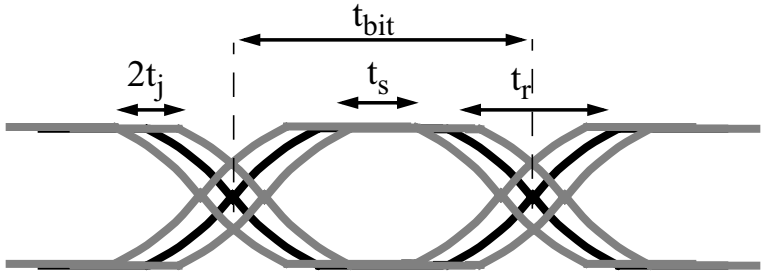

Figure 2.3: Eye diagram showing limitations on signaling rate

time, and timing uncertainty [12]. The time for the bit cell, t_{bit}, must exceed t_r (the time required for the signal to slew from one logic level to another), summed with t_s (the time

for the receiver to sample this signal while stable), and $2t_j$ (the jitter between the signal and the sampling clock).

As the maximum electronic limits scales with the speed of devices, it varies with process technology, temperature, and supply voltage. Thus, it is preferable to use a metric for bit rate that is independent of technology and operating conditions, such that the performance number stated can be compared with other systems. An adequate metric in CMOS technology is the delay of an inverter driving a capacitive load (fan-out) that is four times larger than its input capacitance (FO4). The following paragraphs relate all these factors to the FO4 delay of the process technology. For a typical 0.35μm CMOS process, FO4=150-200ps, and this delay time scales linearly with gate length.

In the current mode signaling scheme, shown in Figure 2.2, the transition time of the signal, t_r, is determined by the time to turn on a current source completely. Using a current source switched by logic gates it is easy to generate a slew time of 0.6FO4 delay, or ~100ps, in a 0.35μm process. The RC time constant of the terminated line (50Ω||50Ω=25Ω) and the capacitance of the driver (~1pF) is smaller than the signal risetime, effectively not limiting the minimum rise time. Hence, the signal rise time still scales up with process improvements (FO4 delay reduction).

The *aperture time*[1] of the receiver can be made as small as 0.2FO4, using a single transistor passgate [31], or gate-isolated sense amplifier [12]. The peak-to-peak jitter is caused by jitter in the sampling clock, jitter in the receiver clock, and delay variation in the signal path. In most practical systems, these three sources of jitter are primarily due to

1. Aperture time of a receiver is the minimum time interval during which the receiver can sample the input data signal successfully (for more details see Chapter 3).

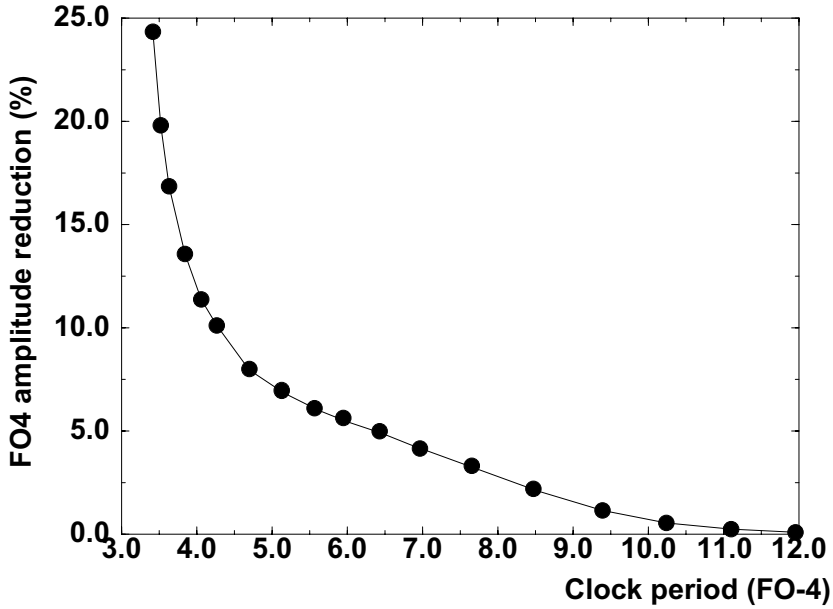

Figure 2.4: Clock amplitude reduction (%) with clock period (in FO-4 delays).

power-supply modulation of delay, and crosstalk-induced delay variation. Both of these factors scale with delay and hence with process gate length. Using circuits with good power supply rejection, and employing layout techniques to avoid crosstalk, Sidiropoulos in [98] shows a peak-to-peak jitter of less than 0.3FO4 is achievable. Of course, factors such as inherent device noise (e.g. thermal and flicker noise) sets the lower limit for jitter.

In addition, in a basic link design as discussed in [32], the maximum data rate has a practical limit set by the on-chip clock period, which in turn is determined by the bandwidth of the clock buffer chain. For minimum propagation delay, a fan-out of roughly four is often used at each stage of the buffering. Because of the limited bandwidth of FO-4 buffers, propagating a clock with frequencies above the bandwidth causes the clock amplitude to reduce, thus imposing a limit on the maximum clock frequency. Figure 2.4 illustrates the reduction of the clock amplitude as the frequency is increased. Therefore, to maintain a reasonable clock amplitude, the clock period is constrained to roughly 6 FO-4.

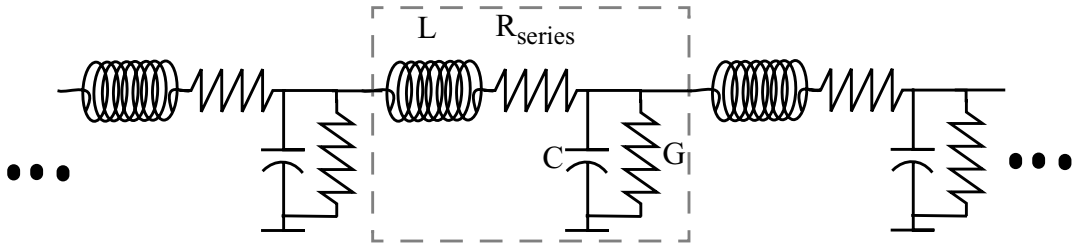

Figure 2.5: lumped LCRG model of a transmission line

However, parallel multi-phase I/O architecture, as described in Section 2.3, can generate and sample multiple bits in one clock period, relaxing the clock frequency constraint on the link bit rate.

Thus, with a careful design, the electronics for an incident-wave signaling system can achieve bit times of 1-2 FO4 delay. More important, because there is no need to charge up the signal line, the signaling rate scales with the technology, independent of the wire[1] length. Unfortunately, for most systems this bit rate exceeds the bandwidth of the transmission medium (e.g. wires), which is discussed next.

2.2.2 Transmission medium limitations

PC board traces and coaxial, twisted-pair cables behave as transmission lines that store and propagate signal energy. These lines can be modeled by a series of lumped *LCRG* elements as shown in Figure 2.5. The loss in transmission is primarily due to the series resistive component of the copper (*R*) and parallel conductive component of the dielectric (*G*). For wires with imperfect shielding or coupling, signal energy can also be

1. Assuming a lossless wire with a flat frequency response.

lost due to radiation. However, this loss is negligible compared to the two latter types, for the limited range of distances used in this work (~10-m coaxial cable or ~1-m PCB trace).

The major frequency-dependent loss in any electrical link using conductors is due to skin-effect resistance. This effect can be modeled as an increase in the series resistance of the wire as shown in Figure 2.5. Higher frequency signals propagate closer to the surface of the conductor; therefore, the resulting signal current conducts within a limited depth, the skin depth, on the conductor surface, which is defined as [12]:

$$\delta = 1 \Big/ \left(\sqrt{\frac{\pi \cdot \mu_o}{\rho} \cdot f} \right) \qquad [2.1]$$

where ρ is the resistivity of the conductor, and f is the frequency of the signal. Hence, the effective series resistance (R_{skin}) of the cable corresponding to this depth increases with square root of frequency [72]:

$$R_{skin} = \frac{1}{2 \cdot r} \cdot \sqrt{\frac{\rho \cdot \mu_o}{\pi} \cdot f} \quad , \qquad [2.2]$$

for a round conductor with radius ***r***. This relation shows a frequency dependent resistive loss. Note that the above equation is valid only for frequencies well above the skin frequency, f_{skin}, where the skin depth is smaller than the radius of the wire:

$$f_{skin} = \frac{\rho}{(\pi \cdot \mu \cdot r)}, \qquad [2.3]$$

for a round conductor. Well below this frequency, the current density is nearly uniform throughout the cross section of the conductor, and the wire resistance is close to its DC resistance.

With some insulating materials, dielectric absorption also causes frequency dependent attenuation. This loss can be modeled as a conductance *G* between the signal wire and ground. This effect can be mitigated by using a low-loss dielectric material. However, because of certain restrictions on PC board materials, the choice of dielectric is limited; thus PCB traces usually demonstrate a considerably higher dielectric loss compared to cables. Dielectric loss for each material is usually expressed in terms of a parameter, called the loss tangent defined as [12]:

$$\tan\sigma_D = \frac{G}{\omega \cdot C}, \qquad [2.4]$$

where C is the capacitance per unit length as shown in Figure 2.5. This quantity is approximately constant with frequency, therefore the dielectric loss, G, typically increases linearly with frequency. Based on the above equations, the approximate cable frequency response in dB accounting for both skin effect and dielectric absorption is given by [12]:

$$H(f, l)\big|_{dB} = -(h_s \cdot \sqrt{f} + h_d \cdot f) \cdot l, \qquad [2.5]$$

where l is the length of the cable, h_s and h_d are the skin-effect and dielectric loss coefficients respectively. For most wires (e.g. coaxial or twisted pair cables), h_s is considerably larger than h_d, so dielectric loss may be neglected in them. For most of the PCB materials h_s is marginally larger than h_d, so it needs to be considered.

Figure 2.6 shows the frequency response of a copper cable, with both skin-effect and dielectric loss components shown separately. At lower frequencies, skin effect is the

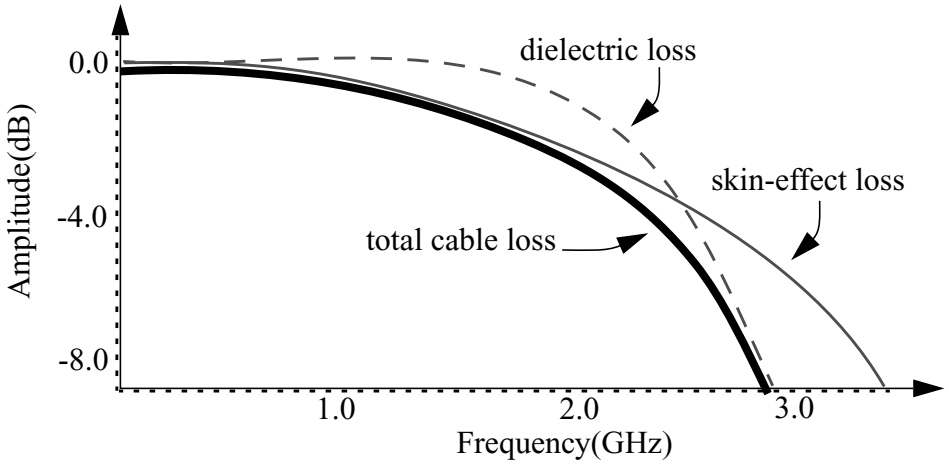

Figure 2.6: Frequency dependence of cable loss

dominant loss; however, dielectric absorption, having a larger slope, dominates at higher frequencies.

Figure 2.7 shows the time-domain response of a cable to an ideal trapezoidal pulse. The pulse suffers more attenuation of its high frequency components (transitions) than that of its low frequency components. The attenuation above the pulse frequency causes the output pulse amplitude to be reduced by more than 4dB. Moreover, lower frequency attenuation results in the long settling tail exhibited by the signal. This long settling tail results in ISI that corrupts the future transmitted symbols and reduces their effective width and amplitude. Because of this effect, transmission without equalization is usually limited to frequencies where the narrowest transmission pulse of the system experiences an attenuation of 2dB or less. Hence, unequalized transmission requires extremely high quality cables, which increase the overall system cost.

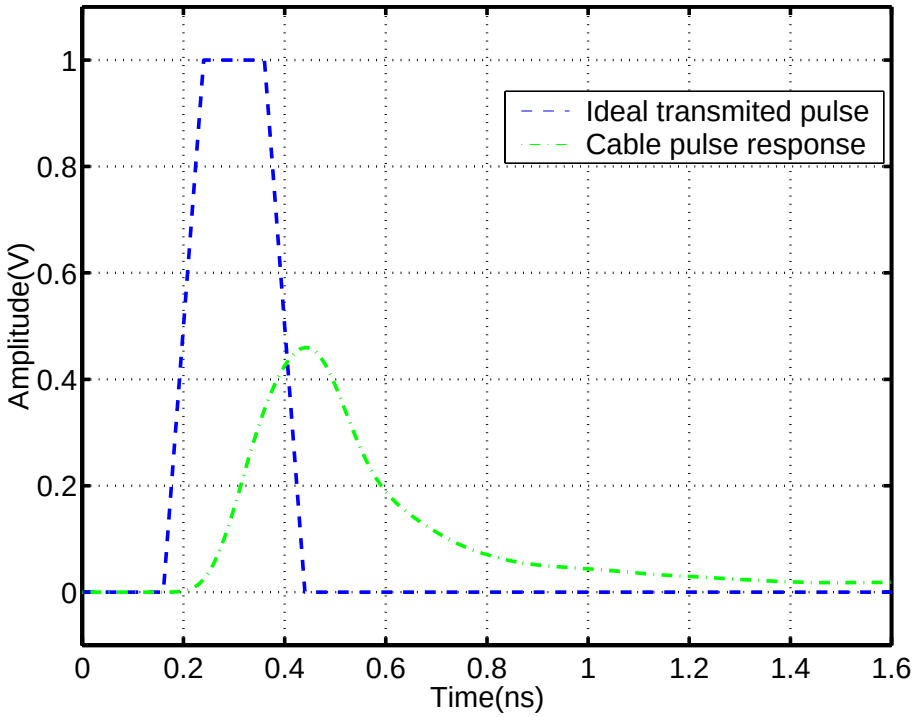

Figure 2.7: Pulse response of a 10-meter coaxial cable

2.3 Parallel I/O Architecture

Since the maximum clock speed on a chip is limited, high data-rate systems transmit more than one bit per clock cycle. In the simplest and most common system the transmitter multiplexes two bits in each cycle, one on each edge of the clock (2:1 multiplexing). The receiver in such a system performs 1:2 demultiplexing on the serial stream prior to the digital processing so that the digital logic can operate at a lower rate than the off-chip bandwidth (e.g. [87] and [89]). This low-degree parallelism improves the bit-time to roughly 3 FO-4. At this point, factors such as the bandwidth of the multiplexer and the minimum cycle time of the receiver emerge as bandwidth-limiting factors. In addition, since both rising and falling edges of the clock are used to transmit data, a duty-cycle error can further degrade performance. A transmitter clock with duty-cycle error

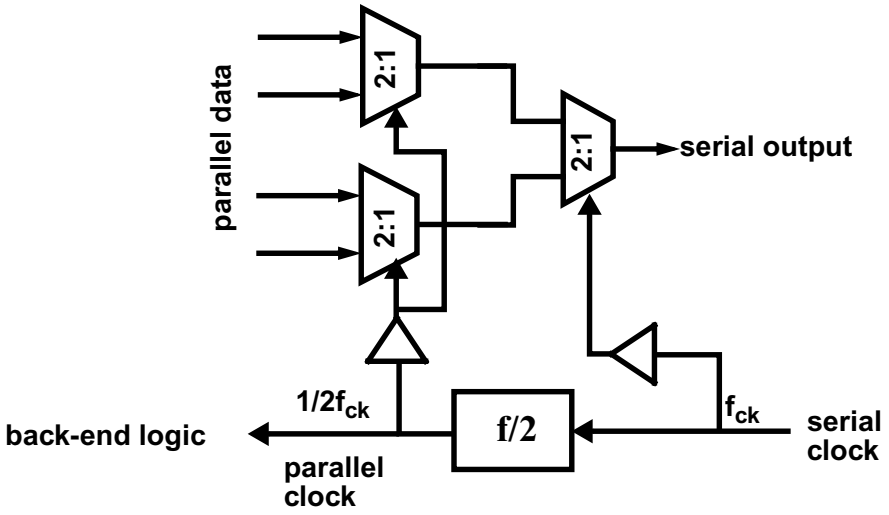

Figure 2.8: Tree type 4:1 multiplexing using 2:1 multiplexer

causes unequal bit-widths for the odd and even data bits. With duty-cycle error in the receiver clock, both receivers cannot simultaneously sample at the middle of the data-eye for both edges, causing a static phase offset for one or both data bits. Links that use 2:1 parallelism (e.g. [67]) typically use additional circuits that perform duty-cycle correction.

A natural strategy to reduce the bit-time further is to increase the degree of parallelism. One common approach is to increase the number of 2:1 multiplexing levels to achieve 2^n:1 multiplexing. Figure 2.8 shows an example of this technique for 4:1 multiplexing that serializes the parallel data from the lower rate to the data rate. One of the main problems with this approach is that synchronization of the additional levels, operating at different clock rates and delays, imposes timing constraints that can limit the cycle time. The other limitation of this scheme is that the multiplexing is restricted to only powers of 2. Therefore, in certain systems that require other multiplexing ratios (e.g. 10:1), when using 8b/10b codes, one needs to add an extra complex stage to perform *N:M* (e.g. 8:10) multiplexing and demultiplexing [97].

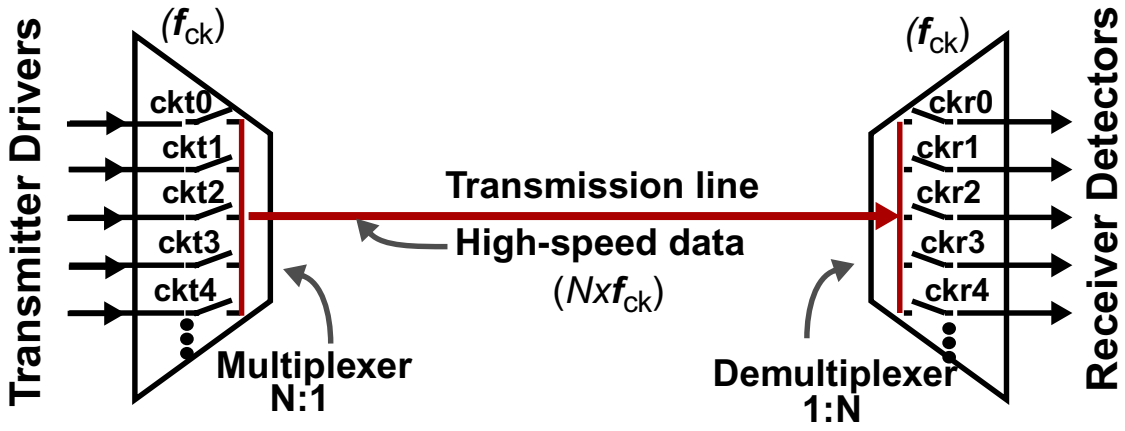

Figure 2.9: Higher order multiplexing, (a), and demultiplexing, (b)

This dissertation uses the approach originally proposed by Yang [10], where a higher degree of parallelism is achieved by using an N:1 multiplexer and 1:N demultiplexer for the transmitter and receiver (Figure 2.9). The on-chip clock rate can be reduced to 1/N of the data rate, thus enabling higher data rates.

In the transmitter, a larger fan-in multiplexer can be used with each input selected by consecutive phases of clock. To apply parallelism in the receiver, multiple clock phases are used with a different phase for each sampler/comparator, Figure 2.9. Each phase of the lower frequency clock samples a different data bit. Because of the longer clock period, the minimum cycle-time of a sampler/comparator is no longer a limitation.

In this technique, one can achieve bit times of 1 FO4 (or less) without the timing limitations involved in the basic link or multi-level multiplexing design. Also, the multiplexing ratio (N:1) can be arbitrarily selected based on the system requirements. Even though the greater parallelism allows smaller bit times, three performance

limitations still exist: the multiplexing and demultiplexing nodes have greater capacitance which limits the bandwidth, the select signals of the multiplexer switches need to be short pulses with pulse-width on the order of one bit-time, and the receiver multiple sampling clock phases need to be precisely spaced from each other to minimize the sampling phase error.

The timing of the select pulses for the transmitter and the clock edges for receiver are critical to properly transmitting and receiving the data bit. The variation in device parameters such as Vt will lead to a fixed mismatch in the delay elements used to build the timing circuits and will result in fixed-pattern phase error (effectively jitter). Fortunately, it is possible to measure these fixed errors and compensate for them. In this design, perfectly spaced clock phases are assumed to be supplied by a multiple clock phase generator. Chapter 3 discusses the circuit design details of this architecture, and how much the data rate will improve despite the high capacitance at the input and output nodes. It will discuss different methods of generating narrow select pulses for the transmit multiplexer and also the impact of non-ideal multiple clocks phases on actual system performance.

In this design a multiplexing ratio of 5 is selected. Therefore, in each clock cycle 5 symbols are transmitted, allowing reduction of the on-chip clock frequency to 1/5 of the symbol rate. As each 4-PAM symbol carries two bits of information, 10 bits of data (5 symbols) are transmitted per clock period. Therefore, the effective bit rate will be an order of magnitude larger than the on-chip clock frequency, or the bit time is reduced to about 0.5 FO4.

2.4 Cable Equalization

The simplest way to achieve higher data rates is to actively compensate for the channel transfer function. For multi-Gbps transmission channels, analog equalization is an attractive alternative to digital approaches. The main advantages of analog equalization are lower power, less silicon area and, most important, capability of performing equalization at high speeds.

Unfortunately, analog compensation of the channel frequency response comes at the expense of reduced signal to noise ratio (SNR). Analog compensation for the channel low-pass effects attenuates lower frequency components of the signal (the amplitude of the wider pulses) to match the attenuation at the higher frequency components, so that the resulting channel frequency transfer is flat within the signal range. This process results in an attenuation in the overall signal power, but most of high-frequency noise sources are not affected. Therefore for systems wih analog equalization, careful system design should be perform to minimize such deterministic high-frequency noises such as crosstalk and supply induced noise.

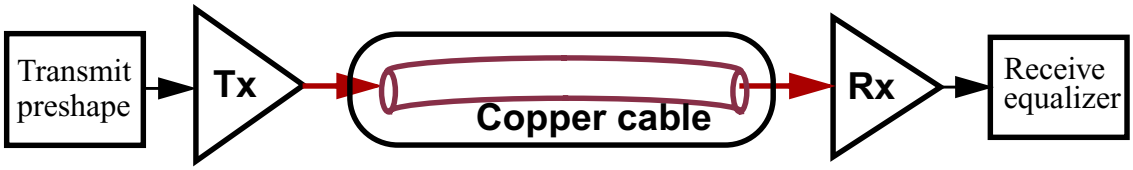

Figure 2.10: Two methods of compensating for channel response

As shown in Figure 2.10, channel compensation can be performed either through a predistortion of the driven signal by the transmitter, or by an equalization filter at the receiver. This work takes advantage of both a transmit preshaping filter and a receive equalizer. The following two sections describe these two filtering techniques as used in this work.

2.4.1 Transmit filter

Transmitter preshaping uses a symbol-spaced finite impulse response (FIR) filter integrated into the line driver, specified by the following equation:

$$V_o(n) = V_i(n) - a_1 \cdot V_i(n-1) - a_2 \cdot V_i(n-2) - \ldots\ldots \quad [2.6]$$

The inputs to this filter are the present and past transmitted symbols. The coefficients of the filter depend on the channel characteristics. The length of the filter *N*, (i.e. number of filter taps), depends on the number of symbols in the past that affect the present symbol. Therefore, the output of the filter is the present symbol value plus the weighted values of the past *N* symbols. Effectively, the FIR filter output, which is the sum of the present symbol and weighted values of *N* former symbols, will no longer have the distinct signal levels of the unfiltered signal stream, and the output driver in fact becomes a digital to analog (DAC) converter that operates at the symbol rate.

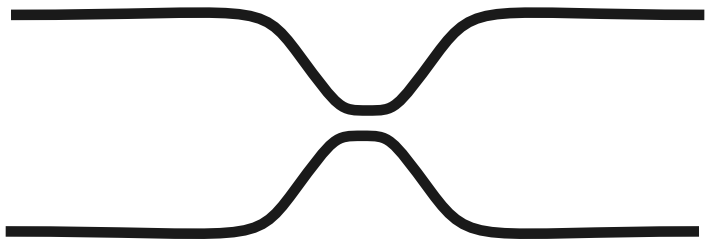

(a)

(b)

Figure 2.11: An isolated pulse in the field of zeros at the end of channel: a) Without pre-emphasis, b) With 1-tap pre-emphasis

In the simplest case, the FIR filter effectively suppresses the power of low-frequency components by reducing the amplitude of continuous strings of same-value data on the line. At the same time, it keeps the power of high-frequency components constant by increasing the signal amplitude during transitions. Figure 2.11 shows an example of the effect of a single tap transmit filter (a2,3....=0) on an isolated pulse in the field of zeros at both ends of the cable. The undistorted waveform on the top shows zero eye opening. With predistortion, as shown on the bottom, a clean eye opening is visible.

The main disadvantage of this approach is that the transmitter uses part of the signal amplitude budget to generate the preshaping symbols following the main symbol. Therefore, the actual transmitted signal amplitude is:

$$Signal\text{–}amplitude = Total\text{–}amplitude \bullet (1 - |a_1| - |a_2| \cdots\cdots), \qquad [2.7]$$

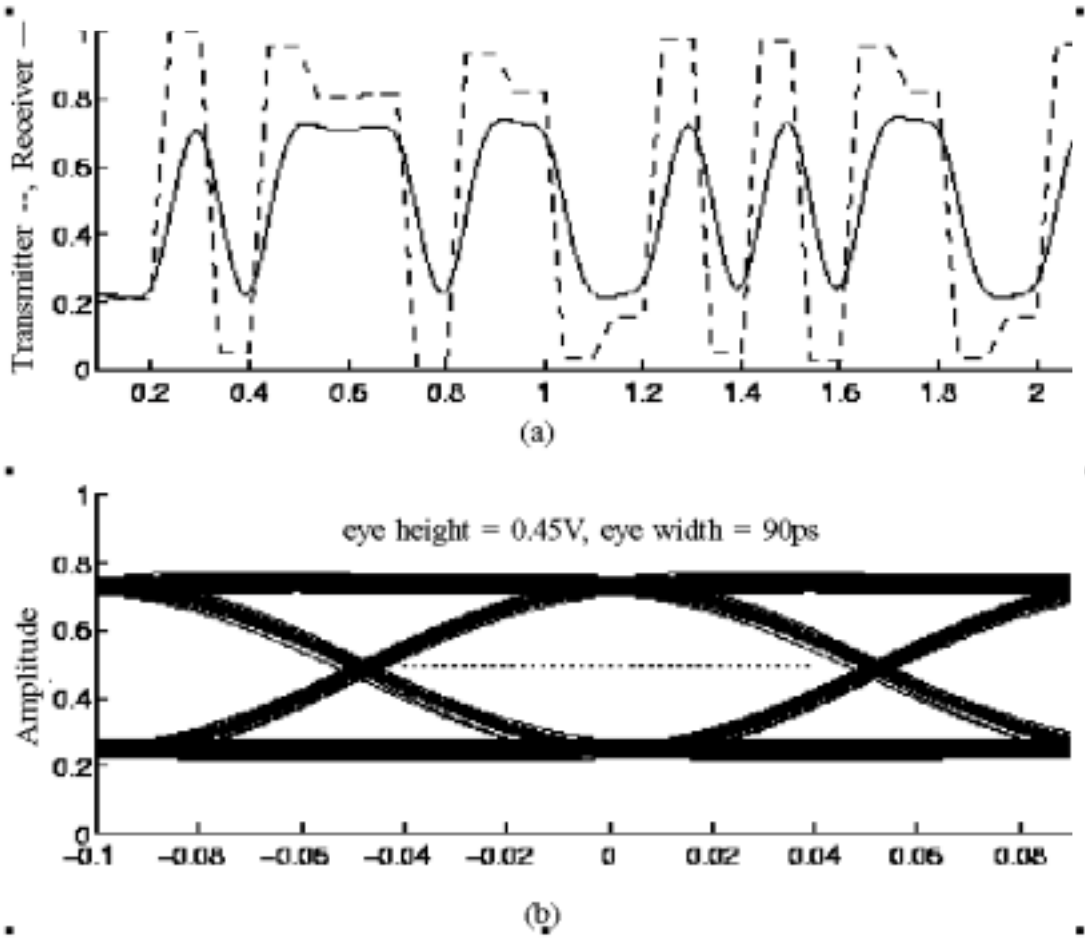

Figure 2.12: a) Effect of 1-tap pre-emphasis on transmitter and receiver, b) Received eye diagram

which effectively reduces the received signal amplitude, and therefore decreases the signal to noise ratio. Figure 2.12 shows transmit preshaped waveforms (dashed line) versus the synchronized received signal (solid line), and the received eye diagram in the lower plot. This illustrates the reduction in received eye height due to transmit pre-emphasis even for a channel with small loss. The transmitter has a nominal range of 1V where 0.4V of it is used for the preshaping (0.2V top and 0.2V bottom), thus the received signal only has a range of about 0.5V after further attenuation in the channel.

One of the important advantages of a transmit preshaping filter is the ease of implementing long FIR filters at very high speeds using this scheme. The reason is that all the previously transmitted symbols are already available to the transmitter. It is only necessary to add or subtract the weighted values of these known symbols from the present symbol value and send the result to the fast transmitter DAC. Hence, the preshaping technique does not require that the transmitter use a faster technology to operate properly. However, because the whole process is performed in the transmitter, the filter generally does not have any information about the signal shape at the receiver end of the cable. Determining the optimal FIR filter coefficients adaptively therefore can be challenging. The systems using this technique usually either need some feedback information from the receiver about the signal shape or require the user to characterize the channel in advance and set the taps manually.

The first implementation of a *N*-tap transmit filter was implemented by Dally and poulton [15]. In this architecture, all the FIR filter calculations are done by digital adders and a digital-to-analog converter (DAC) generates the output pulse. The architecture of this approach is shown for a *5*-tap filter in Figure 2.13. Driver parallelism (see Section 2.3) is used to overcome the speed limitations of the process, and therefore each branch requires separate digital logic. The digital logic modules, having the present and the five previous bits, calculate the amplitude of the current pulse that should be transmitted. In this scheme, each driver module is a 6-bit DAC to reduce the quantization error of the FIR filter. Therefore, the combination of these digital filters and high-resolution DACs operating at high-frequency makes the transmitter circuit quite complex and power

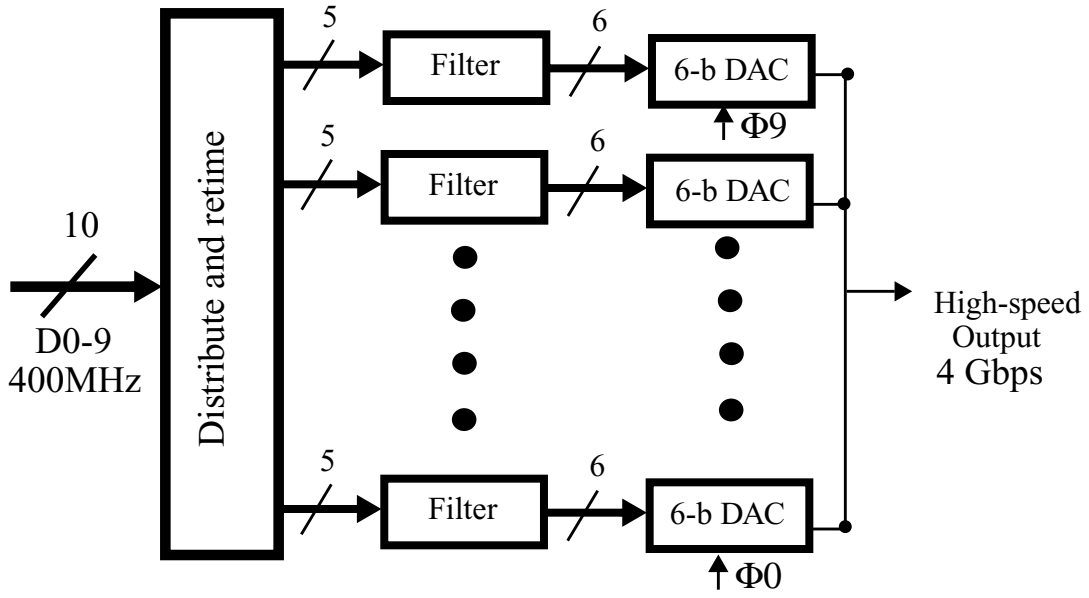

Figure 2.13: Block diagram of a preshaping parallel transmitter

hungry. Moreover, this complexity can be limiting when using multi-level schemes. For example, in 4-PAM modulation all the digital logic will be doubled, and each DAC should be modified to have at least 7 bits of resolution.

The design described in this thesis uses parallelism as well, but with a completely analog technique to realize the pre-emphasis filter. In this technique, the FIR *N*-tap filter is integrated into each driver module of the parallel transmitter. Each module uses the present and *N* previous bits to generate the complete preshaped pulse over *N*+1 bit periods directly and independently of all previously transmitted symbols. Finally, the individually preshaped symbols are summed at the transmitter output in an analog fashion. The top view of this scheme is shown in Figure 2.14.

The need for the complex digital filter is removed as weighing and summing the input filter values is performed by modulating (multiplying) and summing the analog

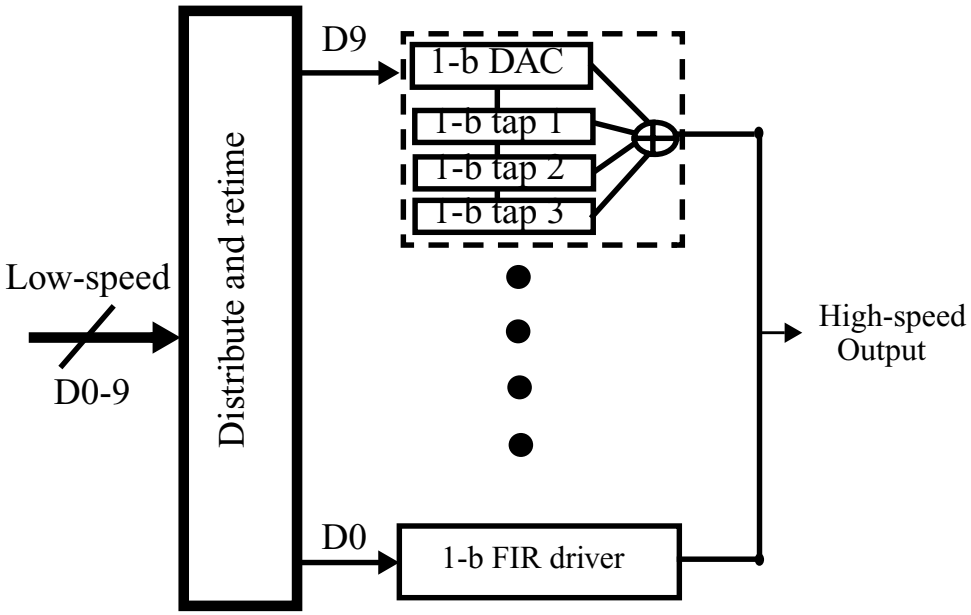

Figure 2.14: Block diagram of the proposed preshaping transmitter

currents in the analog domain. The high-resolution DAC drivers are also replaced by 1-bit drivers (or $log_2(M)$-bit drivers if using *M*-PAM) for each filter tap. To define the tap weight, the *N*-tap FIR filter requires only *N* (one for each tap) fairly low-resolution DACs that are shared among all similar filter taps in each parallel branch. Therefore for a 3-tap filter, one requires only 3 low-resolution DACs, independent of the number of parallel branches. More details of this architecture are discussed in Chapter 3.

High-level simulations of the *N*-tap FIR preshaping filter over a 10-m coaxial cable possessing a -3dB bandwidth of ~1GHz, were performed to find the optimal number of filter taps and their values. These simulations showed that, for 5Gsym/s signaling (200ps symbol width) over the chosen cable, a FIR filter with just 3 taps can successfully cancel the long tail of the cable pulse response to less than 3% of the final peak amplitude. To illustrate the effect of this 3-tap preshaping transmitter for 5Gsym/s signaling, Figure 2.15

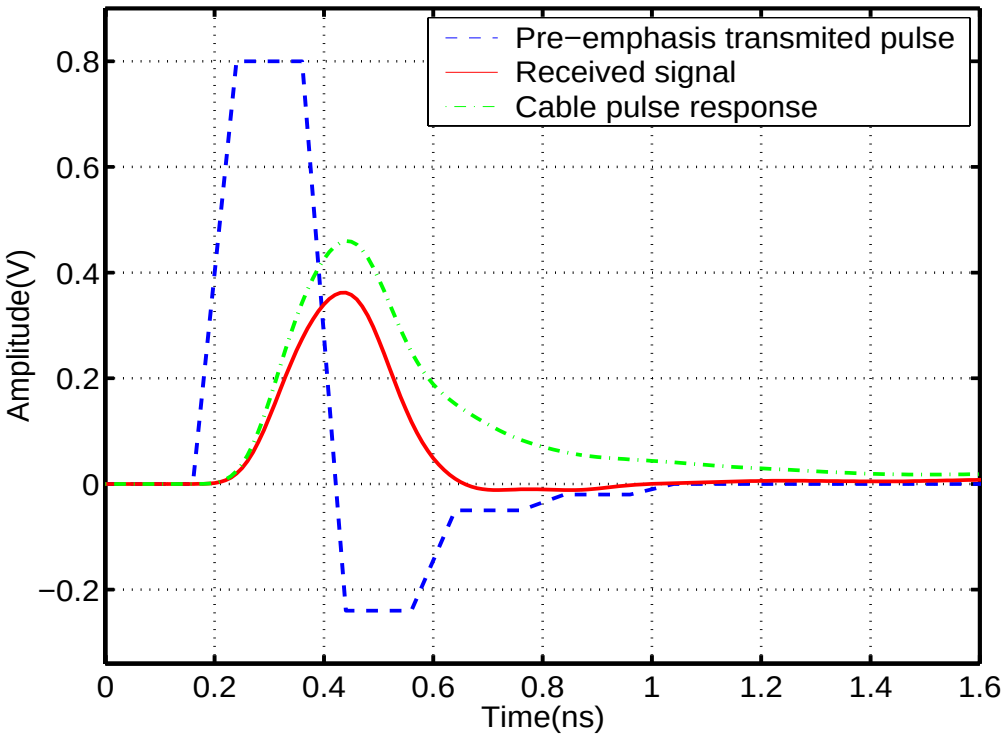

Figure 2.15: Effect of 3-tap transmit filter at the receiver end of cable

shows a preshaped 200-ps pulse, and the resulting pulse at the end of the 10 meter coaxial cable under test as compared to the original pulse. The result of this preshaped pulse at the receiver is a waveform with zero tail amplitude at 200 ps after its peak (corresponding to the ideal sampling point), where the sampling of the next pulse occurs.

2.4.2 Receive equalizer

To mitigate the low-pass effects of the channel, this work uses both a pre-emphasis filter at the transmitter and an equalizer at the receiver. As explained in the previous

section, a multiple-tap FIR filter can be implemented at the transmitter with low complexity. This filter is used to cancel the long tail of the pulse response that extends across multiple symbol periods. However, this pre-emphasis filter has two shortcomings. First, as mentioned in the former section, is the loss of signal power in the preshaped transmitted symbol. Second is the limited frequency range over which this symbol-spaced filter can be effective. As known from basic signal processing theory [19], a FIR filter with T_s spacing between its input samples can effectively compensate frequencies up to:

$$f_{max} = \frac{1}{2 \cdot T_s} \text{ symbol-spaced FIR,} \qquad [2.8]$$

Therefore, the transmitter symbol-spaced filter ($T_s = T_{sym}$) does not compensate for the loss of higher frequency components that exist in the signal spectrum due to fast transitions. On the other hand, a subsymbol-spaced filter is effective for a larger frequency range. For example, a half-symbol-spaced filter ($T_s = \frac{1}{2} T_{sym}$) can compensate frequencies up to twice that of a full-symbol-spaced filter, and such a filter can result in sharper transition slopes.

At this point we should note two important points about sharp signal transitions in a system. First, in order to minimize the BER of a symbol-detecting transmission system, the received eye opening should be maximized. The eye opening can be increased through means such as ISI cancellation, signal power increase, and noise and jitter reduction. One can also increase the eye width by increasing the slope of signal transitions. Figure 2.16 shows two 4-PAM eye diagrams with the same signal swing, and zero jitter and ISI, but different transition slopes. Clearly, the eye diagram with a sharper transition results in

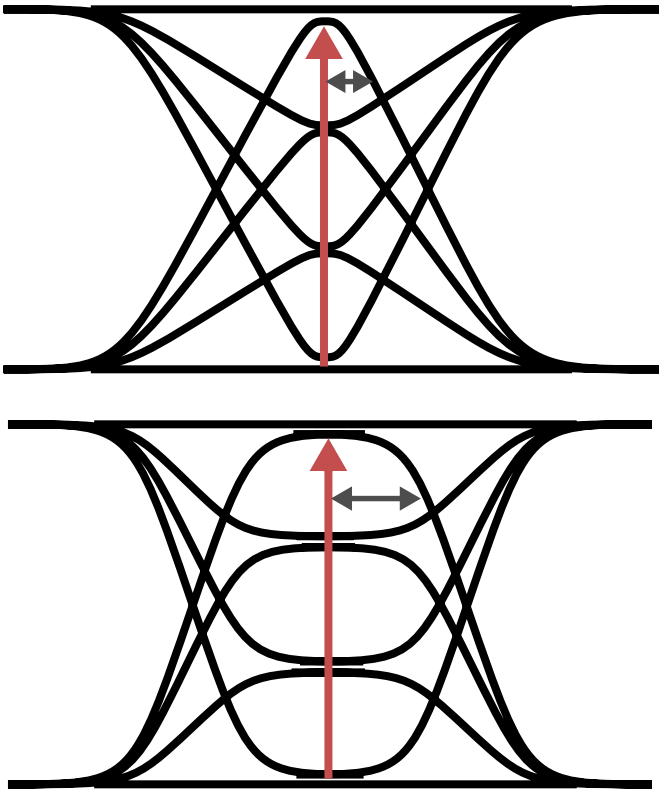

Figure 2.16: Two 4-PAM eye diagrams, a) slow transition, b) sharp transition

larger timing margins, which makes the system more tolerant of sampling phase errors. The effect of sharper transitions becomes more considerable especially in multi-level schemes, where the top and bottom eye width is seriously affected by the signal transition slope.

The second point concerning sharp transitions is that if they are generated at the transmitter, they can excite the high-frequency modes of ever-present line discontinuities and cause more ringing, and reflections, plus crosstalk between adjacent traces that can reduce the final eye opening considerably. Therefore, it is always best to adjust the transition slope of a signal to the minimum required when transmitting the symbols [12]. In addition, further signal preshaping is required to obtain sharper edges at the transmitter.

As discussed in the previous section, this preshaping function consumes part of the transmit power budget, which effectively reduces the final received signal power. Therefore, sharper transitions, i.e. more pre-emphasis, means less SNR at the receiver. These considerations show that implementing a subsymbol-spaced filter, which amplifies the high-frequency components, at the transmitter will not only degrade the signal quality, but also moves the transmitted signal power further into the more lossy regions of the channel that reduce the total received power.

Ideally, the best approach to cancel the lossy channel characteristics is to perform equalization at the receiver, since not only no effective signal energy is wasted for pulse preshaping, but also sub-symbol-spaced filters can be used to sharpen the transition edges of the signal “after the channel”. However, it is not obvious how to design and implement such a receiver equalizer at high frequencies, especially for a multi-tap FIR filter.

Fully digital receiver equalizers, using FIR filters, require high-resolution sampling ADCs that run at GHz speeds, which is quite a challenging task in present CMOS technologies. On the other hand, the disadvantage with fully analog *continuous-time* equalization is that it also needs very wide-bandwidth front-end receiver circuits that run at the same speed as the input data. The low f_t of transistors in present CMOS technologies makes the receiver equalizer design quite challenging at multi-Gbps rates. Input equalizers reported to date in CMOS technology all operate at data rates below 1.5Gbps [96],[21],[20].

This work uses a semi-digital receiver equalizer that is a 1-tap half-symbol-spaced analog FIR filter, represented by the following equation:

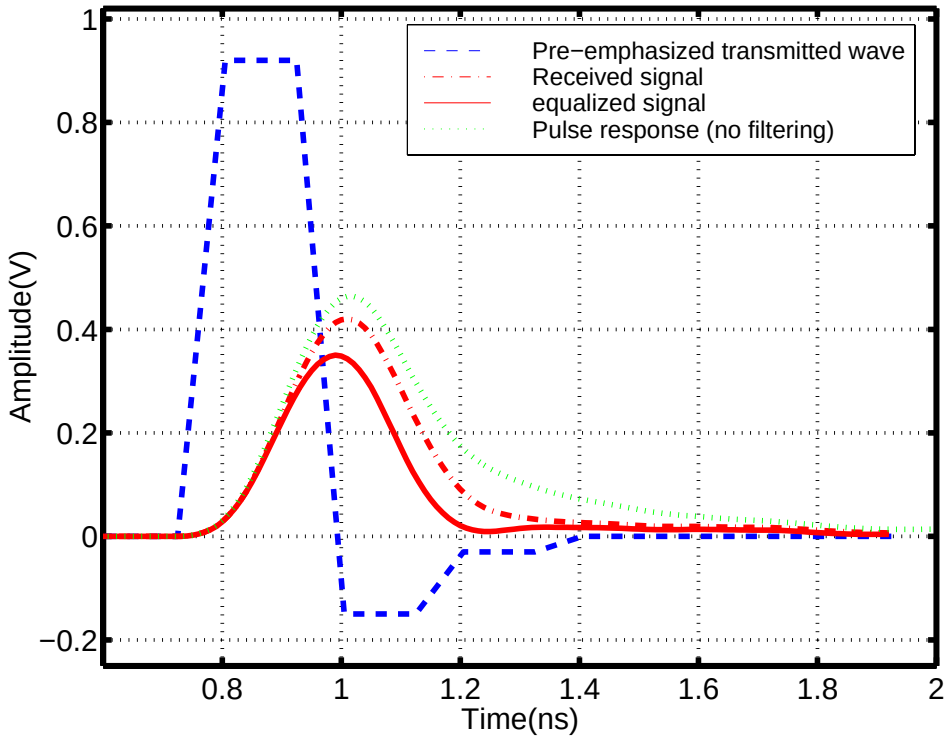

Figure 2.17: The effect of the pre-emphasis and equalization on the transmitted pulse.

$$Veq(t) = Vi(t) - \alpha \cdot Vi\left(t - \frac{T_s}{2}\right), \qquad [2.9]$$

where T_s is the symbol period or sampling interval. This equalizer, using half-symbol-spaced sample values, can equalize the signal over a frequency range that is double that of the transmitter filter without aliasing. Thus, the high-frequency components of the signal that are not compensated by the transmitter filter can be equalized in the receiver.

The effects of the receiver and transmitter filters for a 0.2-ns pulse (5Gsym/s signaling) at the near and far end of the 10-m channel are shown in Figure 2.17. The unfiltered pulse response remains at a large value 0.2ns after its peak (next symbol sample point), while the preshaped and equalized signal is zero at that point. All the filter tap

weights can be programmed for different channels. Comparing Figure 2.15 and Figure 2.17 shows that using a receive equalizer helps reduce the amount of pre-emphasis while obtaining the same ISI reduction in the system, effectively decreasing the transmitted power.

2.5 Selected Modulation

Once channel equalization is achieved, a higher data rate is possible by sending more complex symbols representing multiple bits during each bit-time. Although many modulation schemes can be used, multi-level pulse amplitude modulation (*M*-PAM) was chosen in our designs for its simplicity and higher bits/Hz compared to conventional 2-PAM (binary). As signaling frequency approaches the bandwidth of the cable and circuitry, multi-level PAM becomes attractive since it has smaller signal bandwidth and larger symbol period than binary signaling for the same data rate. By increasing the symbol period, the system has more tolerance to timing errors, and suffers less signal attenuation in the channel. In this section, the reasons for the choice of the modulation scheme used in this work are described in more detail.

2.5.1 Transmission symbol waveform

This work uses a trapezoidal pulse as the transmission waveform with a period equal to the inverse of the maximum symbol rate that can be generated and detected with modest complexity using symbol-by-symbol detection. At symbol rates much above the channel bandwidth, however, the trapezoidal pulses result in severe intersymbol interference (ISI) that limits the maximum communication speed.

This work uses different techniques to combat ISI, which is the main problem of using the trapezoidal pulses in this system. These techniques include channel equalization, as discussed in the last section, and multi-level pulse amplitude modulation, which is described in the following section.

For optimal detection of these trapezoidal pulses, *matched filters* should be used. Matched filters practically integrate the energy of each symbol over its duration, and therefore result in maximum signal to noise ratio [1]. In short-distance data transmission over cables, which is the goal of this dissertation, random noise power, which is mostly due to resistor thermal noise and device shot noise, is negligible compared to the actual signal power. Thus, a matched filter will not improve the system performance considerably, while simultaneously increasing the complexity, unless the received amplitude levels are quite small (<1mV). For example, thermal noise power density for a resistor is $4kTR$. In an environment with a 50-Ω termination resistor, the noise spectral density is $G_n \sim 1\times10^{-18}$ V^2/Hz. This noise density for a frequency range from DC to 5GHz (the frequency range of interest in this work) result in a RMS voltage of [Carlson]:

$$V_{rms} = \sqrt{\int_0^{5GHz} G_n df} = 72\mu V \qquad [2.10]$$

Assuming a Gaussian distribution for thermal noise voltage, one can approximate the peak noise voltage as $V_p \sim 3*V_{rms} = 210\mu V$, which is negligible even at very small signal amplitudes. As a result, to reduce the complexity of this design, hard detectors were used to sample the data pulses, without matched filtering of the waveforms.

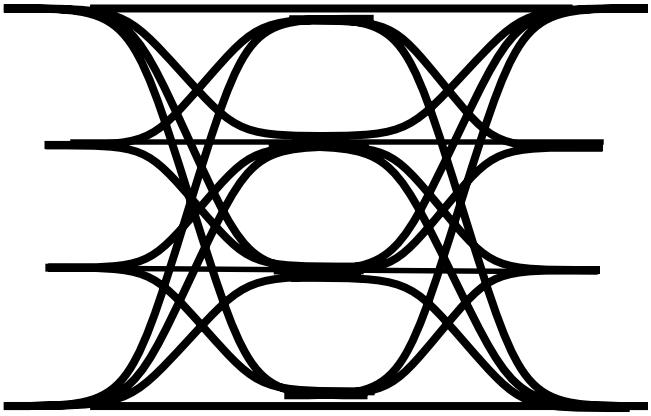

Figure 2.18: A sample 4-PAM eye diagram with finite transition times.

2.5.2 4-level Pulse Amplitude Modulation (4-PAM)

By transmitting multiple bits in each symbol time, the required bandwidth of the channel for a given bit rate decreases and system channel efficiency increases. The simplest multi-level transmission scheme is *M*-level pulse amplitude modulation (*M*-PAM), where each pulse conveys $log_2(M)$ bits of information. For a given data rate, *M*-PAM modulation reduces the effective symbol-rate by a factor of $log_2(M)$ compared to a conventional binary (2-PAM) system. This symbol rate reduction reduces not only the ISI in the channel, but also the maximum required on-chip clock frequency. Note that, however, the more complex transitions to multiple levels, together with finite signal transition times, prevent the data eye width from scaling as $log_2(M)$. Hence, the gain in the eye width is not the same as the symbol width increase for multi-level data. This condition is shown in Figure 2.18 for a 4-PAM eye diagram. This scheme also requires

more complex circuitry, including digital-to-analog (DAC) and an analog-to-digital (ADC) converters to generate and detect the multi-level signal.

This work uses differential 4-level PAM for transmission, which decreases the symbol rate by a factor of two compared to 2-PAM. An *M*-PAM scheme with larger *M* was avoided due to limited signal resolution of the receiver, especially at high speeds, and maximum transmitter output swing, both of which constrain the PAM level spacings. For a fixed transmitter swing budget, level spacings decrease with *M*:

$$LevelSpacing = \frac{XmitterSwing}{M-1} \quad . \qquad [2.11]$$

The smaller the level spacings, the more vulnerable the signal will be to noise from a variety of sources. The differential nature of the 4-PAM signal improves the link perfomance by eliminating the effect of common-mode noise and increasing the total signal swing by a factor of two compared to single-ended signaling.

The receiver amplifier and sampler combination in a high-speed system typically suffers from a large input voltage offset due to input device mismatch. The input devices cannot be made arbitrarily large to reduce mismatch, since their capacitive effects can limit the input bandwidth. The input referred offset voltage (typically $3\sigma \sim$ 30mV-100mV for reasonable device sizes in a modest CMOS technology) plays a major role in limiting the maximum signal resolution of the receiver. Conventional offset cancellation techniques using feedback cannot be used at GHz speed due to limited feedback bandwidth of circuits in the present CMOS processes. A more complex digital offset cancellation method has been implemented by Ellersick [33], which measures and trims the amplifier voltage offset by digital logic, obtaining final offset voltages of <10mV. This

scheme proposed in [33], achieves a maximum resolution of ~10 levels for the same maximum transmitter swing (~1V) that is used in this work. However, due to the extra complexity of this latter scheme, this work does not perform any receiver offset cancellation and therefore is limited in part by input offset.

Crosstalk from other high-speed signals in the system will also have a more detrimental effect on the low-amplitude signals. Moreover, since this design does not attempt to equalize for reflection ISI, reflection interference of large signals due to imperfect terminations and line discontinuities can overwhelm subsequently transmitted low-level signals. Therefore, in this work 4-PAM is chosen to avoid the low-amplitude symbols that can significantly increase the bit error rate of the system

The complex 4-PAM transitions result in an increase in eye-width of only about 40% (v.s 100%), after the 10-meter coaxial cable, compared to 2-PAM signaling. Although the nominal eye height is reduced by a factor of three (as there are three eyes), the eye height is effectively decreased by less than 50%, as there is less attenuation in the channel for the lower frequency 4-PAM signal than for 2-PAM. Therefore, a 4-PAM scheme typically has larger eye width, but smaller eye height, compared to 2-PAM at the same data rate.

On the other hand, multi-level phase modulation, [19], would result in larger eye heights, but about the same eye widths and timing error tolerance as 2-PAM. Previous experiments, [10], have shown that timing noise is more of a problem than amplitude noise in 2-PAM multi-Gbps links (signal level > 100mV), i.e. signal jitter limits the BER rather than the amplitude noise. This fact further suggests that 4-PAM will result in better

overall performance (BER), and thus is a better choice of modulation for the purpose of this project.

2.5.3 Coding

As discussed in Chapter 4, the timing recovery schemes in serial links need to extract the timing information from the received data signal. Therefore, the transmitted signal needs to have enough transitions to convey the necessary timing information to the receiver. This design uses a 4sym/5sym coding scheme that guarantees a transition in each 5-symbol group. This coding is performed by simply taking groups of four 2-bit data symbols, and appending a fifth symbol at the end that is the invert of the last (4th) data symbol of the group. The reason for such special coding will be described in Chapter 4 when the data recovery scheme for 4-PAM serial data is discussed. Because this coding does not provide DC balancing, the link must be DC coupled. More complex coding schemes can be devised to add DC balancing to such a code with 20% overhead. However, a complex coding is not implemented as it was not the main focus of this dissertation.

Since 4-PAM hard decision decoding is used in the receiver, a fixed one-to-one mapping of every two input bits to a constellation point must be chosen. Six distinct mapping exists for 4-PAM. However, only a Gray-code mapping (Figure 2.19) guarantees that every nearest symbol error results in only one bit error. Thus, the expected bit error rate is reduced to that of the linear mapping.

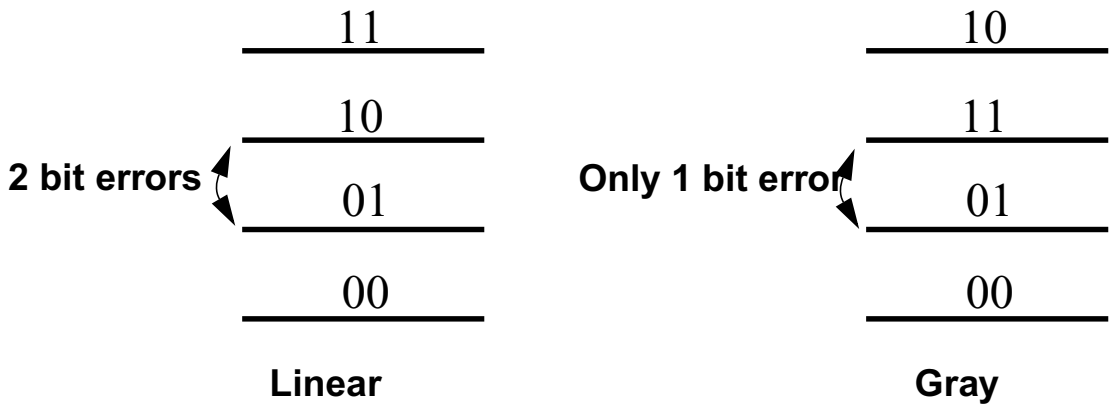

Figure 2.19: Linear versus Gray code mapping of levels

2.6 Summary

This chapter starts with the shortcomings of conventional CMOS links and the limitations they impose on maximum transmission rates (~100MHz) and minimum energy per transmitted bit (~50pJ). It is shown that this speed limitation is inherent in traditional CMOS link architectures, and that it doesn't improve with the process technology. It discusses how low-swing, incident-wave signaling overcomes the limitations of traditional CMOS signaling to achieve over an order of magnitude higher bandwidth (>1Gbps).

At rates exceeding about one gigabit per second, speeds are limited by other factors such as process technology and the bandwidth of cables. In a low-swing incident-wave link design, the bit time (symbol time) can effectively shrink to about one 1-2 FO4 delay (150-300ps in typical 0.35μm) and thus scales with technology. The bandwidth of the interconnection is now limited by the low-pass characteristics of the line caused by skin-effect resistance and dielectric loss, which increase with the length of the cable. This

frequency dependent attenuation results in intersymbol interference, which reduces the maximum data rate.

This chapter introduces a number of techniques to overcome these limitations to enable higher data rates. A 4-PAM modulation is used that enables the system to transmit two bits of information per symbol, effectively doubling the data rate at a fixed symbol rate and clock frequency, which helps reduce the frequency content of the signal and operate in the less lossy portions of the transmission channel. To decrease the channel ISI further, FIR filters are used at both the receiver and transmitter. The transmitter has a multi-tap FIR preshaping filter that cancels the long non-zero tail of the cable impulse response. The receiver equalizer is a 1-tap half symbol-spaced filter that sharpens the signal transitions, and hence improves the eye opening further.

Parallel I/O architecture is used to reduce the symbol time (bit time) from one clock period to a fraction of a clock period. In this work, the on-chip frequency requirement is further reduced to 1/5 the symbol rate (1/10 bit rate) by performing 5:1 multiplexing and 1:5 demultiplexing directly at the chip pads, allowing 5 symbols to be transmitted every clock cycle. The five 4-PAM symbols correspond to 10 bits of data that include 4 data symbols (8 bits = 1 byte) and 1 symbol for line coding (2 bits).

Finally, by using all the techniques discussed, bit times of about 0.5FO4 are shown to be achievable in CMOS technology. The following chapter details the circuit implementation of these techniques in CMOS.

Chapter 3

Full Transceiver Data Path

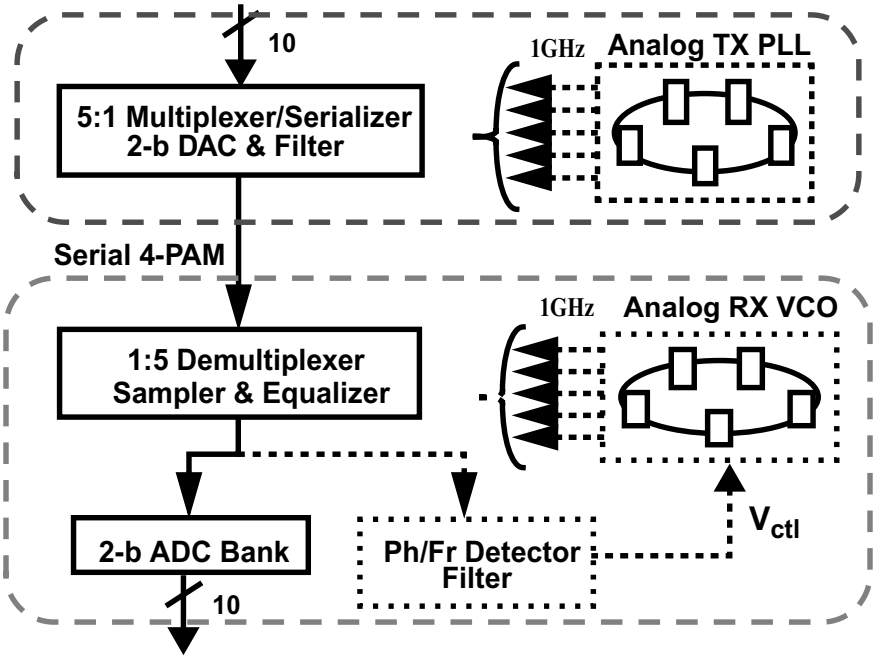

Figure 3.1: Transceiver top-level architecture

This chapter describes the transceiver data path for a multi-Gbps 4-PAM link. The next chapter describes the clocking for this data path. Figure 3.1 shows the block diagram of the link, where the data path elements are in solid boxes. It starts with the design

challenges for high-speed symbol generation, and proposes a 5:1 high-speed multiplexer architecture that places the 5Gsym/s stream onto the transmission line using a 1GHz clock. This multiplexer uses five high-speed 2-bit DACs in parallel to generate a 4-PAM output. The section continues with the design of a mixed-signal *N*-tap pre-emphasis filter and its integration into the high-speed driver to minimize complexity and power. However, the speed of a high fan-in multiplexer can be limited by the high capacitance of the multiplexing node, this problem is also discussed in the first part of Section 3.1.

The focus of Section 3.2 is the input receiver architecture, where parallelism is also used to perform 1:5 demultiplexing. Five evenly-spaced phases of a 1GHz clock are used together with a set of five samplers to sample a 5Gsym/s data stream. The 1-tap input equalizer is integrated with the input samplers using an analog approach to allow high speed operation of the filter. This section also addresses the design challenges of the input sampler and equalizer combination, and explores the factors that limit the bandwidth.

3.1 Transmitter Design

Figure 3.2 shows the top level architecture of the transmitter. To facilitate eye-diagram generation and BER measurements, a 2^7- 1 pseudorandom bit stream (PRBS) encoder is built on-chip. The PRBS encoder generates the pseudorandom data at a rate of up to 1Gbyte/s. To avoid running the PRBS logic at very high frequencies, the encoder is built two bytes (16 bits) wide and clocked at 500MHz [59], [60].

The 4/5-sym encoder performs line coding for the PRBS sequence by appending two bits at the end of each PRBS output byte. In serial transmission, line coding is

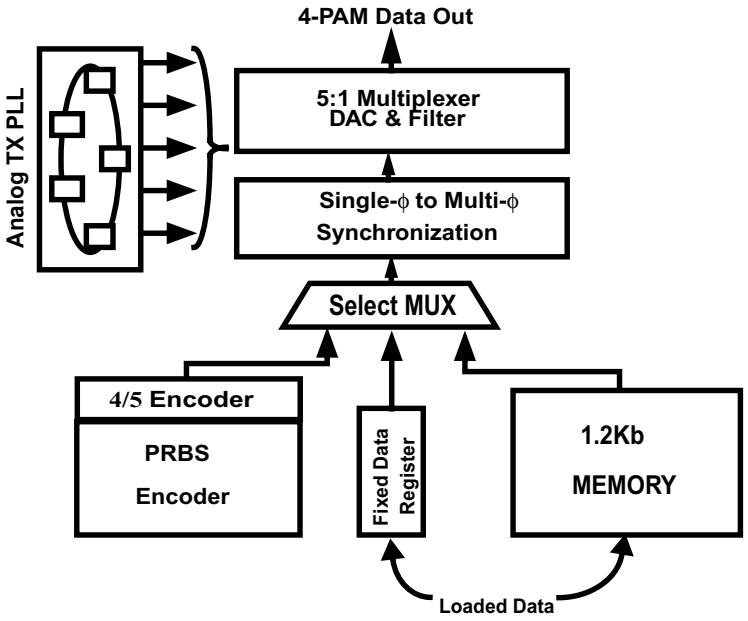

Figure 3.2: Transmitter top level architecture

required to guarantee enough transitions in the symbol stream for clock recovery. This line code is similar to 8b/10b coding except for DC balancing of the serial data, [93]. The final output of the PRBS and 4/5-sym encoder is two 10-bit parallel data *word*s clocked at 500MHz (equivalent to 10Gb/s). The two added bits are the negated versions of the last two bits in the PRBS output byte. As a result, when ten parallel bits are converted into five 4-PAM symbols, the fifth symbol makes a transition to the same magnitude but opposite polarity as the fourth symbol, which results in a zero crossing that occurs exactly at the midpoint between two symbols (Figure 3.3). The importance of this type of coding is discussed in detail in Chapter 4.

A 1.2-kb static memory (SRAM) and a 20-bit data register are implemented on-chip, which enables the user to load and transmit data patterns of different sizes for other verification purposes. To read from the memory reliably at 10Gbps in a 0.4-μm CMOS

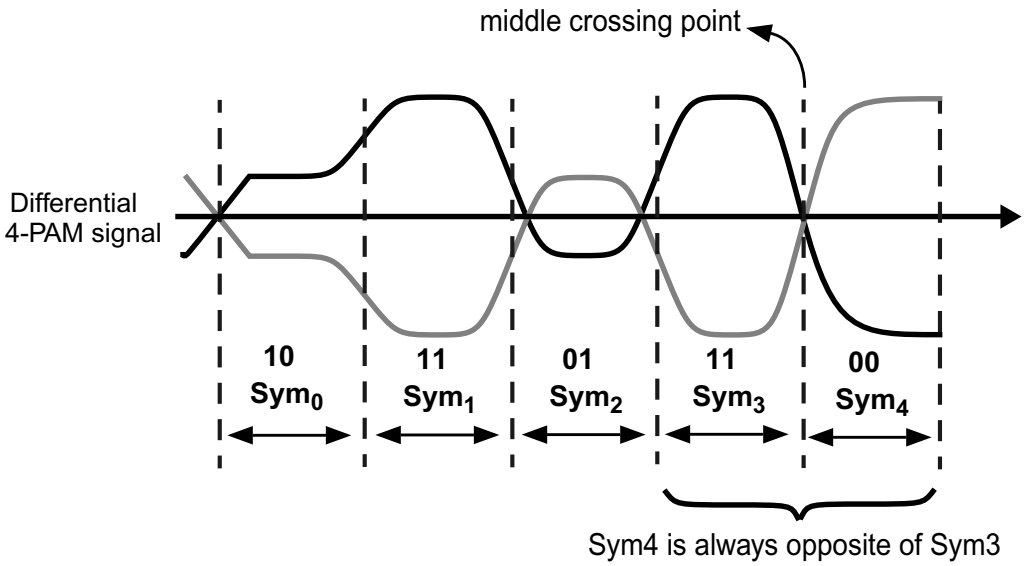

Figure 3.3: 4/5-sym encoder appends a fifth symbol to every four symbols to guarantee transitions in the serial stream.

process, the SRAM is 40 bits wide and clocked at 250MHz (half the frequency of the on-chip logic clock). The memory has a 2:1 multiplexer at its output that converts the 40-bit wide 250MHz output words from memory into 20-bit 500MHz words. Each one of the three data sources can be selected and sent to the 4-PAM driver separately by the 20-bit wide 3:1 multiplexer shown in Figure 3.2.

The selected 20-bit parallel data are further multiplexed, by a 2:1 multiplexer in the synchronizer, into 10-bit 1GHz words. The 10-bit words are then divided into five 2-bit data packs and delayed appropriately by the synchronizer block. The data timings in the synchronizer are carefully adjusted to ensure enough set-up and hold time for each 2-bit pack when transmitted by the 5:1 multiplexing driver at one of the five phases of clock, as shown in Figure 3.4. The synchronizer plays an essential role in enabling reliable data transmission at the speeds intended in this work (i.e. 5Gsym/s), since it minimizes clock-to-data skew sensitivity for every symbol transmitted by the fast output driver.

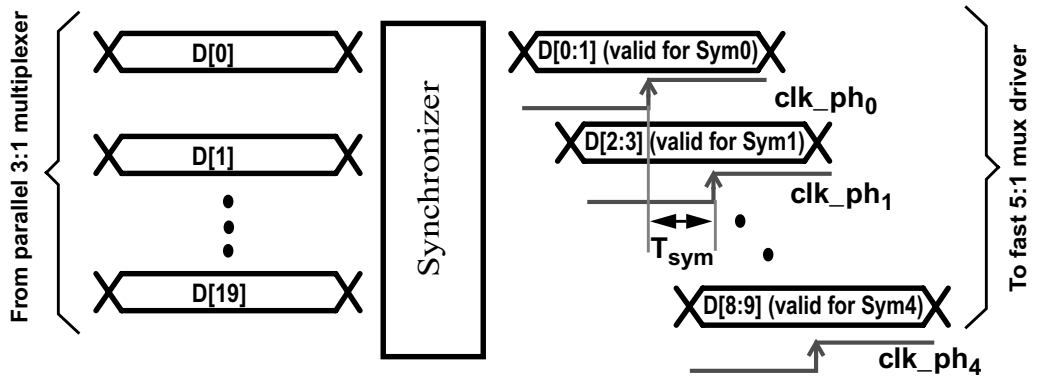

Figure 3.4: Synchronizer aligns each 2-bit pack to one of five clock phase of 5:1 multiplexer driver (clk_ph_i)

As the important components of the transmitter data path are the high-speed 4-PAM output multiplexer and pre-emphasis filter, each of these circuits is discussed in more detail in the following sections.

3.1.1 Proposed Architecture of the Multiplexing Driver

Figure 3.5 shows the general architecture of a 5:1 multiplexing transmitter as was originally proposed by Ken Yang [10]. Parallel low-speed data are applied to separate drivers that are activated consecutively by the narrow select pulses to generate the high-speed output stream.

One simple implementation for this multiplexer is shown in Figure 3.6. However, there are two main difficulties in implementing a multiplexing transmitter at high speeds using this method: the high fan-in switch multiplexer has a large output capacitance, and the select pulses need to have very sharp edges and short widths compared to the clock cycle (on the order of the output symbol period). The combination of these two factors

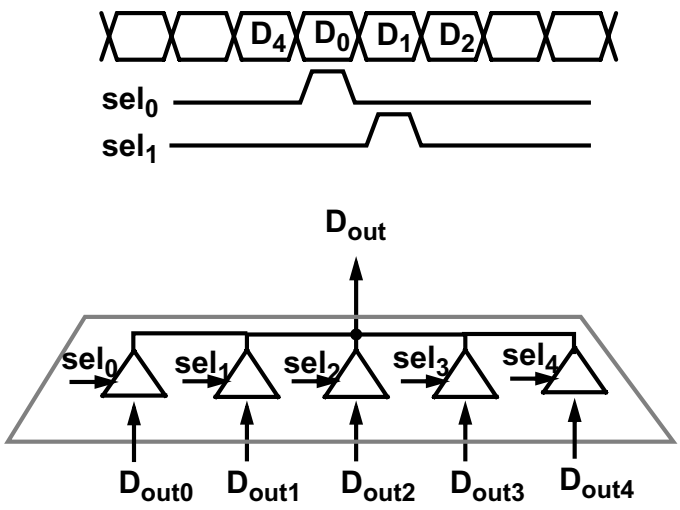

Figure 3.5: Transmitter 5:1 Multiplexer

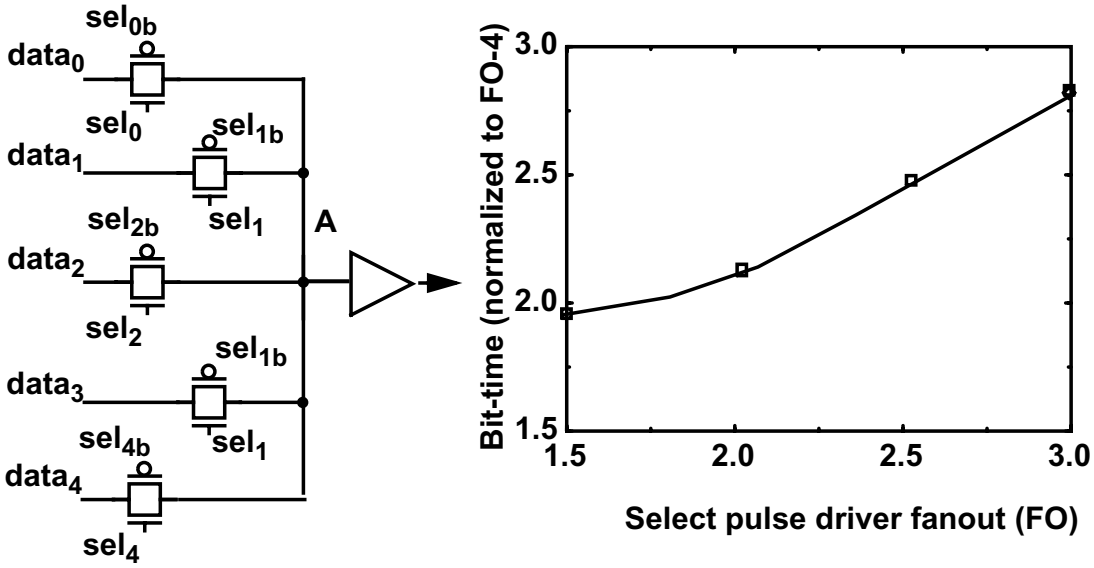

Figure 3.6: A simple 5:1 multiplexer

limits the maximum bandwidth of this multiplexer by generating intersymbol interference at the shared multiplexer node A. Figure 3.6 also shows the minimum bit time that this multiplexer can operate with less than 10% degradation in signal amplitude for different select pulse driver strengths. Clearly, a larger driver provides a sharper pulse that can

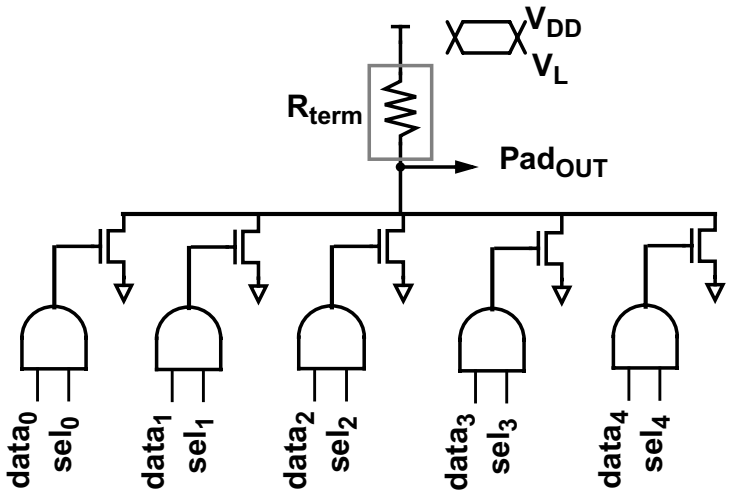

Figure 3.7: NMOS current-mode output multiplexer.

effectively reduce the overlap of the output symbols (ISI), and therefore allow higher transmission rates for the multiplexer. However, the minimum bit time does not decrease linearly with the select driver fan-out (FO) and is limited to a minimum of ~2FO4 due to the large capacitive loading at node A. Note that increasing the sizes of the switches will not significantly improve the minimum bit time as most of the capacitance arises from self-loading.

A NMOS current-mode low-swing multiplexer that is directly connected to the output line, as shown in Figure 3.7, can simply overcome the bandwidth limitation caused by high multiplexer capacitive loading. Since it uses only NMOS switches directly connected to the output pad, this multiplexer is the fastest (lowest output RC) that can be

built because of the low output capacitance of NMOS[1] and the low impedance[2] of the terminated line.

Note that the output NMOS transistors should always stay in saturation to act as high-impedance current sources, otherwise their output impedance could degrade the termination of the line. As a result, the output NMOS devices should be made large enough to supply the required current with the minimum gate overdrive that keeps the device in saturation. Although it seems that a multiplexing driver requires a very large output device with a correspondingly large output diffusion capacitance, one should note that the large driver devices can share their output diffusion with the electrostatic discharge (ESD) diode always needed in a practical design, resulting in little or no extra output loading.

The proposed driver in Figure 3.7 still suffers from a bandwidth limitation as the minimum bit time is now limited by the width of the select pulses. On-chip pulse widths are in turn limited by the finite bandwidth of the logic gates, and in practice, generating pulses less than 1.5 FO4 is very difficult.

The architecture shown in Figure 3.8 overcomes this limitation by merging the narrow pulse generation into the output driver. In this design, the 5:1 multiplexing driver consists of 5 modules, each made of stacks of NMOS transistors. Each of the series stacks receives two phases of a clock that are a symbol period (T_{sym}) apart, ck_0 and $\overline{ck}_1$, from a 5-

1. The capacitance depends on the device size which in turn depends on the output current. NMOS transistors always have a small output-capacitance-to-output-current ratio.
2. For a doubly terminated line, the effective impedance at transmitter output pad is 25Ω(assuming a 50-Ω transmission line).

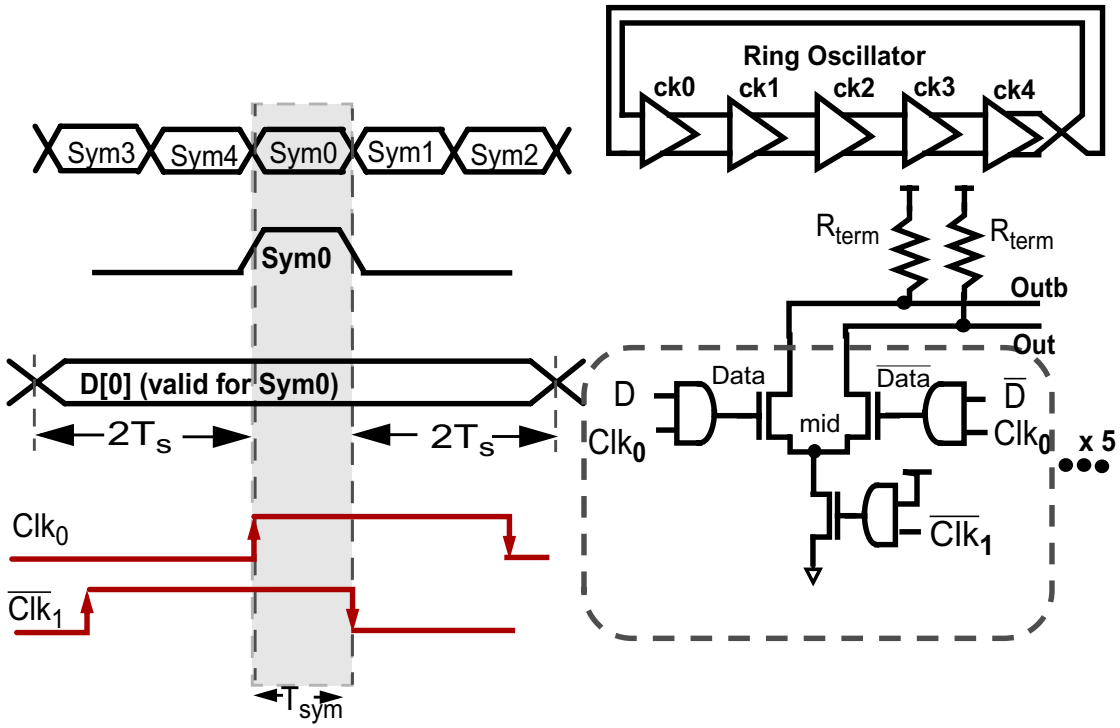

Figure 3.8: The proposed transmitter multiplexer design

stage ring oscillator. The top switch is driven by a *Data* signal that is qualified by a pre-driver AND gate so that input data *D* is evaluated only when clock, ck_0, is *HIGH* (Figure 3.5). During the other half cycle, both *Data* and $\overline{Data}$ are de-asserted low. In operation, the bottom NMOS switch discharges the internal node, *mid,* of the series stack with the rising edge of $\overline{ck}_1$. When ck_0 rises, one of the two driver inputs, *Data* or $\overline{Data,}$ rises and current for the bit flows through one of the branches until $\overline{ck}_1$ falls and turns off the bottom switch. This process repeats five times with each of the five driver legs to perform the multiplexing. The timing of the low-speed input data *D*, which is five T_{sym} wide, is set by the synchronizer block so that its transition edges are two T_{sym} away from the two clock edges that define the symbol boundary. As a result the output symbol has maximum timing margins on either side. Since the output symbol boundary is determined by the

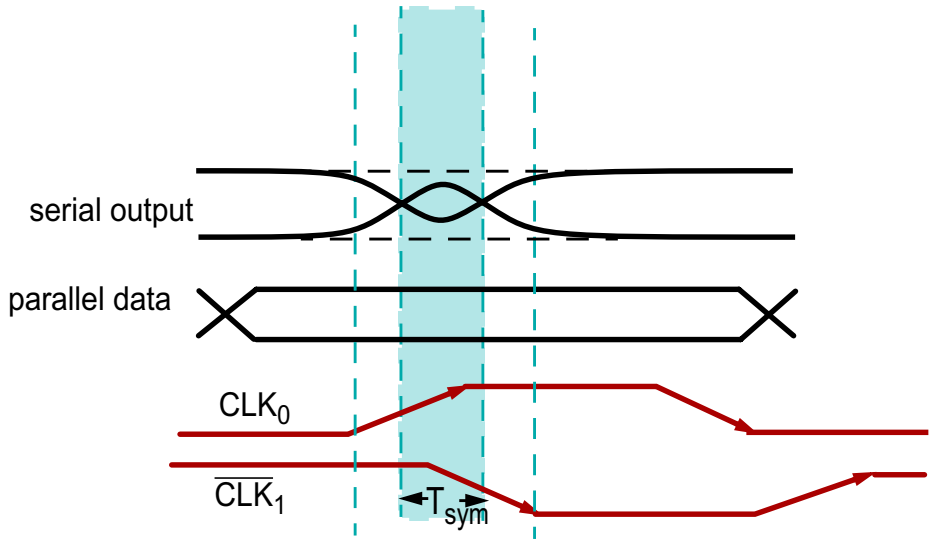

Figure 3.9: Data-eye opening reduction due to ISI cased by pre-driver slew rate.

phase spacing between the clock signals driving each multiplexer leg, care is required to minimize the timing error between the clock phases. These design considerations are discussed in more detail in the last part of this section.

Compared to the proposed design in Figure 3.7, the series devices would need to be larger than a single device to drive the same current. However, due to velocity saturation in short-channel devices, the increase in size is roughly 1.5x. This larger width may not be a penalty because output drivers, especially in submicron processes, require longer channel lengths for ESD protection anyway. The series stacked transistors potentially serve as the longer channel length device, depending on the design rules.

An advantage of the overlapping architecture is that the predriver's signals are at a much lower frequency than the data rate so only the final stage of the predriver chain requires a small fanout. Higher fanouts can be used in the predriver's buffering chain to save power. Although by using overlapping signals the minimum bit time is not limited by the minimum on-chip pulse width, intersymbol interference still occurs when the pre-

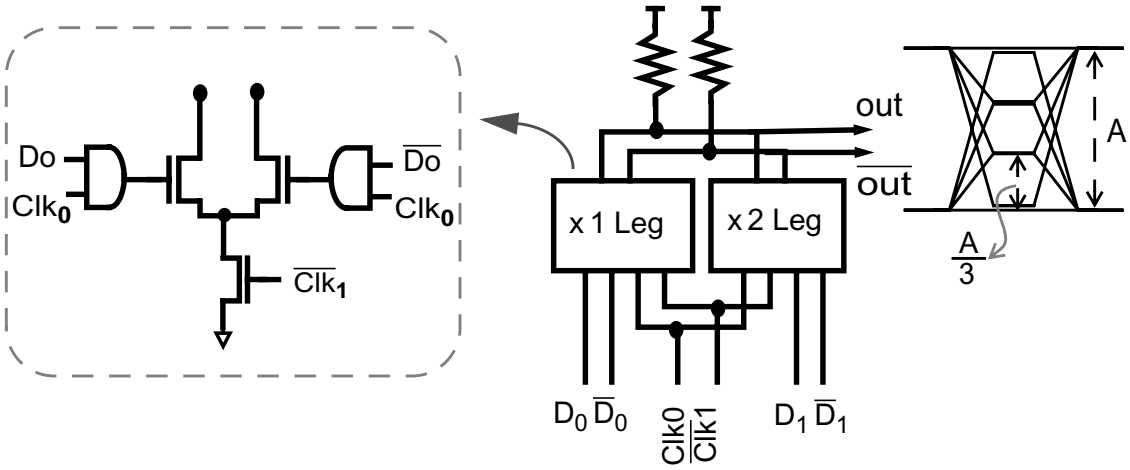

Figure 3.10: Driver 2-bit DAC module

driver transition time exceeds the bit time. Figure 3.9 shows the effect on the data eye from the predriver slew rate reduction. As the slew rate decreases, the current pulse extends into the neighboring bits and fails to reach its peak value before being switched off by the falling transition. As a result, a fanout of one is used for the pre-driver to ensure that the signal transition time is not more than a bit time. Using fanouts lower than one is not efficient because more than half of the total capacitance is from self-loading. Further decrease in the predriver fanout would improve the eye opening only slightly while consuming excessive area and power.

3.1.2 High-speed 2-bit DAC Driver

Using a current-mode driver, the best choice for implementing a high-speed low-resolution digital-to-analog converter (DAC) is a current-summing architecture. A 2-bit DAC driver leg using this method is shown in Figure 3.10. The DAC module contains two

differential driving legs that act as current sources as discussed earlier in this section. The two legs are binary weighted (sized) to generate four selectable equally-spaced current levels according to the 2 bits of input data, D_0 and D_1. In a 5:1 multiplexer, the 10-bit slow parallel data is grouped into five 2-bit sets. Each is used in one of the five parallel 2-bit DAC modules to generate the fast serial 4-PAM stream, using the same method shown in Figure 3.8.

Converting the output driver into a 2-bit DAC raises new design issues. An important one is the reduced noise margin of the transmitted signal compared to that of binary transmission, as discussed in previous chapter. Assuming that there is a maximum amplitude budget *A* for the transmitter, each eye opening of a 4-PAM signal will now be reduced to *A/3* at best. The eye openings are degraded further in practice due to factors including DAC nonlinearity.

The finite output impedance of the output stacked NMOS devices and their mismatch in the two binary weighted DAC legs are the main factors that affect the linearity of the current DAC driver. To minimize mismatch between the two driving legs, the x2 DAC leg is constructed of two parallel replicas of x1 legs, and random current offsets among the x1 legs, caused by transistor mismatch, are further reduced by careful layout.

As discussed earlier, a 4-PAM transmitter typically requires a large output swing to generate a sufficiently large detectable eye opening for the receiver. The large output swing results in a large voltage drop on the driver devices, and due to NMOS finite output impedance and drain-induced barrier lowering, the pull-down output current will not

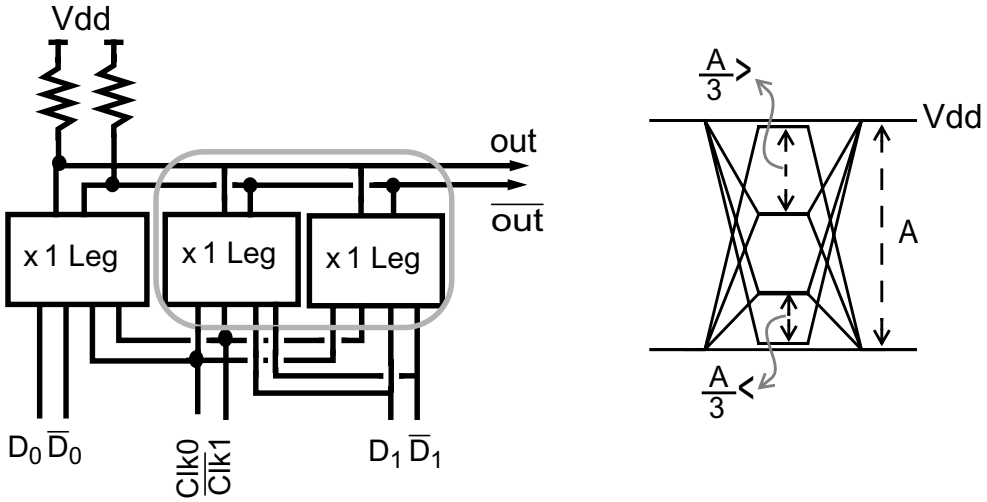

Figure 3.11: Effect of channel length modulation on 4-PAM eye diagram

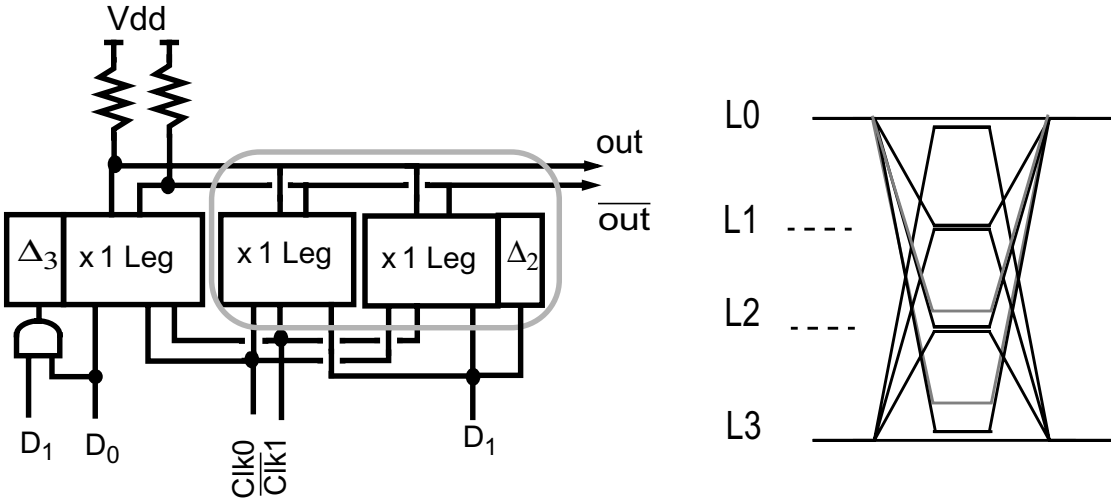

Figure 3.12: Trimming the 2-bit DAC legs to compensate for non-linearity

increase proportionally with the number of the x1 legs. Figure 3.11 shows the effect of channel length modulation on the 4-PAM eye diagram, when a 2-bit DAC with identical x1 legs is used. This nonlinearity can be corrected by trimming the strength of the driving legs for each of the four output levels separately. Figure 3.12 shows this method and the

trimmed 4-PAM eye diagram with the original levels shown in shaded lines. In this design the first x1 leg turns on without any trims to generate the first level L1. For the next level, L2, x2 legs turns on together with its trim driver Δ_2, adjusting the second level spacing equal to the first. Finally for the third level, L3, all the legs turn on including the trim leg Δ_3, which is only qualified when both LSB and MSB bits are one, pushing the last level lower to make all level spacings equal.

Note that the gate voltage of the driver NMOS devices (pre-driver output) should stay low enough to guarantee that these devices always stay out of the resistive region. Otherwise the DAC legs no longer act as current sources, and display poor linearity, and also reduce the effective termination impedance of the transmitter.

3.1.3 Three-tap Pre-emphasis FIR Filter

This section describes the design of the pre-emphasis filter in more detail. Figure 3.13a shows in greater detail the 5:1 multiplexing 2-bit driver as discussed earlier in this section. Each of the five driver blocks includes three filter modules (DAC-Ti) that use the same data as the main block (DAC-M), along with three consecutive clock phases. After the main module's current pulse, the three filter modules turn on consecutively in the next three symbol periods to perform pre-emphasis on the main pulse, as shown by the waveform in Figure 3.13a. To subtract the tap symbols from the main symbol, the differential currents generated by the tap drivers are summed at the output with the main driver current with the opposite polarity as shown in Figure 3.13b.

Tap symbols are generated with the same mechanism as the main symbol, but instead of a complex synchronizer used to provide valid data to each main driver, three

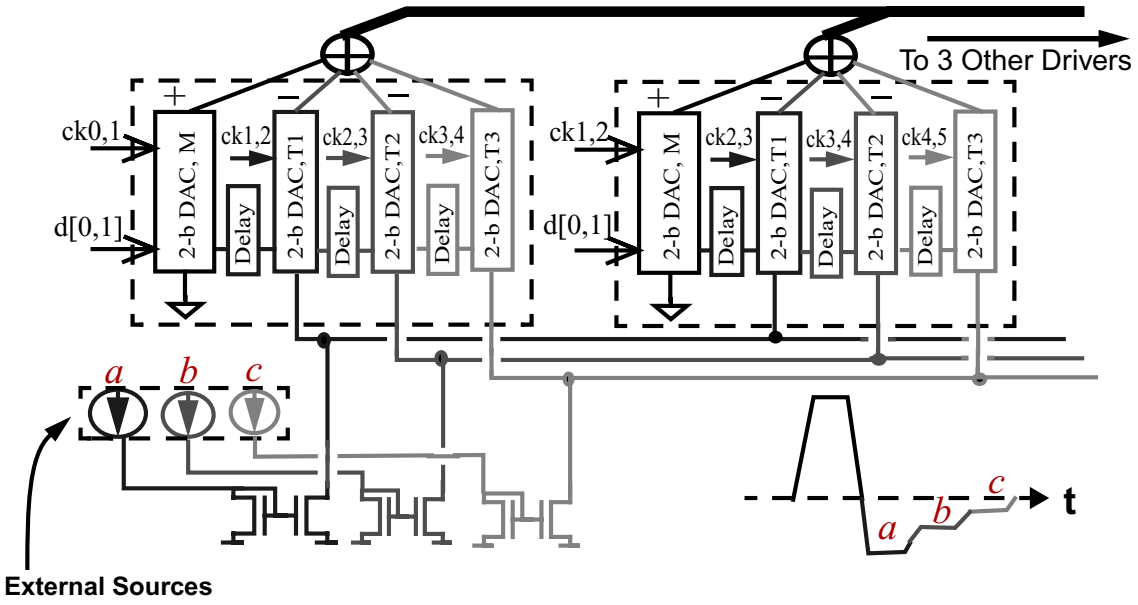

(a)

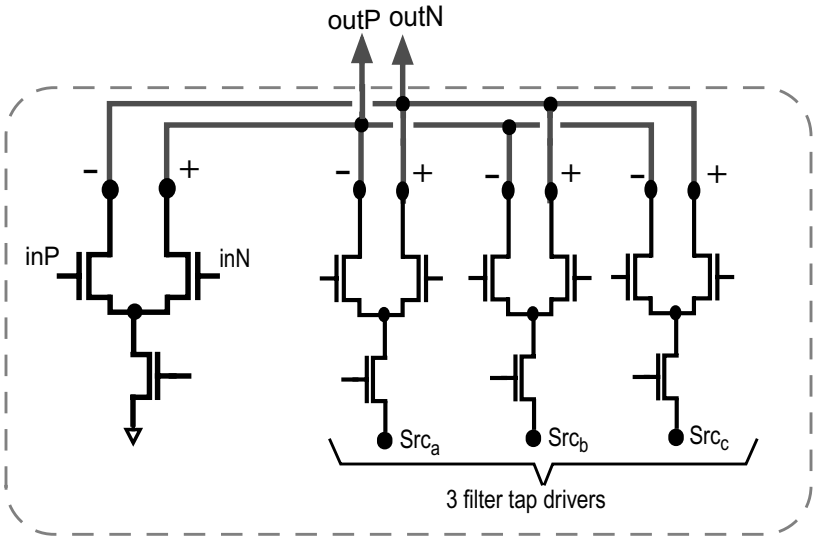

(b)

Figure 3.13: a) The multiplexing driver with 3-tap pre-emphasis filter, and b) circuit diagram of a driving block

series stages with a delay of approximately T_{sym} are used in each block to guarantee enough setup and hold time for the input data while it passes from one module to the other.

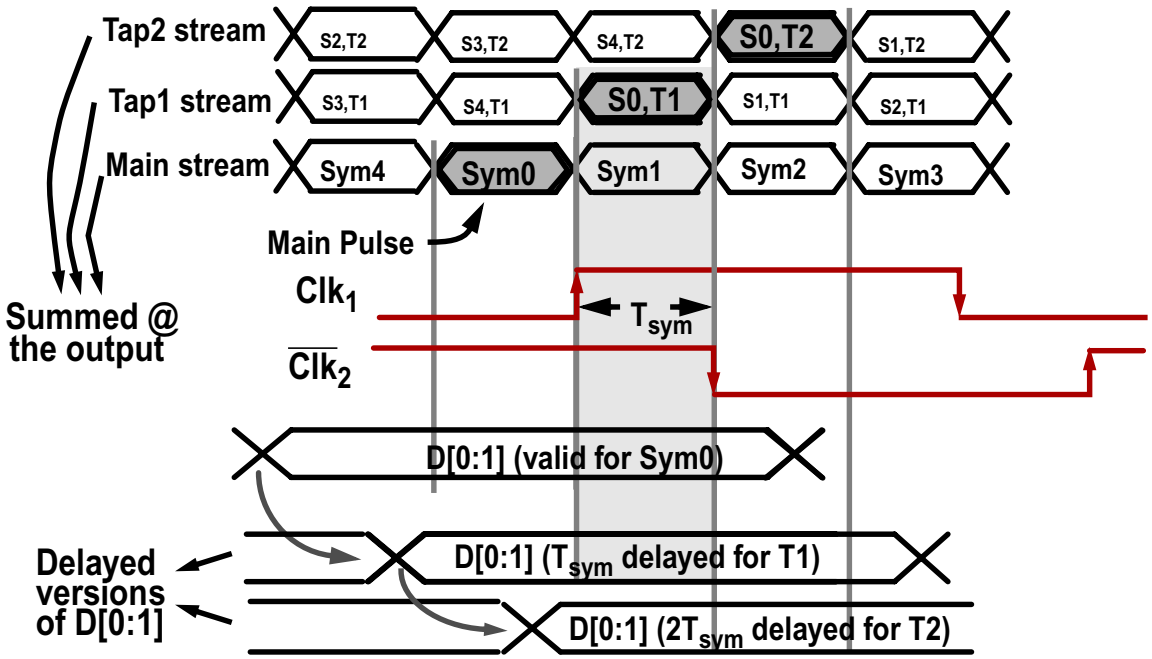

Figure 3.14: Tap symbol generation timing

Figure 3.14 shows this timing for the first two tap symbols. Note that the delays of the these stages do not need to be carefully adjusted to match a symbol period, since the input data to the main stage has two T_{sym} margins on either side, set by the synchronizer (Figure 3.8). Hence, the delay values should be accurate enough to guarantee that the data *D* after the last delay stage remains valid during the transmission of the last pre-emphasis symbol. To satisfy this condition, the maximum acceptable delay offset of the three series stages should be less than two T_{sym} at most. However the maximum acceptable delay offset of the stages reduces due to timing errors caused by delay mismatch of buffers in the data and clock path, and in the 5-stage ring oscillator. For 5Gsym/s signaling (T_{sym}=200ps), the maximum acceptable delay of the three stages is 300ps (2*T_{sym}-100ps)

or 100ps per stage, allowing a budget of 100ps for other timing errors. For typical process, temperature, and supply voltage, the delay of these stages is adjusted to 200ps (T_{sym}), and have a budget of +/-100ps ($T_{sym}/2$) for a delay offset around 200ps. Because the budget is large, these stages can be made of simple inverters without any process, temperature, and supply voltage compensation.

The currents in the filter taps (tap weights) are modulated by three controllable current sources connected to the *Src* node at the bottom of each module (Figure 3.13). To protect the tap currents from on-chip noise, each current source is a mirror whose input current is supplied from a clean off-chip source (Figure 3.13a). Because the corresponding filter modules in each of the 5 drivers is turned on sequentially, only one of the modules pulls current at each symbol time. Thus, each current source is shared among the corresponding filter module in every 5 drivers.

The *Src* nodes of the main modules' legs are grounded to minimize the device size for a given output current (Figure 3.13b). Smaller driver device sizes help reduce the pre-drivers and clock buffers in the transmitter, and as a result reduce the total transmitter power consumption. The device size reduction also prevent excess parasitic diffusion capacitances at the output from limiting the overall bandwidth.

3.1.4 Circuit Implementation

This section describes the implementation details and design techniques that improve symbol generation and signal integrity in the transmitter.

Two main sources of jitter and phase error in the transmitter are the 5-stage ring oscillator, which is discussed in more detail in Chapter 4, and the pre-driving buffers. To reduce jitter in the pre-drive buffering, the delay of these stages should be minimized as

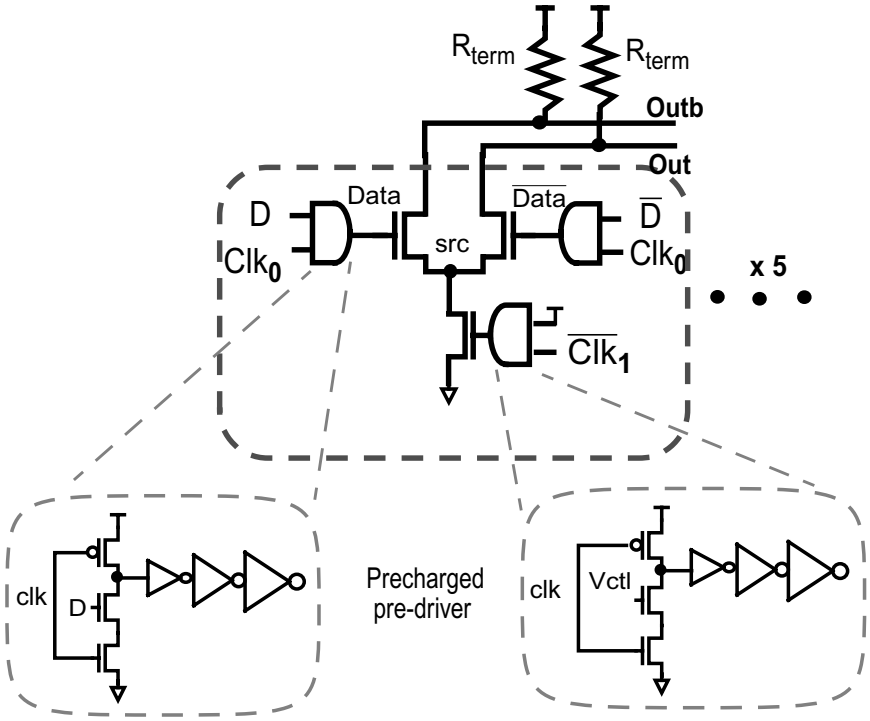

Figure 3.15: A driver with its pre-driver circuitry

the delay of the CMOS buffers is sensitive to supply noise. Figure 3.15 shows a driver leg and its pre-driver circuitry. To minimize the delay in the AND gates, they are implemented using pre-charged NAND gates, such that valid data passes through the gate on the output falling transition that is produced by the strong NMOS devices while the PMOS pull-up devices are already off. This pre-charged NAND stage also minimally loads the previous stage. As a simple inverter is the most efficient CMOS buffer in terms of speed and power dissipation, the pre-charged NAND is placed at the beginning of the buffer chain where CMOS inverters buffer the signal strength up to the large output driver. Thus, this design for the pre-driver AND buffers results in minimum delay and power dissipation. The implemented AND buffer has a maximum delay of 3 FO4. In the 0.4-μm CMOS used in this work (from LSI logic), this delay is equivalent to ~450ps (FO-4=150ps) in the worst case. With a supply sensitivity of 1%/% (percentage of delay variation per percentage of supply change), a 10% supply bump would only cause ~45ps (10% x 450ps) timing error, which is <25% symbol period reduction for 5Gsym/s signaling.

Even though only the data signal needs to be qualified by a precharged NAND, an

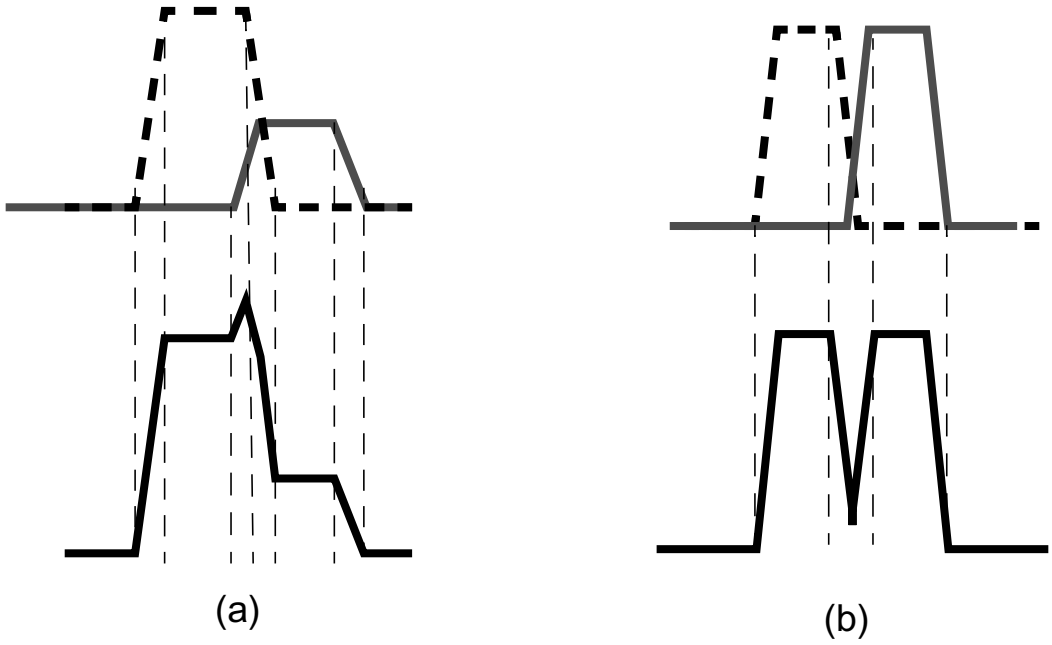

Figure 3.16: Two extreme cases of reduced effective symbol width

identical precharged stage also drives the clock input to the lower driver NMOS device in order to maintain the phase alignment between the data and the clock (Figure 3.15).

There is a separate power supply for the transmitter pre-driver from the rest of the chip. This power supply seperation allows the user to adjust the transmitter output current by varying gate overdrive of the driver NMOS transistors. However, the supply difference between the predriver stage and its previous stage results in duty cycle error in the signal waveform. In addition, Variations in PMOS-to-NMOS strength ratio in the buffering stages from oscillator to the final driver can result in duty cycle error in the clocks. This error reduces the effective width of the final output symbol, since the symbol boundaries are determined by both the falling and rising edges of the clocks (Figure 3.8).

Figure 3.16 shows two extreme cases of the problem caused by duty cycle error in

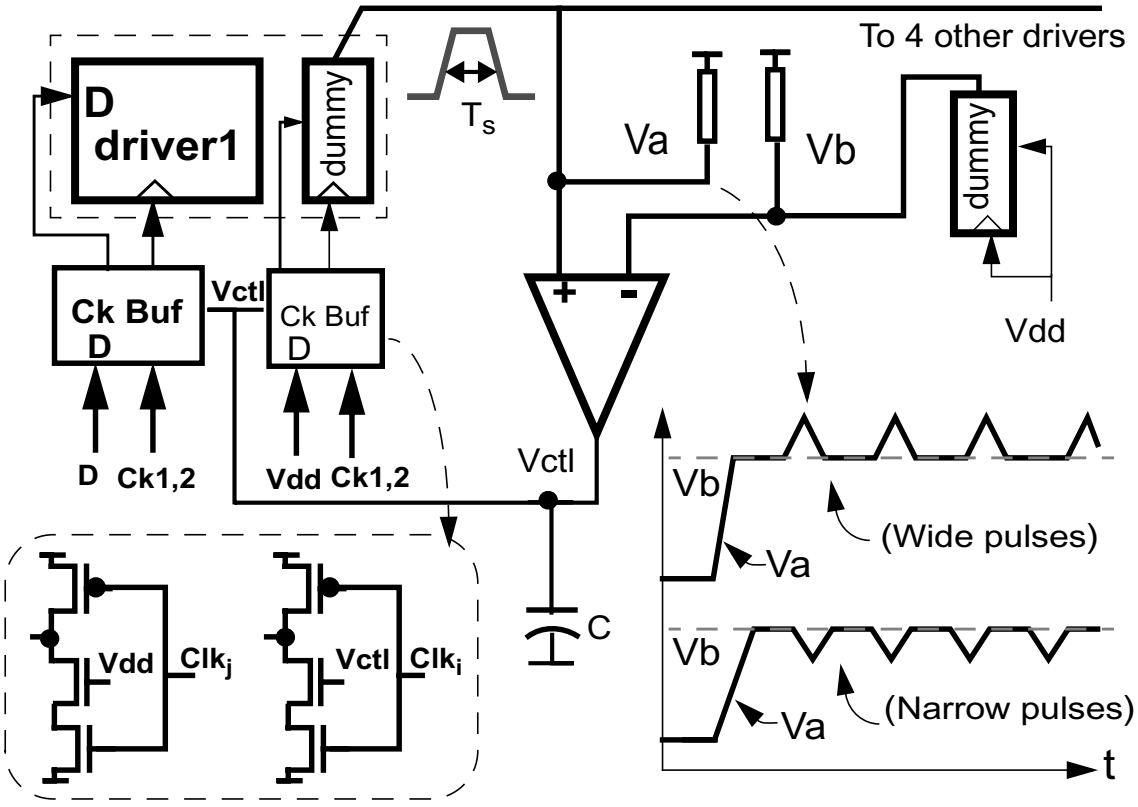

Figure 3.17: Symbol-width control loop

the predriver. In case (a) each pulse is longer than the optimum symbol period, thus corrupting the next symbol. In case (b) each pulse itself is shorter than a full period. However, the position of the clock falling edge (phase of $\overline{\text{Clk}}_1$) can be controlled by controlling the NMOS gate voltage *Vctl* in the pre-driver NAND gate (Figure 3.15). Therefore, the control loop shown in Figure 3.17 is employed to eliminate this problem by finding the optimum *Vctl* automatically. Each of the five driving blocks includes a dummy driver with the same topology and clock inputs as the main driver, but of a smaller size. These five dummy drivers are used to generate current pulses with fixed amplitude that are summed at node *Va* (in a manner similar to main symbol generation) to form one of the two waveforms shown in Figure 3.17. The top waveform occurs when pulses with longer period overlap the next

pulses, resulting in an average value greater than a reference *Vb* generated by a matching dummy driver, but drawing a fixed DC current. Similarly, the bottom waveform shows that shorter pulses result in an average value smaller than *Vb*. The error voltage between *Vb* and average value of *Va* is measured by a comparator whose output is a control voltage (*Vctl*) to servos the symbol width error toward zero for both the dummy and the main drivers. Mismatches among the dummy drivers and between the two pull-up resistors at *Va* and *Vb* nodes, in addition to random and systematic offsets in the comparator, result in phase error between the final two clocks arriving at the driver. Therefore, these blocks are carefully designed and laid out to minimize offsets. Capacitor *C* at node *Vctl* serves two purposes: decreasing the noise coupled to *Vctl*, and acting as the compensation capacitance for the control loop.

As shown earlier by Figure 3.9, the last stage of predriver should have a sharp rising edge to avoid reduction of output eye opening. Using the *Data* signals instead of the $\overline{\text{Clk}}_1$ to enable the current pulse has the advantage of speeding up the initial pull down (Figure 3.8). In this driver configuration, the charge sharing speeds the output transition and relaxes the size of the pre-driver PMOS because the data input pre-discharges the internal node *mid*.

Voltage noise on the gates of the stacked NMOS in the driver transforms directly into output amplitude noise. In order to reduce this amplitude noise, the last pre-driver should have a quiet supply voltage. A quiet supply on the complete pre-driver buffer chain also reduces the jitter in this path. Therefore, a dedicated supply with a large bypass capacitor to ground is used for the pre-driver stage. Fortunately, since the transmitter has differential current drivers, the net supply current stays almost constant.

The differential outputs are connected to 50-Ω on-chip resistors to eliminate line reflections. These resistors are built of non-salicided poly in parallel with a trimming PMOS operated in the triode region. The trimming PMOS resistor is used to compensate for the 10% tolerance of the poly resistor over process and temperature. As the PMOS resistor can only reduce the total resistance, the poly resistor is built nominally 10% larger (55Ω) than the target so that its resistance allegedly varies between 50Ω and 60Ω over corners. Hence, the maximum PMOS resistor should be 300Ω (since 300Ω||60Ω=50Ω) over all corners. To achieve good linearity and almost constant 50-Ω termination for varying output voltages, the output NMOS devices must remain in saturation to ensure a high output impedance compared to the 50-Ω line. Due to short channel effects, these devices can have a maximum output swing of 1.1V (2.2V differential) while maintaining a minimum output impedance of 500Ω (>>50Ω). This large output impedance results in a total nonlinearity of less than 2%, which is much smaller than the resolution of a 2-bit DAC. The good linearity of the stacked NMOS structure allows the 2-bit DAC drivers to be implemented without the trimming drivers as shown in Figure 3.12.

The transmitter also uses bondwires as series inductors L_{dd}, connected between Vdd and termination resistors, to increase the output bandwidth by "shunt peaking" (Figure 3.18). According to [23], the inductance value (L) should be $\frac{R^2 \times C}{2}$ to increase bandwidth by a factor of 1.8 with less than 3% frequency response peaking. With a total output capacitance (diffusion and interconnect) of approximately 1.4pF and a termination resistance of 50Ω, the optimum peaking inductance is 1.8nH, which corresponds to ~2mm of bondwire. This small length of bondwire is achieved by careful pad arrangements based on the pin configuration on the package.

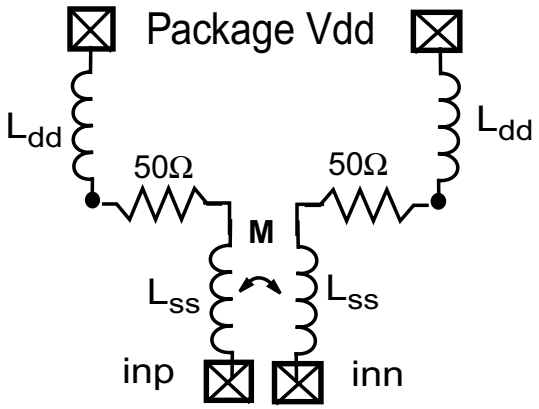

Figure 3.18: Shunt peaking to increase the I/O bandwidth

On the other hand, the inductance of the high-speed I/O bondwires L_{ss}, connected between the on-chip signal pads and package traces, can limit the link speed. Even with careful design to minimize the lengths of I/O bondwires, 2mm of bondwire (equivalent to ~1.8nH) still shows a considerably large impedance at multi-GHz frequency. For example, at 5Gsym/s signaling, where the frequency of the main harmonic is 2.5GHz (half the symbol rate), the impedance of a 1.8-nH bondwire is:

$$Z = 2\pi \cdot f \cdot L = 28\Omega \qquad [3.1]$$

An impedance of 28Ω between the transmitter output and the 50-Ω transmission line considerably degrades high-speed signal integrity. However, one can take advantage of the negative mutual inductance between the two differential output bondwires to reduce their effective series impedance by positioning them close to each other. In order to build a uniform transmission medium for the transmitter output, the spacing between the two differential bondwires can be adjusted in such a way that the effective inductance and capacitance (self and mutual capacitance) of the lines form a 100-Ω differential (50-Ω single-ended) transmission line (Figure 3.19). Applying this technique, 2mm long

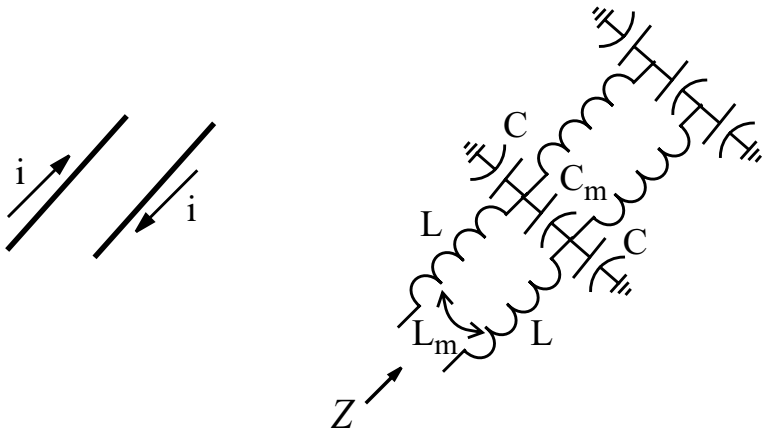

$$Z = 2 \cdot \sqrt{\frac{L - L_m}{C + C_m}}$$

Figure 3.19: Impedance of two adjacent wires carrying a differential signal

bondwires can be used without introducing a large discontinuity in the high-speed signal path.

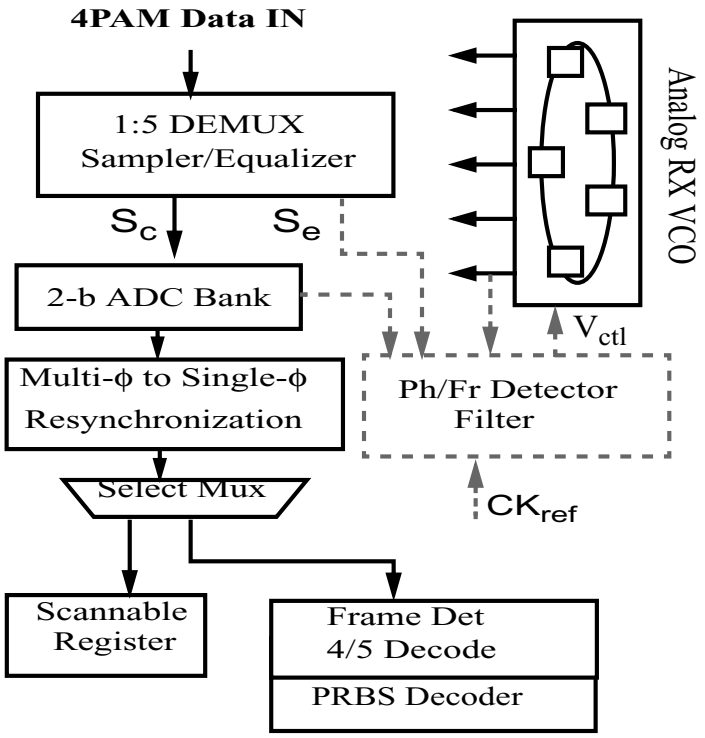

Figure 3.20: Receiver top-level architecture

3.2 Receiver Design

Figure 3.20 shows a top level view of the receiver architecture, with the main data path in bold. The receiver performs 1:5 demultiplexing at its input pads by sampling the signal with 5 out of 10 clock phases from a 5-stage differential ring oscillator (RX-VCO). As described in Section 2.3, this input demultiplexing helps relieve the speed limit set by the minimum cycle time of the receiver comparators as the sampling clock frequency reduces by a factor of 5. The 5 additional alternate clock phases allow 2x oversampling for timing recovery, and provide the required samples for the input equalizer with half-symbol-spaced taps. After equalization, the recovered data samples are converted to bits

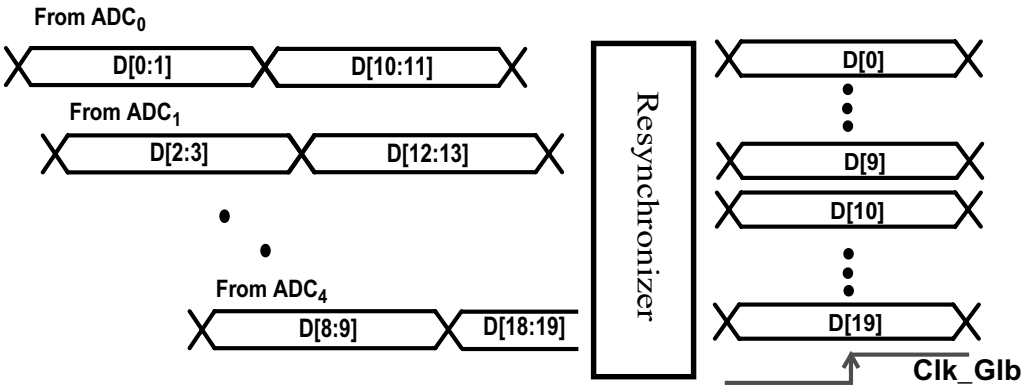

Figure 3.21: Resynchronization of the data packs from the five ADCs to a global clock.

(binary data) by a bank of five 2-bit ADCs. Finally, the bits in each pack of 10-bit data are pipelined properly, and resynchronized to a global clock. The function of the resynchronization block is exactly the opposite of that of the transmitter synchronizer, as shown in Figure 3.21. To halve the frequency of the on-chip global clock, the resynchronizer also performs 10:20 demultiplexing, generating 20-bit wide parallel data. Therefore, for the target 10Gb/s (5Gsym/s) signaling, the 5-stage RX-VCO generates 1GHz multi-phase sampling clocks, while the global clock runs at 500MHz to relax the timing constraints on the back-end logic.

A 2:1 multiplexer directs the parallelized data either to a 20-bit long scannable register or to the 2^7-1 PRBS decoder block. The scannable register, which is accessible through a low speed serial port, is used to read out a 20-bit long snapshot of the received data. The PRBS decoder block, which also includes a frame detector and 5/4sym decoder, checks the received PRBS sequence from the transmitter and flags an error whenever a wrong bit is detected in the sequence.

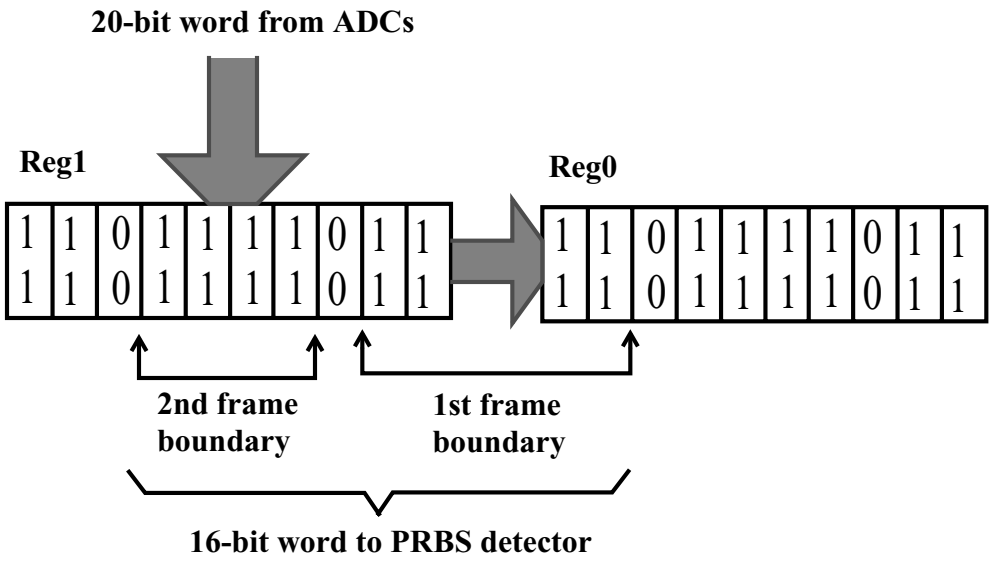

Figure 3.22: Frame detection scheme

As explained in Section 3.1, the generated PRBS sequence is 4/5sym encoded before transmission. The appended symbols are identified and removed from the received sequence such that the PRBS decoder recognizes it as the valid sequence. To facilitate frame detection, the transmitter starts with an initial training sequence before actual data transmission. The training sequence is generated by simply applying an all-ones data to the 4/5sym encoder, resulting in a sequence with four consecutive "11" symbols, and one "00" symbol appended by the encoder (i.e. 11,11,11,11,00). During the receiver initialization sequence, the present and previous 20-bit words are stored in two registers, ***Reg1*** and ***Reg0*** respectively (Figure 3.22). Then the frame detector looks for "00" commas to identify the boundary between received data bytes in these two registers. Once the byte boundaries are detected, the frame pointers are frozen and the recovered bytes are send to the PRBS decoder. This operation is shown in Figure 3.22.

The most important components of the receiver data path are the high-speed input sampler, equalizer, and 4-PAM comparator (2-bit ADCs), so each of these circuits is discussed in more detail in the following sections.

3.2.1 High-speed Input Sampler

Although comparator cycle-time is no longer an issue when a higher degree of input parallelism is used, the minimum time interval during which the receiver can sample the data remains an inherent limitation. This minimum time interval is also known as the *sampling aperture* (or *aperture window*). One can show that a sampled value is a weighted average of the sampler input signal over the aperture window, where the weight function is called the *sampling function*. Therefore, the output voltage of a sampler, V_{smp}, as a function of input signal, V_{in}, and sampling function, H_{sf} is

$$V_{smp} = \int_{-\infty}^{\infty} V_{in}(\tau) \cdot H_{sf}(\tau) \cdot d\tau. \qquad [3.2]$$

The width of the sampling function that includes the most significant area under curve, determines the sampling aperture. These definitions are shown in Figure 3.23. The aperture is effectively the time-domain dual of the sampling bandwidth. The smaller the aperture, the higher the input data bandwidth that can be sampled. However, note that the maximum sampling rate is also limited by the *aperture uncertainty* or *jitter*, which is the uncertainty of the sampling point (Figure 3.23). The above definitions are based on the terminology in [25]-[30].

The sampler used in this design is a PMOS pass transistor switched by a clock, as shown in Figure 3.24. The voltage on the sampling capacitor tracks the input data when

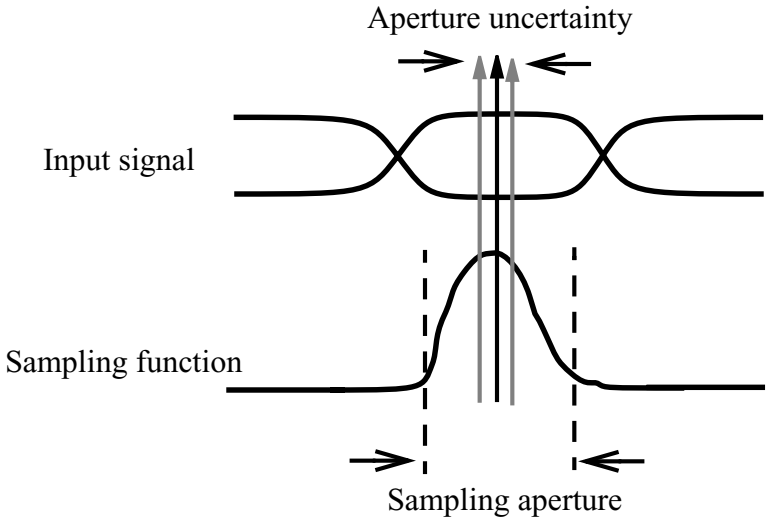

Figure 3.23: Definitions of some properties of a sampler.

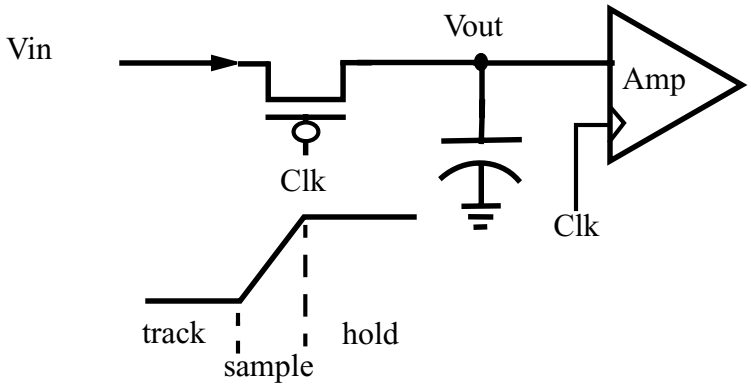

Figure 3.24: A PMOS switch sampler

the clock of the PMOS switch is *LOW*. When the clock is *HIGH*, the sampling capacitance holds the data voltage last seen. This sampled voltage is amplified by the comparator that follows while the voltage is held. Alternatively, many designs (such as [17] and [11]) merge the two functions into a single clocked comparator. This design was chosen because

it has one of the best sampling apertures [24] and has an analog output required for receiver analog equalizer (Section 3.2.2).

Note that an NMOS switch is generally a better sampler for high-frequency signals than is a PMOS switch, due to its smaller channel resistance. However, since no input ac-coupling capacitor is used in this design, the voltage level of the received signal is determined by the transmitter driver architecture. The transmitter driver, consisting of open-drain current sources with pull-up termination resistors (Section 3.1.4), generates an output waveform that swings between the supply and a few hundred millivolts below the supply. Thus, an NMOS sampling switch cannot be used as a receiver sampler here. On the other hand, a PMOS switch has the least resistance for input voltage levels close to the supply, and as a result can best track the high-frequency received signal for the conditions of this design.

Two characteristics of the sampling switch are of importance in determining the time resolution. First, the tracking speed of the switch must be high enough to let the output follow the input, and second, the transition from tracking to hold must be fast enough to minimize the sampling window width. The following discussions show the dependency of the aperture window on both of these characteristics.

Minimum aperture due to switch RC

The tracking speed of the PMOS switch depends on its RC time constant, which in turn limits the sampling aperture time. We initially assume a perfectly switched resistance

with infinitely sharp switching clock so that the aperture is determined only by the switch RC.

The aperture can be measured and defined based on the sampler's sampling function as discussed by Johansson [31]. The sampling function can be used to determine the sampled response to any input. Based on the definition of sampling function, the output of the switch depends on the input as shown in Equation (3.2). By treating the sampler as a linear filter[1], with an impulse response of *h(t)*, the output voltage at time t_s, where t_s is the sampling moment, is

$$V_{smp}(t_s) = \int_{-\infty}^{\infty} V_{in}(\tau) \cdot h(t_s - \tau) \cdot d\tau \quad . \qquad [3.3]$$

Comparing the Equations (3.2) and (3.3), one concludes that

$$H_{sf}(\tau) = h(t_s - \tau) \quad . \qquad [3.4]$$

Therefore, one can simply find the sampling function from the impulse response of the switch.

The impulse response for a sampling circuit can be obtained by using a step as the input at time t=0 and a clock sampling the input with varying times, t_s, where t_s is the time difference between step rising edge and clock. The resulting values that the switch samples at each time difference, t_s, forms the step response, which is then differentiated to

1. For a discussion of linearity and time-invariance assumptions for a sampling switch, refer to Appendix A in [32].

obtain the impulse response. Figure 3.25a shows the steps to obtain the sampling function for the switch with infinitely sharp clock.

Johansson [31] arbitrarily defines the aperture time as the width of the peak of the sampling function which contains 80% of the sensitivity (area). However, different values may be appropriate for different systems. For binary signaling, the minimum aperture width at the optimum sample point is that which would result in a sampled voltage larger than the logical switching threshold (1/2 the signal maximum amplitude) [32]. For 4-PAM signaling, which is the focus of this design, this threshold is as high as 5/6[1] the signal maximum amplitude, and as a result the aperture width should cover a minimum of ~84% (5/6) of the sampling function area. Of course this minimum aperture is only valid in the absence of amplitude and phase noise during sampling, and for a practical implementation this minimum may be over 90%. Hence, the following discussion explores different ways to improve the aperture time of a PMOS switch limited by the *RC* time constant.

When the PMOS transistor is in triode, the small signal channel resistance R is given by

$$R = \frac{1}{k \cdot \frac{W_{eff}}{L_{eff}} \cdot \left|(v_{GS} - V_{tb} - v_{DS})\right|}, \qquad [3.5]$$

where k is the process transconductance, and v_{GS} and v_{DS} are the large signal voltages. To obtain a small resistance both v_{GS} and W_{eff}/L_{eff} should be large, and v_{DS} should be low. Figure 3.26 shows the effect on the aperture time of increasing W_{eff} (decreasing R). The

1. In 4-PAM signaling, three comparators are used with three different voltage thresholds of 5/6, 1/2, and 1/6 of maximum amplitude.

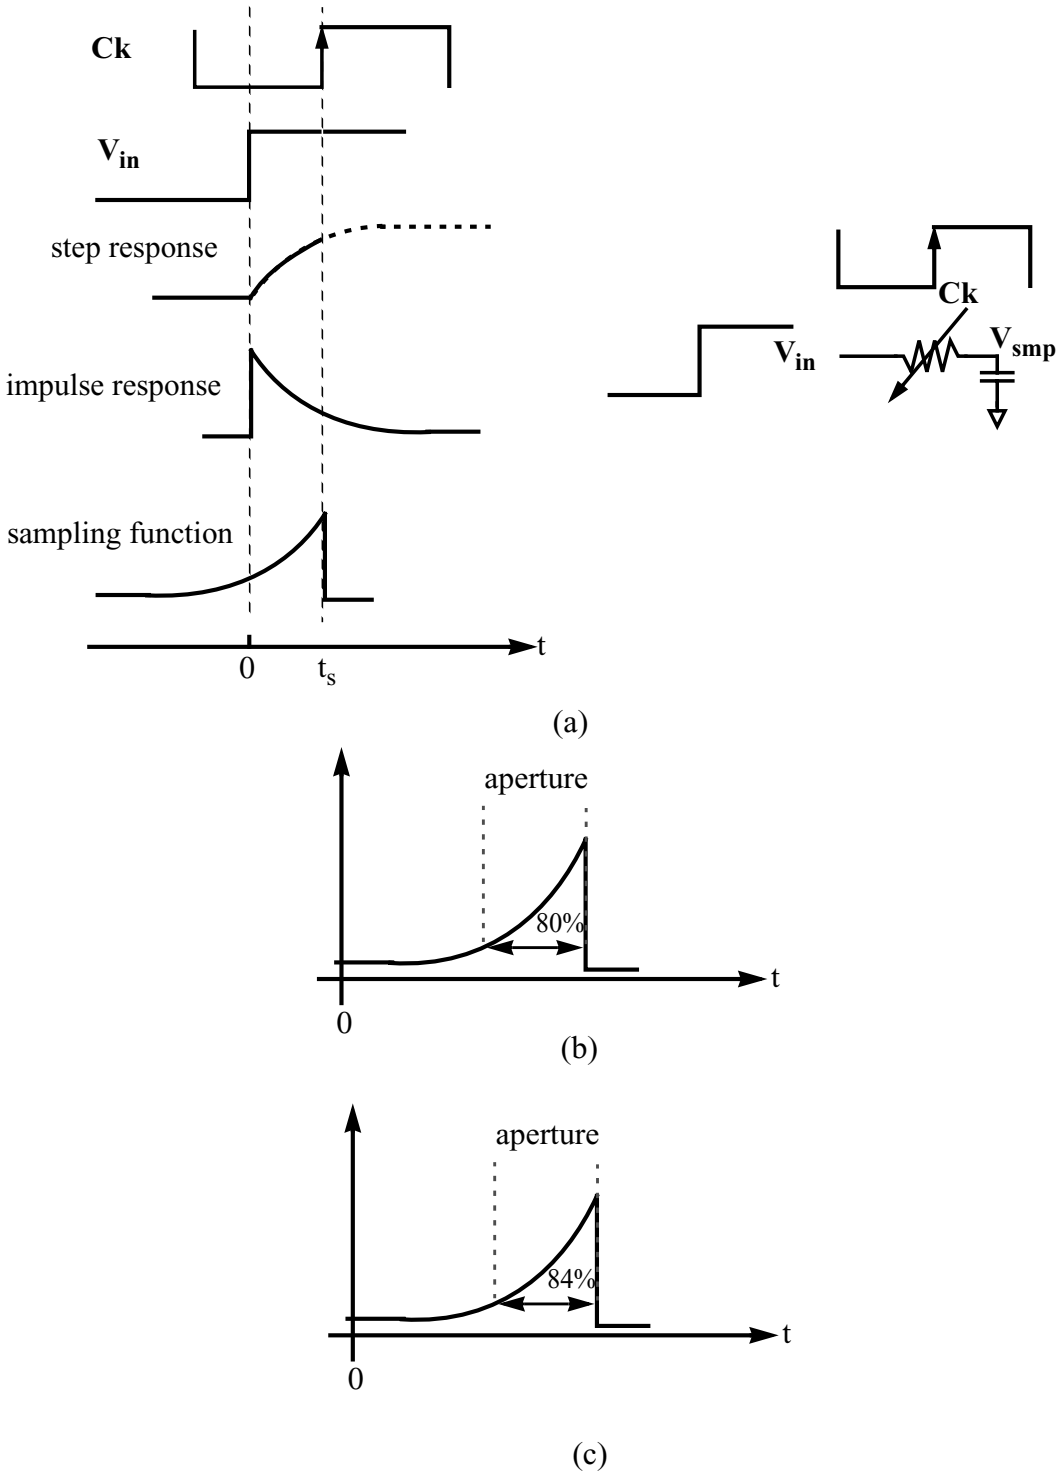

Figure 3.25: a) Steps to obtain the sampling function, b) aperture definition by Johansson, c) minimum aperture for 4-PAM signaling

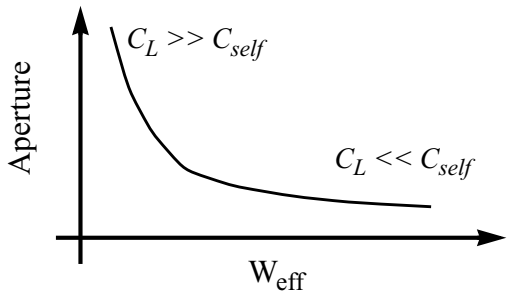

Figure 3.26: Aperture improvement with switch size

aperture window decreases with increasing W_{eff} until the switch self capacitance C_{self} dominates the load capacitance C_L. Beyond this point increasing the switch size does not change the effective RC ($C = C_L + C_{self}$). Hence, an optimum switch size that neither degrades the aperture nor increases the clock loading considerably corresponds to where the switch capacitance equals the load capacitance, $C_{self} \simeq C_L$.

The self capacitance consists of the drain diffusion C_{diff}, gate overlap C_{gdo}, and half of the channel capacitance C_{ch}.

$$C = C_L + C_{self} = C_L + C_{diff} + C_{gdo} + \frac{1}{2} \cdot C_{ch}. \qquad [3.6]$$

Although decreasing the switch size reduces the total capacitance, the resultant switch RC increases for a fixed load C_L. However, one can reduce the effect of self loading without changing the switch size, using the following techniques and considerations.

To minimize the drain diffusion capacitance, the PMOS switch should be folded so that the drain diffusion is shared between the two half transistors, as shown in Figure 3.27a. Although folding the transistor further reduces the source diffusion area, it also

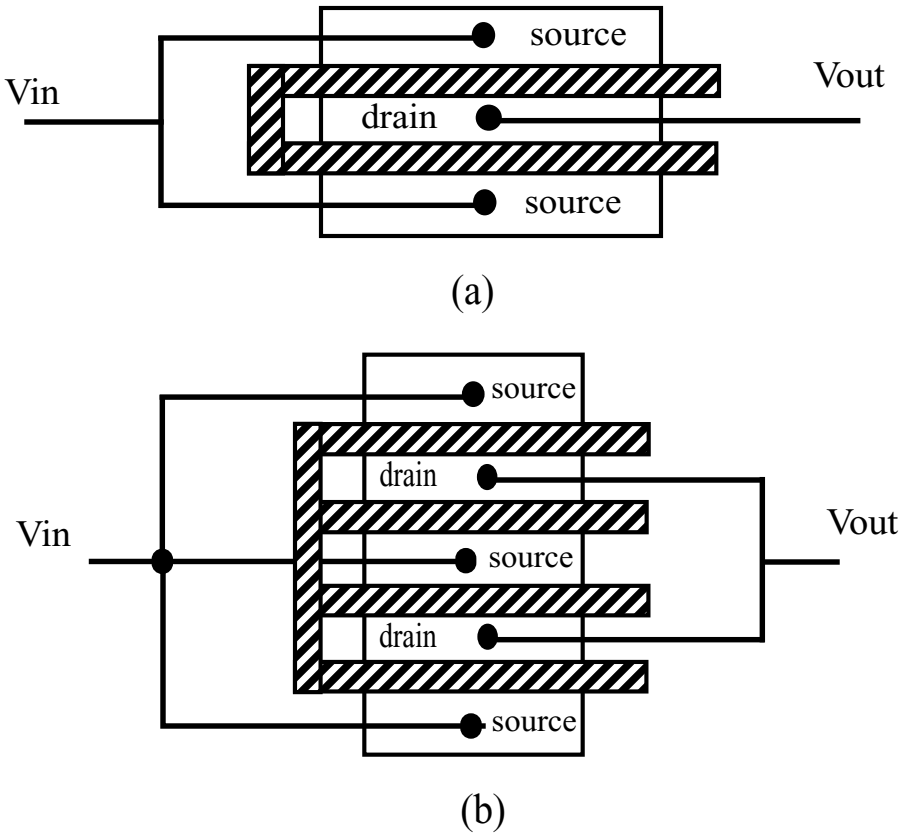

Figure 3.27: Different topologies for the sampler switch layout

increases the drain perimeter and thus the drain diffusion capacitance (Figure 3.27b). Therefore, there is a trade-off for each case and the optimum topology depends on the application. For the receiver samplers that connect to the chip pads and practically require an ESD protection circuit between the actual PMOS switch and the pad, Figure 3.27b may offer better results. This is due to the fact that ESD rules usually require a resistance on the order of ~200Ω in series with the switch, causing the input capacitance of the switch (source diffusion capacitance) to also affect the sampling bandwidth.

Standard ESD protection devices are one of the major bandwidth limitation in the high-speed path, and unfortunately do not improve with process technology because they

must always tolerate some certain amount of ESD energy. The foundry ESD circuit in the 0.4-μm process used limits the symbol rate to a maximum of 2Gsym/s before introducing significant ISI. As a result, no ESD circuit is used[1] at the high-speed front ends in this design, and samplers are directly connected to the receiver pads with a low 25Ω effective parallel impedance. Therefore, the switches use the topology in Figure 3.27a to minimize drain capacitance.

The other two capacitors, C_{gdo} and C_{ch}, are connected between the switch output and the gate of the PMOS device. One can reduce the gate capacitive loading by increasing the impedance between the PMOS gate and ground, such that the net loading is limited by the added gate impedance and not the gate capacitance. One way to implement such a circuit is to use a clock buffer (inverter) with a large pull-down resistor in series with inverter NMOS, so that when the switch gate is pulled low the impedance to ground of the pull-down path is large (Figure 3.28). The obvious problem with this approach is that the weak pulldown results in a slow falling transition that limits the maximum clocking frequency.

To improve the falling transition of the clock buffer in Figure 3.28, the buffer in Figure 3.29 is proposed. This circuit ideally results in infinite gate impedance during tracking mode while providing a strong pulldown. When Clk_1 rises, the stacked NMOS buffer turns on the PMOS switch by rapidly pulling its gate down to ground. A quarter of a cycle later, Clk_0 falls and cuts the pulldown path of the clock buffer, resulting in an open

1. Note that when no ESD device is used at the chip I/O, special care is required in the lab to guarantee that an ESD event does not happen during testing. In the case this design, the 50-Ω termination resistor and drain diffusion of the receiver sampling switches and transmitter drivers also absorb some of the ESD energy.

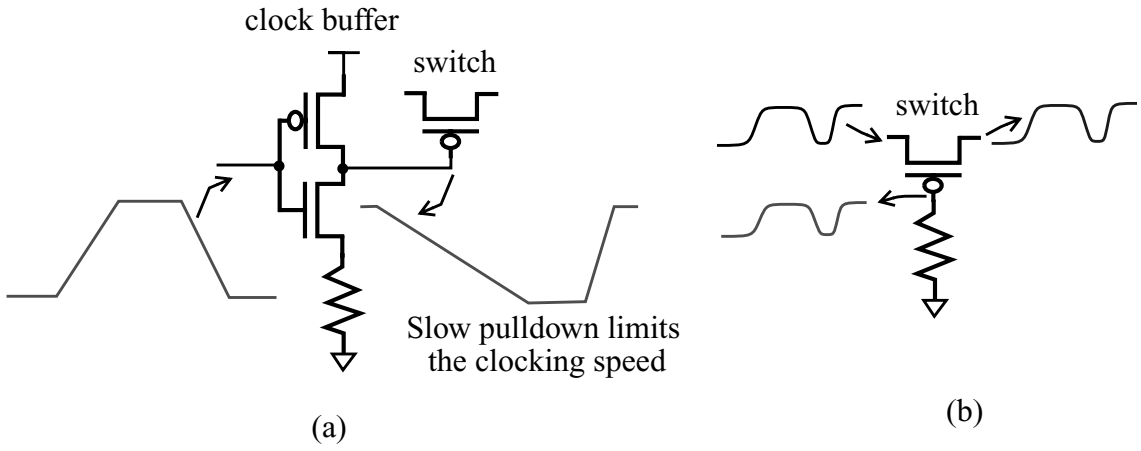

Figure 3.28: a) High-impedance pulldown clock buffer to reduce switch self-loading, b) PMOS gate tracks the input signal, thus reduces the gate capacitive loading

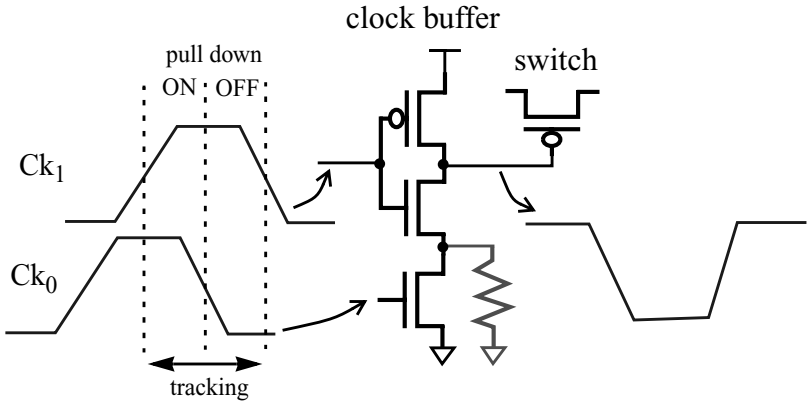

Figure 3.29: Modified high-impedance clock buffer to overcome clock speed limitation.

circuit (infinite resistance) path between ground and PMOS gate. This buffer also has a faster pull-up than a normal CMOS inverter, since during pull-up the NMOS stack is off

and the pull-up PMOS charges up the output (gate of the switch PMOS) without fighting the NMOS in the beginning as in a normal CMOS inverter.

However, note that this buffer design causes the PMOS gate to float for a quarter of a cycle. In very low speed operation (e.g. functionality testing), this floating node may get charged to some unknown value due to leakage currents and result in sampler malfunction. Therefore, it is recommended to add a large pulldown resistor to the clock buffer as shown in Figure 3.29 to guarantee proper operation at low speeds. One other important consideration for using this switch is the limitation on the input swing at high-speed operation. The gate of the PMOS switch is originally pulled down to ground by the stack of two NMOS transistors. After the lower NMOS is truned off, the PMOS gate tracks the input and its voltage drops below ground fr negative input transitions. Therefore, the input swing should be limited such that the PMOS gate voltage drop does not result in minority carrier injection through the upper NMOS drain diffusion (i.e. gate voltage drop should be below a diode drop).

Aperture degradation due to finite switching speed

In addition to what we have considered so far, a practical sampler has a finite switching speed due to limited clock slew rate at the gate of the PMOS. As a result, the aperture time is not only a function of the clock waveform, but also of the voltage dependent resistance of the PMOS switch, R. Figure 3.30 shows a plot of R versus the gate voltage (clock signal) with an input DC level of 2.5V for a PMOS 1μm wide and 0.3μm long. This plot shows that PMOS resistance is a weak function of clock waveform in the range 0V to 1.5V, and a very strong function from 1.5V to 2.0V. This property of the

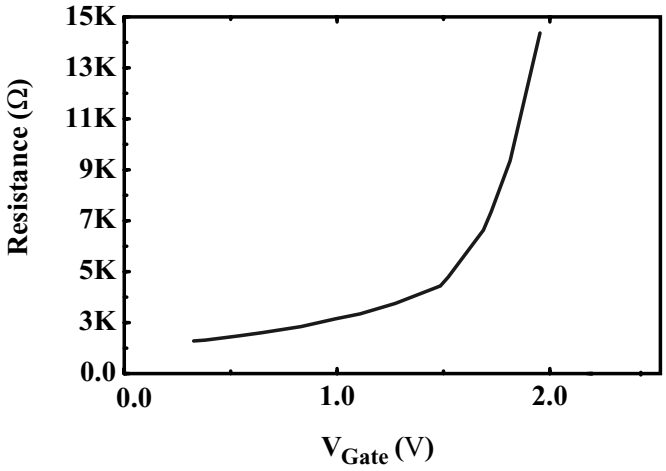

Figure 3.30: Plot of PMOS resistance vs. gate voltage for $V_{Drain}=V_{Source}=2.5V$, $V_T=0.5V$

PMOS transistor is highly desirable, because it means that the switching occurs only over the small range (0.5V) of clock waveform that defines the switching time.

To obtain the aperture time for the case of non-ideal switching, circuit simulations are required to account for the varying resistance of the PMOS during the switching event. Following the steps described in Figure 3.25a, the sampling functions for two different values of clock waveforms, $t_{rise}=0$ and t_{rise}[1]$=80ps$ (equivalent of a FO1 clock buffer), are obtained as shown in Figure 3.31. The finite clock slew rate of a FO1 sampling clock enlarges the sampling aperture by over 70% to 110ps (~3/4FO4), based on the definition of aperture for 4-PAM signaling.

In large-swing differential signaling, limited clock slew rate degrades the sampling performance further even with ideal switches ($R_{ON}=0$), since each one of the differential signals is sampled at a different time due to different gate-source voltages (V_{GS}) for each

1. t_{rise} is defined as the clock transition time from 10% to 90% of the full swing.

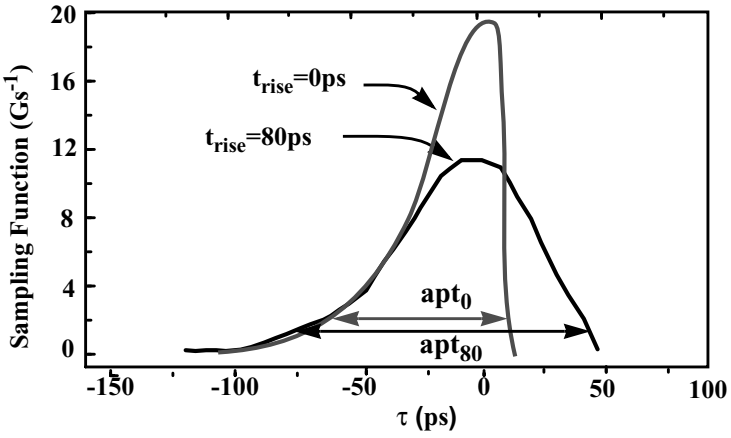

Figure 3.31: Effect of finite sampling clock slew rate on sampling function.

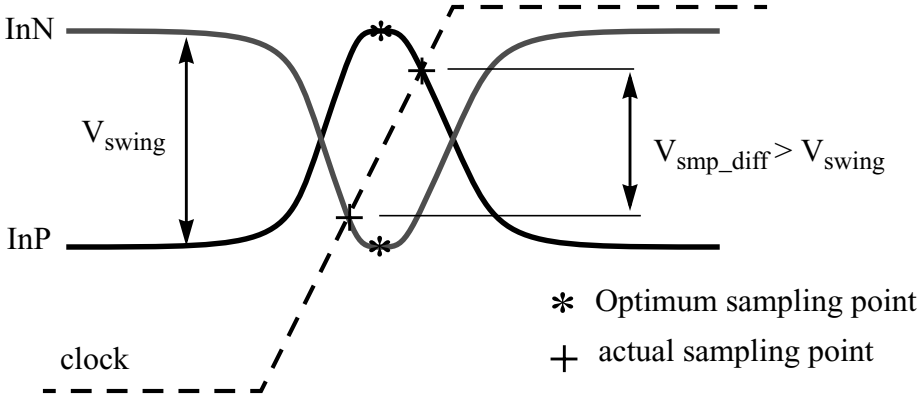

Figure 3.32: Effect of limited clock slew rate on large-swing differential signal for an ideal switch

PMOS switch at any time. This effect makes it impossible for the differential switch to sample each of the two signals at its own optimum point.

Figure 3.32 shows this effect when the sampling point is adjusted at the center of the input symbol. For simplicity the threshold voltage (V_t) of the PMOS and the ON resistance of the switch are assumed to be zero such that the sampling function is a delta at

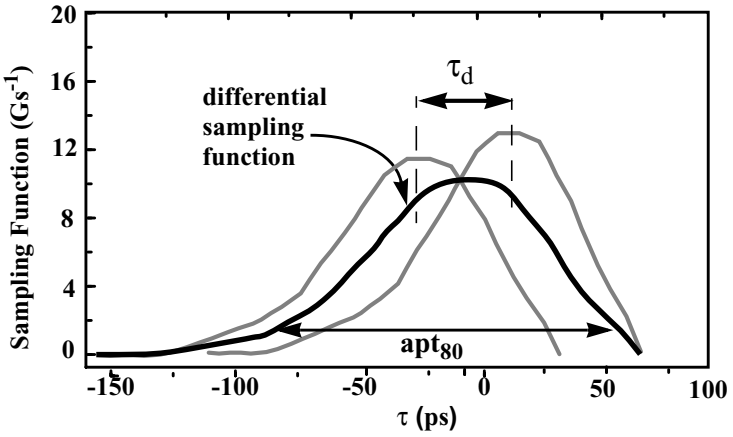

Figure 3.33: Spreading of differential sampling function due to clock limited slew rate

the crossing point of the clock and input signals. Figure 3.32 clearly shows that even with ideal switching the sampled differential value is smaller than the maximum differential voltage. Following the steps shown in Figure 3.25a, one can show for a non-ideal switch that the differential sampling function is the average of the single-ended sampling functions staggered in time by $\tau_d = V_{swing}/R_{slew}$[1]. Figure 3.33 shows the sampling function for a differential input swing of 1.0V (1.5V to 2.5V). A switch with a FO1 clock buffer (80ps rise time or 31V/ns slew rate), is assumed, hence τ_d = (1.0V)/(31V/ns) = 32ps. The differential sampling function is clearly wider than each single-ended sampling function, resulting in a larger aperture of ~130ps (0.8FO4) for the differential signal.

As shown in the forgoing discussions, sampling clock slew rate directly affects the sampling aperture, and therefore the clock buffer should have a small fan-out. As explained in Section 3.1, reducing the fan-out to much less than one does not improve the

1. R_{slew} is the sampling clock slew rate.

slew rate significantly as the buffer self loading dominates the total loading. Therefore in this design, sampling clock buffers with FO1 are used.

One source of concern in these sampling circuits is the amount of charge that is kicked back to the input from both the previously stored data and the switching of the clock. Since different samplers operate on different edges of the clock, the kickback of one sampler can interfered with the signal of a different sampler when it starts to track. An advantage of having a very low impedance at the receiver input is that any voltage error caused by samplers kickback is very small. In addition, a small sampling capacitance compared to the input capacitance results in a small amount of interfering charge. The switching of the clocks is a less significant concern because it introduces a common mode noise to the signal.

3.2.2 1-Tap Receiver Equalizer

The desire for a 1-tap half symbol-spaced receiver equalizer was discussed in Chapter 2. This section discusses the trade-offs and implementation details of this equalizer.

Since equalization has to be performed at a very high-frequency (symbol rate) on each data sample, speed limitations of the process make it impractical to implement this equalizer as a fully digital FIR filter. For a digital filter to generate accurate results, it requires the analog samples to be digitized with sufficient resolution. Therefore, this approach requires the receiver to possess high-resolution front-end ADCs with 1GHz[1] conversion rate. The only ADC architecture in a 0.3-µm process technology that can potentially operate at these speeds is the flash architecture [28]. However, the number of

1. For 5Gsym/s, the sampling clock frequency of receiver 1:5 demultiplexer is 1GHz.

comparators in a flash topology grows exponentially with the resolution[1] which makes implementation of ADCs above 4 bits of resolution impractical at multi-Gbps rates [33]. In a binary signaling system, a 4-bit ADC may be able to provide the necessary resolution for the digital filter to generate accurate results. However, the same accuracy in a 4-PAM system requires 6 bits ADC resolution, because the number of data levels is four times that of the binary signaling. In addition, the downstream digital logic consumes excessively large power and area [14] due to the high data throughput (10Gbps). Each one of these considerations can introduce major bottlenecks in the digital implementation of the receiver equalizer.

To avoid these problems, this design performs equalization directly on the sampled values completely in the analog domain before they are used by other blocks. The resolution of the input data to the filter is thus no longer an issue, and the need for high-resolution ADCs at the front-end is consequently removed. As will be seen, analog equalization is performed in one simple stage using previous and present analog samples.

Figure 3.34 shows the architecture of the 1-tap half symbol-space FIR equalizer, which operates on 2x oversampled differential samples provided by the receiver. As explained in Chapter 4, 2x oversampling is required anyway by the data recovery loop to extract clock information from the incoming serial stream, so this filter exploits an already present resource.

Having the present and former differential samples (e.g. S_0 and S_1 in Figure 3.34), the equalizer subtracts the weighted value of the former sample (S_0) from the present sample

1. A *m*-bit flash ADC requires 2^m comparators

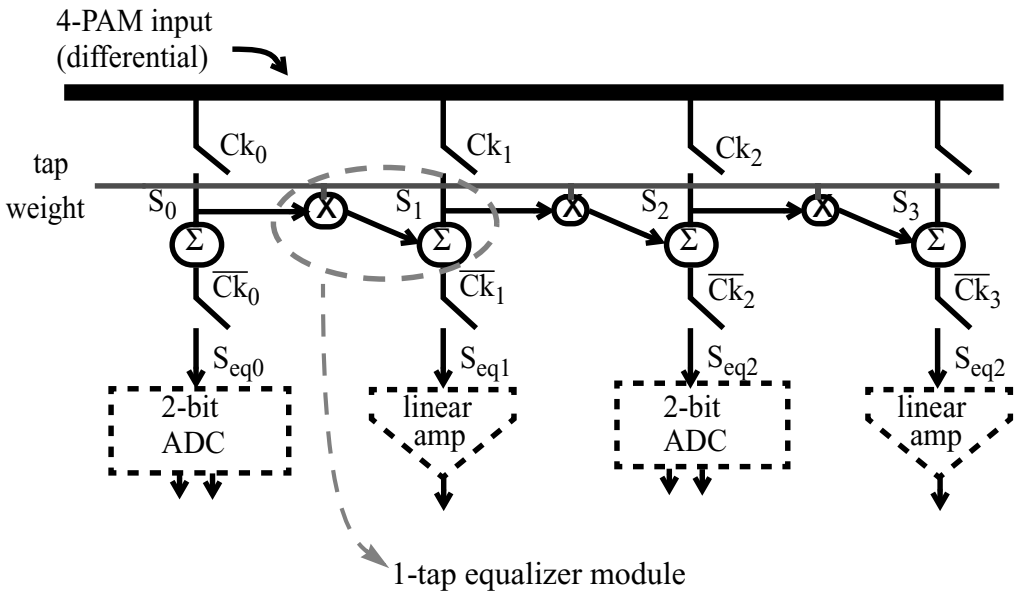

Figure 3.34: Block diagram of the half symbol-spaced 1-tap equalizer integrated into the receiver demultiplexer.

(S_1) to implement the 1-tap FIR equation

$$S_{eq1} = S_1 - \alpha \cdot S_0 . \qquad [3.7]$$

The equalized sampled value is re-sampled and held by a second switch, which is clocked by $\overline{Ck}_1$, so that its value is preserved when the input sampler starts tracking the input signal in the next half cycle. Therefore, the filtering process must complete in less than half a clock cycle to provide enough set-up time for the second switch to sample a valid value for S_{eq}. At the target sampling clock frequency of 1GHz, the multiplication and subtraction operation has to be performed in less than 500 ns. The simplest way to perform this operation fast enough in the 0.3-µm CMOS technology is with a current-mode approach. In this technique, the equalizer uses two differential voltage-to-current converters (V-to-I) where the current in one is modulated by the filter tap weight, and the subtraction operation is done by summing the two differential output currents with opposite

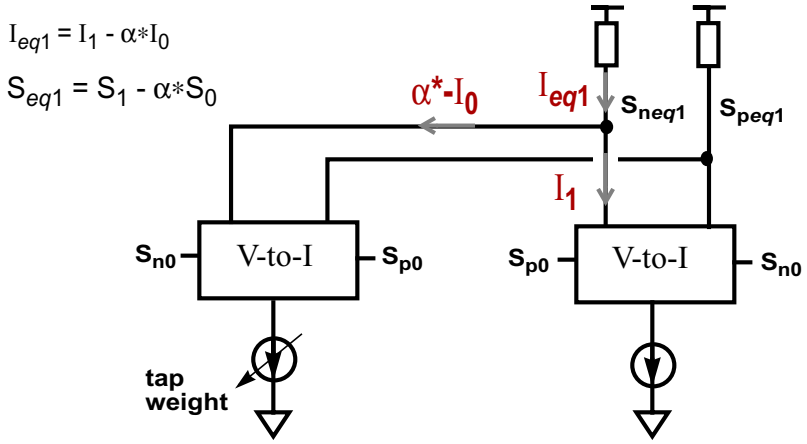

Figure 3.35: Current-mode 1-tap equalizer

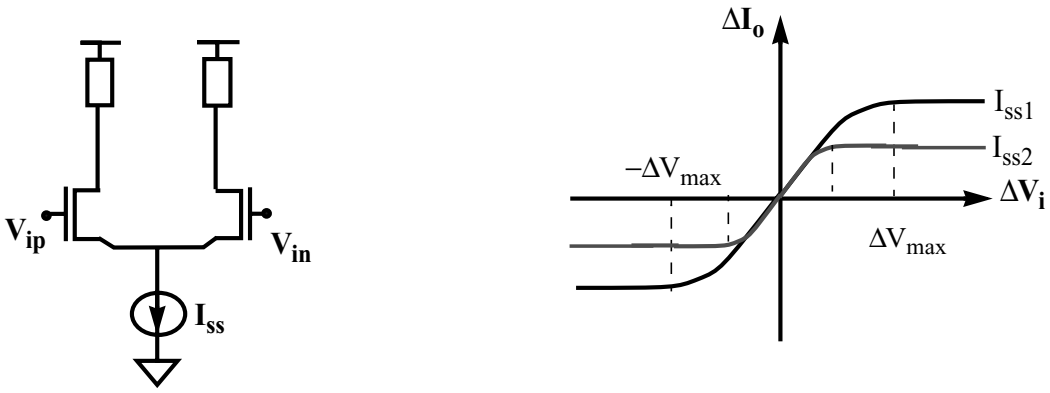

Figure 3.36: Coupled-source differential pair

polarity, as shown in Figure 3.35. One important requirement of the V-to-I converters is a large linear input dynamic range, since for the equalizer to function properly, the generated currents should be proportional to any received input voltage. As explained earlier in Section 3.1.4, the swing of the 4-PAM signal can be as large as 1.0V p-p single-ended (or 2.0V p-p differential),.

The source-coupled differential pair, shown in Figure 3.36, is a common type of voltage

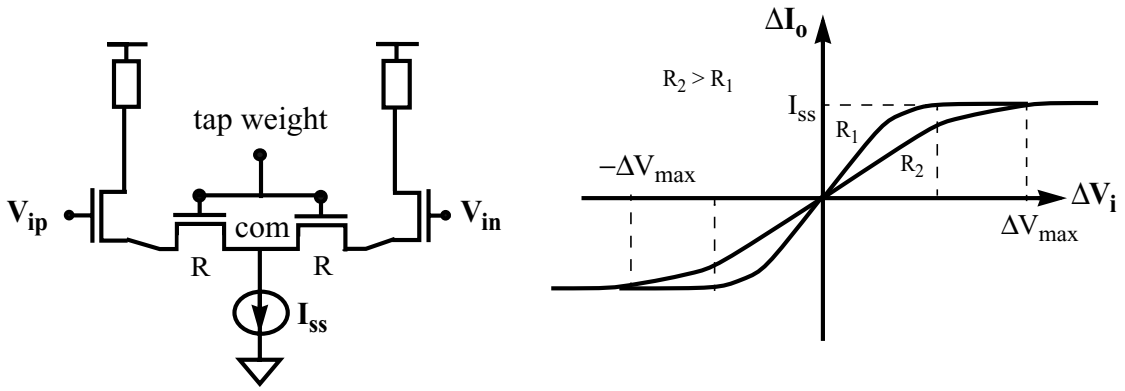

Figure 3.37: Coupled-source differential pair with source degeneration

to current converter. The input dynamic range[1] of a source-coupled pair with short-channel devices is [34]

$$\Delta V_{max} = \frac{I_{ss}}{k(W/L)} \quad . \qquad [3.8]$$

Because Equation (3.8) shows that the dynamic range of the source-coupled pair shrinks with the tail current I_{ss}, it is impractical to use this stage for the V-to-I converter with variable current (tap converter) and guarantee that the converter has large enough dynamic range for all practical values of tap weight (tail current).

In order to avoid input dynamic range reduction for larger values of tap weight, this design uses resistive source degeneration. The resistors are made of NMOS transistors operating in triode region such that their resistance value, and therefore the gain of the stage, is controllable for different tap values (Figure 3.37). The I-V characteristic in Figure

1. Input dynamic range of an stage is defined as the range of input values that the stage is sensitive to.

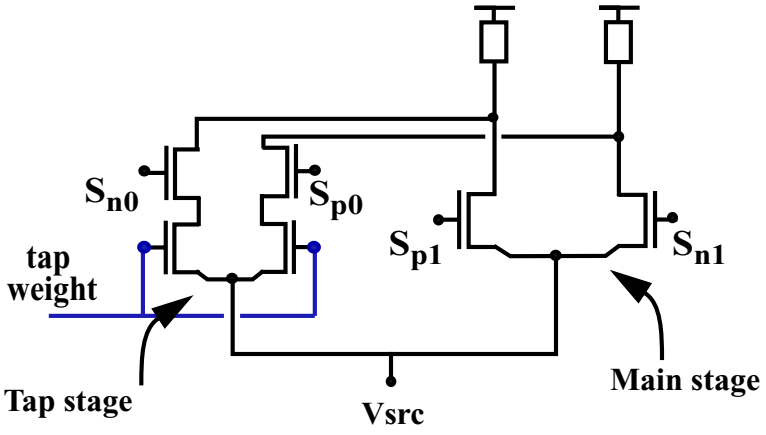

Figure 3.38: The proposed 1-tap equalizer with adjustable input dynamic range using off-chip voltage V_{src}

3.37 shows that the input dynamic range of this stage improves because reducing the tap weight increases the degeneration resistance. One consideration for this configuration is that as the differential stage common node, *com*, follows the input common mode, the tap weight voltage should be designed to follow the input common mode so that the gate overdrive voltage of the tap transistors stays constant and tap weight does not change with common mode variations.

One drawback of controlling the tail current I_{ss} of the differential pair is that the input dynamic range, based on Equation (3.8), varies with process. Hence, unless there is on-chip process compensation control circuitry, it will be hard to adjust the tail current to the right value. To reduce complexity in this test chip, such extra compensation circuitry was avoided, and the simple method shown in Figure 3.38 was used instead. In this scheme, the input dynamic range of the differential pairs can be adjusted by an off-chip control voltage V_{src}. However, note that the common-mode rejection of this design is not as good as sources-coupled stage with a tail current source.

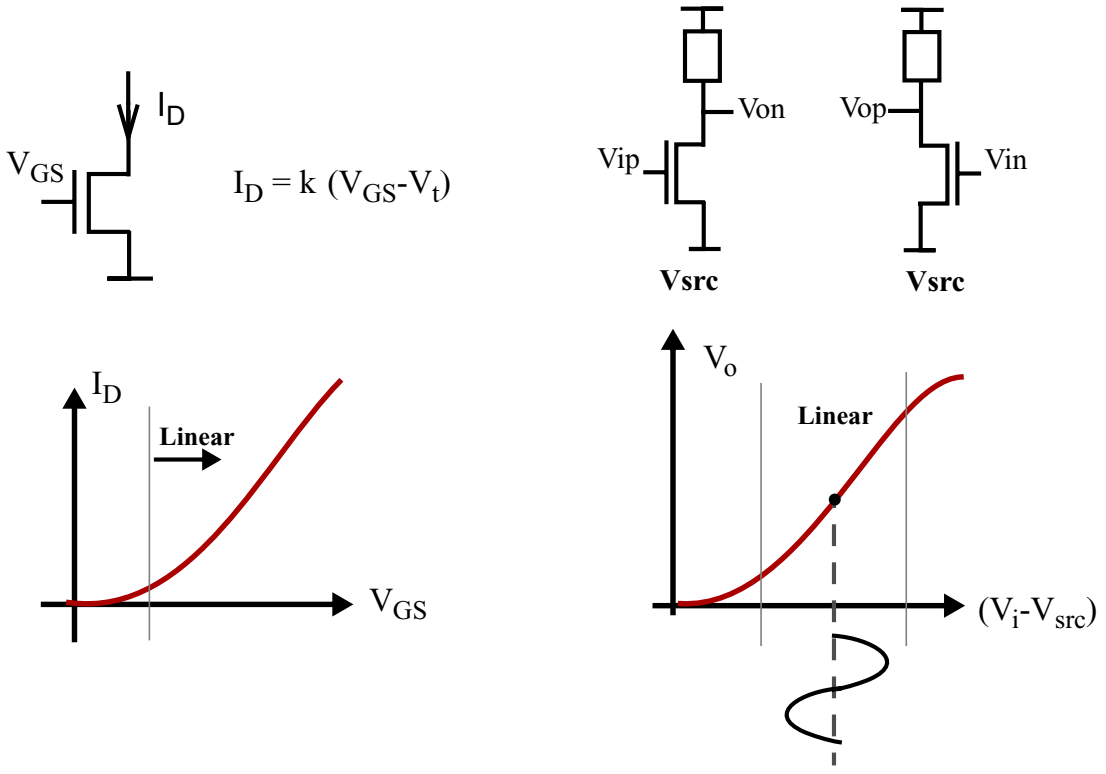

Figure 3.39: a) Linear I_D-V_{GS} curve for short channel NMOS, b) Linear V_o-V_i for common source amplifier due to short channel effect.

In order to maximize the gain of the main stage in Figure 3.38, the input transistors are used without source degeneration resistance. However, to achieve the required linearity range for these devices, the design takes advantage of velocity saturation in short channel devices. In a 0.3-µm long device, the drain current I_D is an approximately linear function of the gate-source voltage when the device operates in the saturation region [35]

$$I_D = \mu_n \varepsilon_c C_{ox} W (V_{GS} - V_{tb}). \qquad [3.9]$$

Hence, knowing the receiver input swing range, one can adjust V_{src} such that the input signal falls in the linear range of the input devices, as shown in Figure 3.39b. Note that although a lower V_{src} provides larger dynamic range due to larger V_{GS}, a large V_{GS} also

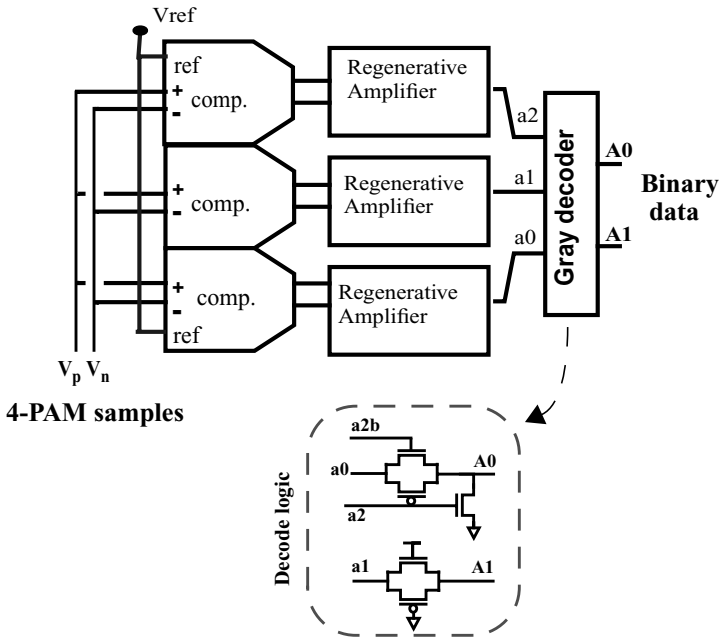

Figure 3.40: 2-bit differential flash analog-to-digital converter

results in large drain currents, which can generate a big enough voltage drop on the resistive loads to drive the input devices out of the saturation region. This effect is shown in Figure 3.39b where the V_o-V_i transfer curve flattens for large (V_i-V_{src}) values. Thus, at the optimum V_{src}, the input devices should be on the edge of turning ON for the minimum input voltage V_{i_min}, or $V_{src} = V_{i\_min} - V_{tb}$.

3.2.3 2-bit Analog to Digital Converters

Flash converters are used to convert the 4-level analog symbols into digital bits at a maximum frequency of 1GHz. Five 2-bit flash ADCs are implemented after the receiver 1:5 demultiplexer. Each ADC consists of three differential comparators and regenerative latches, followed by a thermometer code to Gray code converter, as shown in Figure 3.40.

In addition to improving the system BER, Gray coding has two advantages in receiver data conversion. First, it lowers the probability of metastability in the flash converter, as no

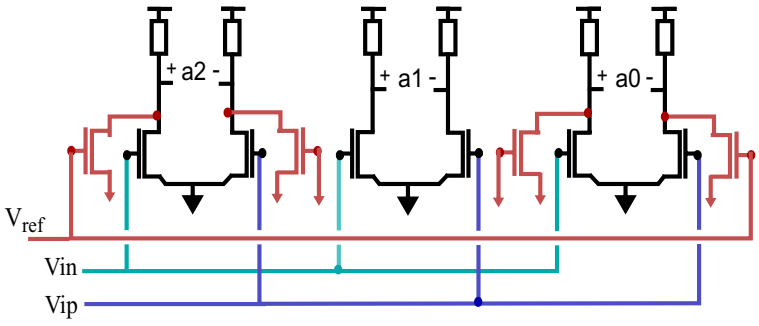

Figure 3.41: Three comparators of a 2-bit differential ADC with one reference voltage

signal is applied to more than one input, thus allowing the use of pipelining to increase the time for regeneration [28]. Increased total regeneration time increases the mean time between failures (MTBF) during data conversion, which can be calculated by [37]

$$MTBF = \frac{T_{data} \cdot T_{clk}}{(T_o \cdot e^{-\Sigma t_{ri}/\tau_i})} \quad , \qquad [3.10]$$

where T_o is a topology and process depend constant, τ_i is the regenerative time constant and t_{ri} is the regeneration time for each stage of pipelining.

Gray coding also makes thermometer to binary conversion logic simpler (Figure 3.40), reducing net delay and loading in the high-speed path and thus allowing more regeneration time in the pipelined stage.

Since a differential signaling scheme is used, only one reference voltage value is required to differentiate among the four input levels, as shown in Figure 3.41. Also shown is how this reference voltage is applied to the three comparators. The middle stage is a balanced differential comparator, and the two other stages are unbalanced by two

transistors, which are controlled by a single V_{ref}. To equalize the output capacitive loading of the unbalanced comparators, identical dummy transistors, with grounded gates, are placed on the opposite branch of the differential pair. As the reference line V_{ref} is a single-ended voltage, it should be well isolated to reduce the effect of noise coupled into this line.

The input referred voltage offset for the ADC increases with the number of stages before analog to digital conversion, because the mismatch in each stage statistically adds to the total voltage offset. According to Pelgrom [36] and more recent studies by Mizuno [38], the standard deviation of the offset is inversely proportional to the area of the device — the larger the device, the smaller the offset. Thus, to limit the total input referred offset to a reasonable limit, the device sizes in the sampler, equalizer, and ADC comparators and amplifiers should increase. This increase in device size translates into a larger total input loading for receiver. However, one can reduce the net mismatch by reducing the number of stage of the analog path. Therefore in this design, each one of the ADC's comparators is integrated into its preceding 1-tap equalizers, as shown in Figure 3.42.

The regenerative amplifier following the comparator is a slight variation of the comparator by Yukawa in [22]. When the clock is *HIGH*, the NMOS transistors in the middle of the stack disconnect the upper and lower cross-coupled devices to reset the latch. During this phase, the sampler is tracking, and the latch outputs are precharged *HIGH,* while the internal nodes, *a* and *b*, are precharged *LOW*. When the tracking phase ends, the cross-coupled devices are connected to enable positive feedback and amplify the difference at the sampled input. This regenerative amplifier is chosen because it has a faster regeneration [22] and a smaller input referred offset voltage than other differential

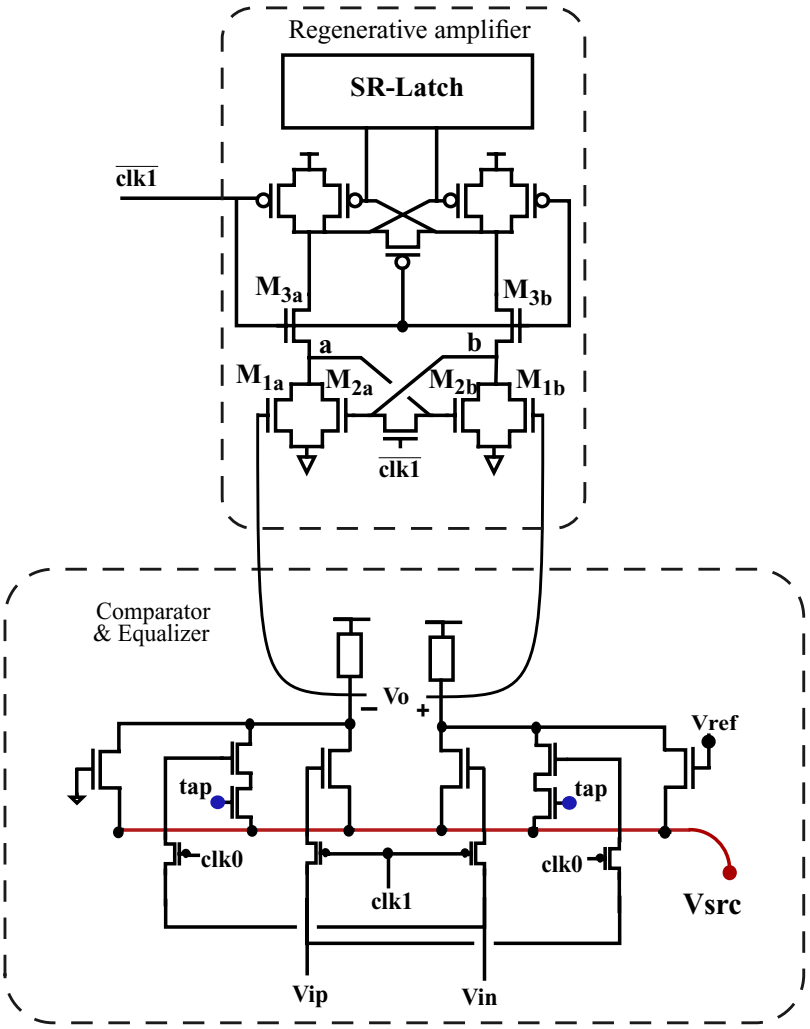

Figure 3.42: Combination of 1-tap equalizer with one ADC comparator.

latches as used in [11]. Faster regeneration time means a smaller τ in Equation (3.10), and thus a larger MTBF.

The input offset voltage is due not only to mismatched input transistors, but also to device and capacitive mismatch within the positive-feedback structure. However, the clock slew rate can be carefully adjusted so that by the time the cross-coupled devices turn on to start regeneration, the voltage difference established at the nodes *a* and *b* (Figure 3.42), is larger than the net offsets of the cross-coupled stage. As a result, the offset of the cross-

coupled stage does not have any impact on the final decision [39]. The capacitive mismatches can also be reduced by careful and symmetrical layout. Therefore, the input offset in this amplifier is mainly determined by the two input devices. The input refered offset measurements for the samplers are shown in Chapter 4.

3.3 Summary

This chapter shows how the low off-chip impedance (25Ω) of doubly terminated lines allows one to use large multiplexing and demultiplexing ratios at the chip pads to relax the on-chip speed requirements. The limitations of the transmitter are the transition times of the on-chip signal, and the inherent filtering due to the output capacitance, while the receiver sampling rate is limited by the minimum aperture of the sampling switches, which is defined by the RC product of the sampling switches and sampling clock slew rate.

Techniques for improving the transmitter and receiver I/O circuit implementation were introduced in this chapter as well. These include: an analog implementation for the transmitter 3-tap pre-emphasis circuit that reduces the complexity (power and area) of the FIR filter; a receiver equalizer that enables filtering the serial symbols at a high sampling rate; and, a differential comparator for the 2-bit front-end ADC that is integrated into the equalizer to reduce offset in the input analog path.

Chapter 4

Clock Generation and Timing Recovery

Phase spacing errors and phase noise (jitter) in the generated clocks, as well as phase alignment of the sampling clocks to the received high-speed data, all affect system performance. Therefore, clock generation and receiver timing recovery are important and critical functions in a high-speed signaling system (shown in bold in Figure 4.1). Section 4.1 describes the multiple-phase clock architecture used to multiplex and demultiplex the high-speed serial data. The clock generation takes advantage of phase locked loops with multi-stage ring oscillators to generate the required high-frequency clock phases from a low-frequency input clock. This chapter focuses on generating low-jitter clock signals with accurate phase spacing.

The sampling clock phases for the receiver must be aligned with respect to the input data stream to maximize timing margins. To perform this phase alignment, timing information must be extracted from the serial input. There are two distinct approaches for extracting timing information from random received data: oversampled phase detection

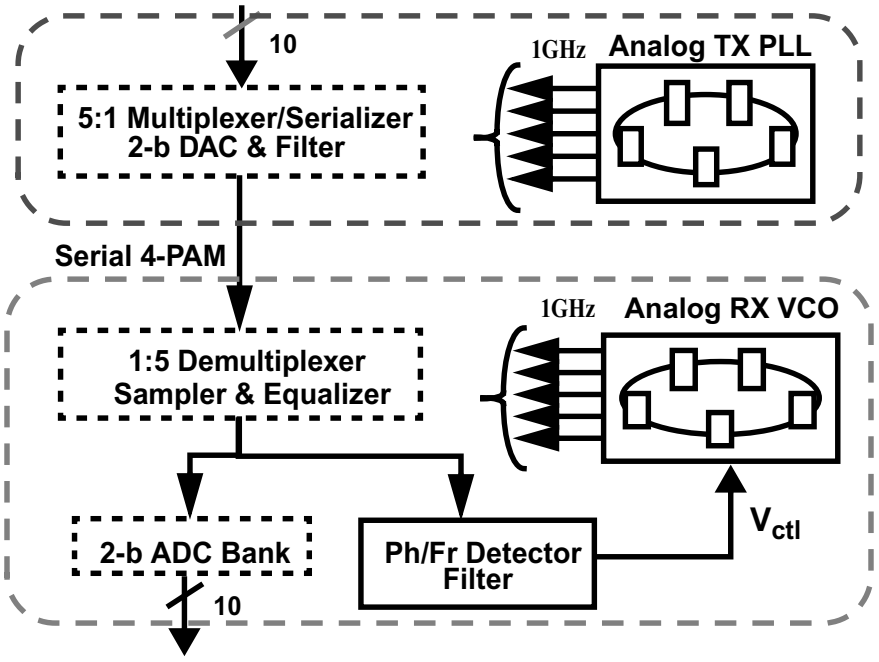

Figure 4.1: Transceiver top-level block diagram showing the clocking architectures

([40],[41],[42]) and tracking phase detection ([44],[45],[17]). This design uses a tracking approach, where the recovery loop adjusts the sampling clock phases continuously to align them to the center of the incoming data symbols. The data phase detector used is a new proportional phase detector that overcomes the main shortcomings of previous tracking methods. The detailed architecture of the proposed clock recovery and its trade-offs are discussed in Section 4.2.

4.1 Multi-Phase Clock Generation

This design uses a charge-pump PLL [48], together with a phase-frequency detector (PFD) and a ring-type voltage controllable oscillator (VCO) both in the transmitter and receiver. The VCO is a 5-stage differential ring oscillator that generates 10 evenly spaced clock phases (5 differential clock phases) for high-speed data multiplexing in the transmitter and demultiplexing in the receiver. The combination of a charge-pump

PLL and a PFD offers two important advantages over the simple XOR/LPF [43] approach: 1) the frequency capture range of the loop is limited only by the VCO output frequency range, not the phase detector range, and 2) the static phase error is zero if phase offset in the phase detector is negligible. The first advantage is important for PLLs with monolithic ring oscillators because their frequency can vary widely over process and temperature. The zero static phase error of this structure is critical in the receiver timing recovery loop, where any phase error reduces the timing margin for sampling.

A block diagram of a basic PLL is depicted in Figure 4.2. The divide by *N* in the feedback forces the on-chip frequency to be *N* times the off-chip frequency. The PLL locks the oscillator output clock to the external reference signal using a phase/frequency detector and a filter in a feedback loop. The phase/frequency detector produces current pulses that have widths proportional to the phase difference. The charge pump filter integrates these phase error pulses and servos the VCO drive the steady-state phase error toward zero. Because the input to the VCO controls frequency and not phase, the integration inherent in an oscillator adds a pole to the loop dynamics. Coupled with the

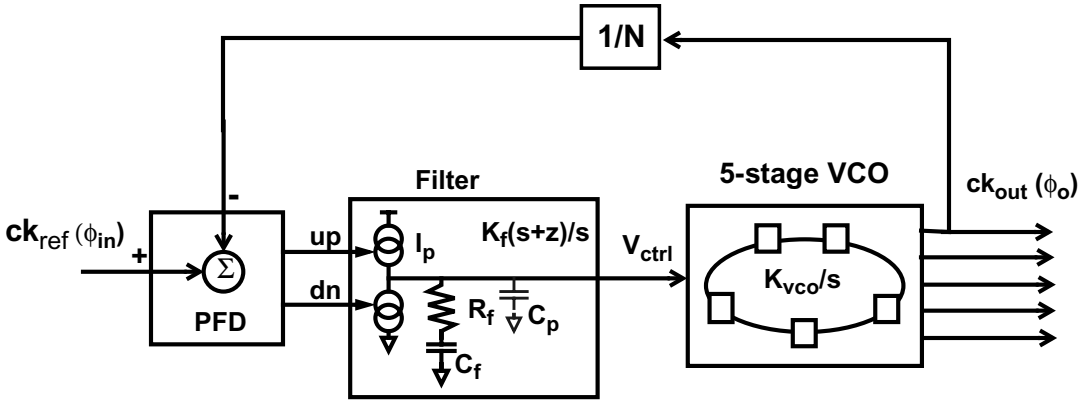

Figure 4.2: Multi-phase charge-pump PLL block diagram

loop filter pole, this results in a second-order system. By locking to a stable external signal, ck_{ref}, the active feedback of the loop compensates on-chip phase/frequency errors caused by temperature and supply variations, or other low-frequency noise sources. In the transmitter, the reference signal is an external low-jitter oscillator (e.g. a crystal oscillator), and for the receiver, the reference signal is the high-speed incoming data, or a combination of the latter with a reference clock ([44],[47]), where a more sophisticated phase/frequency detector is used (Section 4.2).

The primary concerns of this loop are stability and jitter performance. A loop that responds more quickly to noise can reduce the effective jitter of the output clock relative to the reference signal. A high loop bandwidth is thus desirable but achievable bandwidths are limited by the loop stability concerns. The following sections discuss the loop dynamics, its building blocks, and how phase noise depends on loop parameters and design.

4.1.1 Loop Dynamics

The PLL filter (Figure 4.2) integrates the charge pump current into voltage. As explained earlier, the charge pump and loop filter act as an integrator and therefore introduce a pole at DC in the loop, in addition to the DC pole due to the VCO. Therefore, a series resistor, R_f is added to provide a zero in the loop transfer function to increase phase margin [48]. Assuming the loop begins with a phase error $\phi_e = \phi_i - \phi_o$, the average current charging the capacitor is given by $I(\phi_e/2\pi)$. The average change in the control voltage, V_{ctrl}, of the VCO equals

$$V_{ctrl} = \frac{I\phi_e}{2\pi}\left(R_f + \frac{1}{C_f s}\right) = K_f \phi_e \frac{(s/z + 1)}{s}, \qquad [4.1]$$

where $K_f = I/(2\pi C_f)$ and $z = 1/(R_f C_f)$. Noting that $\phi_o = V_{ctrl} K_{VCO}/s$, we obtain the following open-loop transfer function

$$\frac{\phi_o}{\phi_e} = K_{loop} \frac{(s/z + 1)}{s^2} \qquad \text{(open-loop)}, \qquad [4.2]$$

and closed loop transfer function

$$\frac{\phi_o}{\phi_i} = \frac{s/z + 1}{s^2/K_{loop} + s/z + 1} \qquad \text{(closed loop)}, \qquad [4.3]$$

where $K_{loop} = K_f K_{VCO}/N$, is the loop gain of the whole loop. In all potential implementations, the control voltage node of the VCO, V_{ctrl}, has some parasitic capacitance, C_p. This capacitance introduces an additional pole which degrades the phase margin of the feedback loop, and thus potentially limits the maximum achievable loop bandwidth. However, in certain circumstances, this capacitor is used on purpose to produce a third-order pole that reduces the ripple on V_{ctrl} produced by the current switching of the charge pump [49].

The required loop-stability zero can also be implemented through feedforward [50], by adding a fast signal path in parallel with the main charge pump. Shown in Figure 4.3, this technique utilizes an auxiliary charge pump driving a wide-band RC network. This feedforward approach is preferable for certain filter implementations [46] discussed later in this section. Using this filter design, the open-loop transfer function is thus equal to

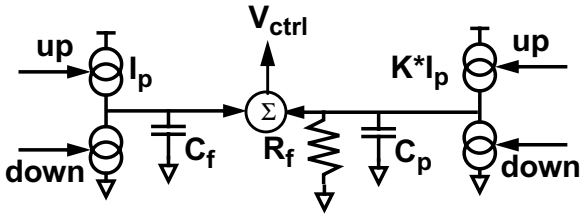

Figure 4.3: Loop filter with feedforward zero

$$\frac{\phi_o}{\phi_e} = K_{loop}\frac{(s/z+1)}{s^2(s/p+1)} \text{ (open-loop)}, \qquad [4.4]$$

where

$$z = \frac{1}{R_f C_p + K R_f C_f}, \qquad [4.5]$$

$$p = \frac{1}{R_f C_p}. \qquad [4.6]$$

The closed-loop bandwidth (BW_{loop}) is set by the loop gain, K_{loop}, and the position of the stabilizing zero. The third order pole, p, should be at least a decade higher than the loop stabilizing zero, and the open-loop unity-gain frequency ($=BW_{loop}$) should be equal to the frequency of maximum phase in order to obtain the necessary phase margin (ph.m. $> 45^o$). Figure 4.4 shows the magnitude and phase Bode plots for a charge-pump PLL with $p \simeq 14z$. Aligning the unity-gain frequency to the peak of the phase response also minimizes the phase margin variations due to the misalignments of the unity-gain frequency in the actual circuit.

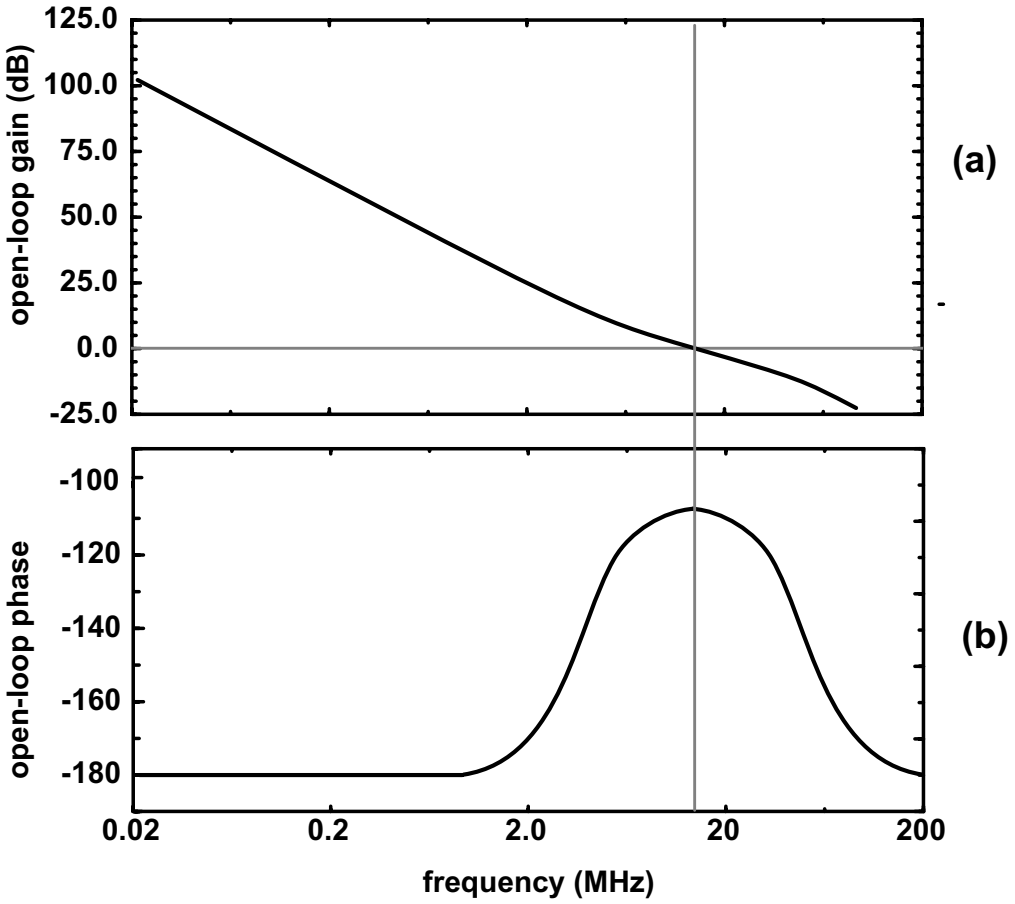

Figure 4.4: Open-loop Bode plots for a charge-pump PLL with third-order pole ($p \simeq 14z$)

Although designers aim for a high loop bandwidth and phase margin, several factors degrade the phase margin and limit the unity-gain frequency. First, because a PLL is a sampled system where new phase information arrives at the frequency of the reference clock (phase comparison rate), the sampling interval introduces a delay in the feedback loop. This delay lowers the phase margin of the system and degrades stability. Gardner [48], using discrete-time analysis, has derived a stability limit that can be reduced to

$$K_{loop} < \frac{\omega_{ref}^2}{\pi(R_f C_f \omega_{ref} + \pi)}, \qquad [4.7]$$

implying an upper bond on K_{loop} and $z(1/R_fC_f)$.

The phase margin can be further lowered by any additional buffering delay in the feedback path in addition to this inherent sampling delay. A total delay of T_d (seconds) in the feedback path with a bandwidth of BW_{loop} (Hz) reduces the phase margin (in degrees) by

$$\Delta\theta = T_d \, BW_{loop} \, 360^{o}. \qquad [4.8]$$

Therefore, the unity-gain frequency should be at least set an order of magnitude less than the reference frequency to avoid instability.

A second source of phase margin degradation is the third-order pole, p. When the crossover frequency is within 1/5 of the third-order pole, the phase margin degrades to less than 45^{o}. To set the appropriate unity-gain frequency, the filter gain K_f, is adjusted along with the zero position.

4.1.2 Phase noise in PLLs

The feedback loop tries to minimize the phase error between the reference clock and the PLL outputs. Therefore, the phased-locked VCO can reject on-chip phase noise at frequencies below the loop bandwidth (BW_{loop}) by tracking a clean input reference [51]. Figure 4.5 shows the VCO power spectral density (PSD) for two cases of free running (dashed), and phased locked (solid) oscillations that track the clean reference ($f_{osc}=N f_{ref}$) for frequencies $(f_{osc}-BW_{loop})<f<(f_{osc}+BW_{loop})$.

The integral of the area beneath the curve is an indication of the amount of noise power and hence jitter [52]. Assuming a very low-jitter reference clock, higher loop

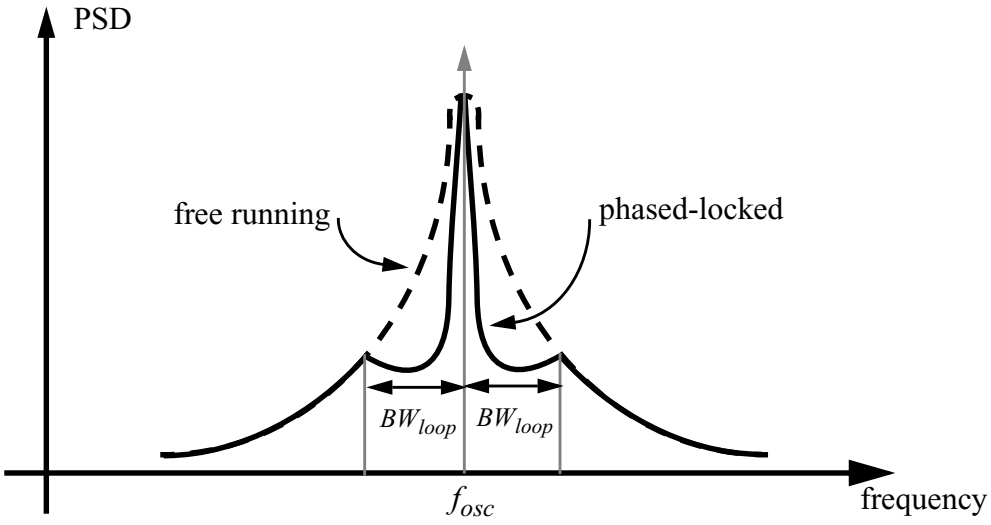

Figure 4.5: Effect of PLL on the output frequency spectrum

bandwidth can reduce timing errors substantially. However, the bandwidth of a sampled-system PLL is limited to below 1/10 the reference frequency (sampling rate), as explained earlier. This limitation forces designers to use additional techniques to cancel higher frequency noise in the loop.

The sources of noise that can introduce phase noise to the PLL output are the supply and substrate voltage noise, inherent device noise (thermal, shot noise, flicker) in the VCO elements, and periodic ripple on VCO control voltage. The following equation shows the output phase ϕ_{nout} (Figure 4.6) as a function of supply noise, v_{nsup}, substrate noise, v_{nsub}, and phase noise, ϕ_{ndev}, due to VCO devices [43]

$$\phi_{nout} = \frac{v_{nsup} s K_{sup}/K_{loop}}{s^2/K_{loop} + s/z + 1} + \frac{v_{nsub} s K_{sub}/K_{loop}}{s^2/K_{loop} + s/z + 1} + \frac{\phi_{ndev} s(s/z + 1)}{s^2/K_{loop} + s/z + 1}, \qquad [4.9]$$

where K_{sup} and K_{sub} (Hz/V) are the VCO sensitivity to supply and substrate noise, respectively. The first two terms exhibit a band-pass behavior where the noise is amplified

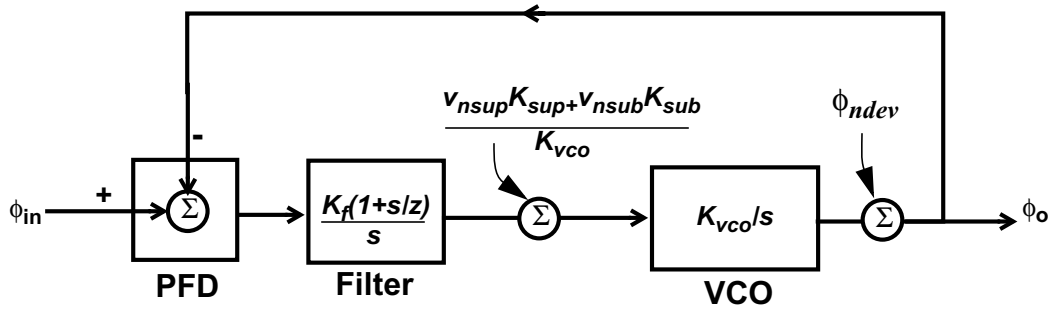

Figure 4.6: Block diagram of PLL with noise sources

at the loop bandwidth, also known as *jitter peaking*. The jitter in the previous cycle contributes to the starting point of the next clock cycle in a VCO. Therefore, for noise frequencies at the loop bandwidth, the phase noise of the previous cycle adds up constructively with the phase noise in the current cycle, resulting in the jitter peaking phenomenon. Low-frequency phase noise within the loop bandwidth is always cancelled by the feedback loop, and high-frequency noise is cancelled by the phase integration property of the VCO (K_{vco}/s). A replica feedback loop, proposed by Maneatis, is used in this PLL design to suppress supply and substrate noise for intermediate frequencies.

A PLL's response to inherent phase noise in the VCO (3rd term in Equation 4.9) is high-pass in nature, in which only noise (e.g. flicker) within the bandwidth of the loop is suppressed and any noise at frequencies above is passed to the output. To reduce this type of noise, the number of devices in the oscillator should be minimized and the quality factor (Q) of the oscillator should be maximized. For example, a multi-stage ring oscillator (Q=1) has a large device noise, while a LC oscillator (Q>>1) has much better performance

[52]. However, for mixed signal applications, the jitter contribution of device noise is negligible compared to that of supply and substrate noise from digital circuits.

In single-ended charge pumps, the resistor, R_f, in the loop filter can amplify ripple in the control voltage when the loop is locked. Since the charge pump's up and down current pulses (Figure 4.3) turn on at every phase comparison instant, any mismatch between the two currents flows through R_f, causing a step at the output that modulates the VCO frequency, resulting in a phase noise. This effect is especially undesirable for PLLs with multiple-phase clocks or frequency synthesizers with frequency multiplying ratios (N) larger than one [53]. For example, in a high-speed receiver where the sampling clock is aligned to the center of the data to best detect the narrow data pulse, such a phase error degrades the detection performance.

4.1.3 PLL Building Blocks

This section describes the circuit for the main building blocks of the multi-phase clock generating PLL, as shown in Figure 4.2, and the techniques used to reduce the phase noise of the generated clock phases.

Five-stage voltage-controlled ring oscillator

The key element in a VCO-based PLL is the VCO. While there are many possible types of VCO, this work uses a 5-stage differential ring oscillator as shown in Figure 4.7 to generate ten evenly spaced phases (5 differential). In a ring oscillator structure, the delay of each ring element is controlled by an external voltage V_{ctrl}. Unfortunately, the delay is also affected by the supply and substrate voltage, whose noise creates jitter

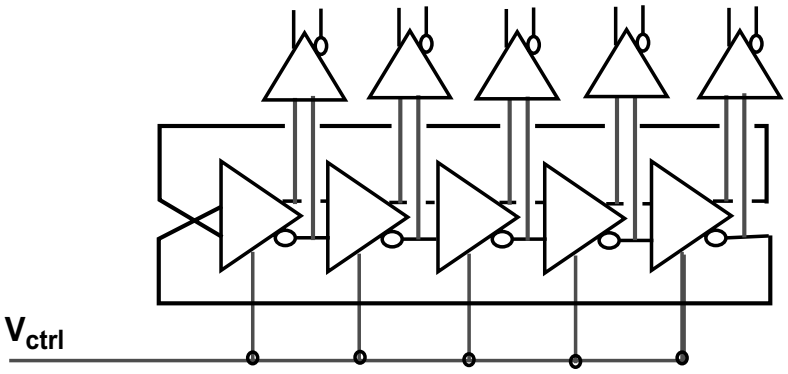

Figure 4.7: Differential 5-stage voltage-controlled ring oscillator.

according to Equation 4.9. In a ring configuration, these voltage variations effectively cause the frequency of the oscillator to change; the phase error introduced is the integral of this frequency difference and accumulates over many cycles. To reduce the jitter, the loop elements are designed for low supply and substrate sensitivity.

The output of the ring oscillator is typically buffered before being used by subsequent circuits. If the VCO design is inherently low in jitter, this buffering can often contribute a significant fraction of the total jitter because buffers such as CMOS inverters can be much more sensitive to supply noise than the source-coupled differential stages used in the VCO. However, using source-coupled differential stages to buffer up the signals is not efficient due to their much larger power consumption.

The differential delay elements of Maneatis in [54] are used in the oscillators of this design. The elements, illustrated in Figure 4.8a, are designed for low sensitivity of the delay to common-mode noise. Supply variations affect delay in these buffers because the current of the tail NMOS current source varies due to its finite output impedance. Similarly, substrate noise causes V_T variations in the NMOS current source device which

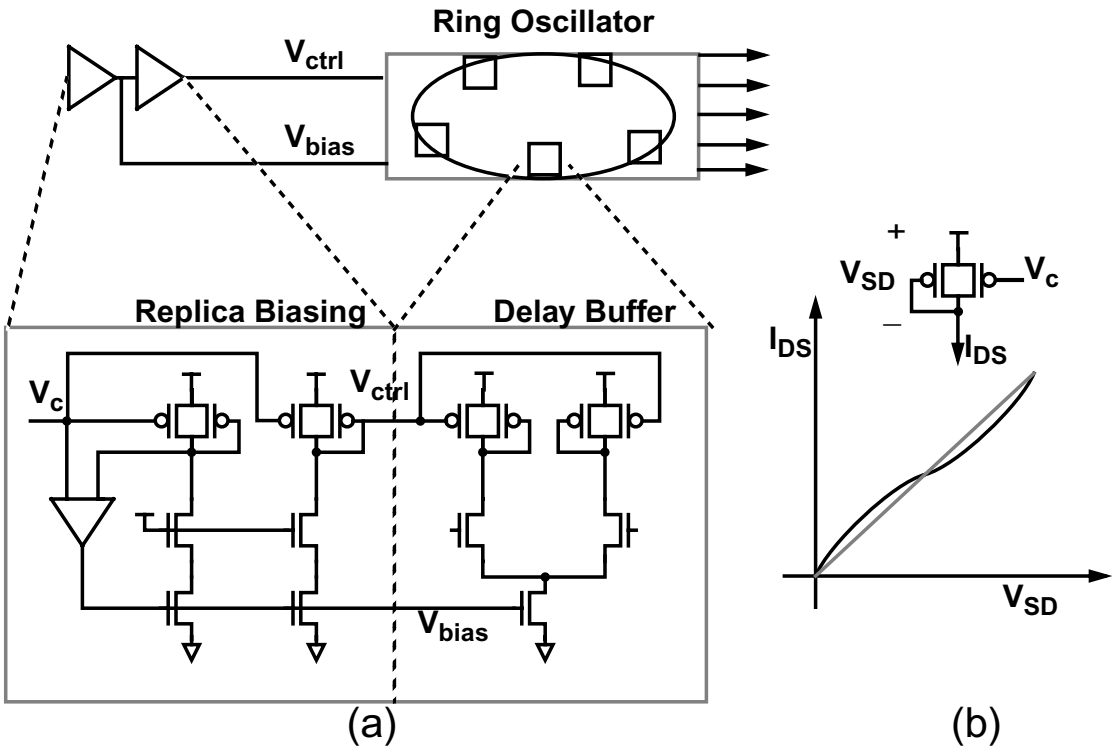

Figure 4.8: a) Delay element and its biasing scheme using a half-buffer replica, b) the load element I-V characteristics.

again affects the buffer current. Current variations change the V_{DS} of the load element and hence change the effective resistance and the delay. The load structure, composed of two PMOS devices, is insensitive to these variations to first order due to the swing limiting rovided by the diode-connected device with symmetric I-V characteristics [54], as shown in Figure 4.8b. The current variation in the tail current also changes the switching threshold of the source-coupled differential pair, affecting the delay through the stage.

Therefore, a replica-bias loop is used to maintain a constant tail current in the presence of supply and substrate noise. The V_{bias} is generated using an amplifier that adjusts the current of a half-buffer replica and the delay elements. The current through the replica controls the output voltage, V_{ctrl}, until it matches the voltage from the loop filter, V_c. This feedback loop effectively improves the output impedance of the current source

and adjusts the current for substrate variations. However, note that substrate noise cancellation for the VCO delay stages is only accomplished successfully if the substrate noise in the replica stage and VCO is exactly the same. Fortunately, this assumption is satisfied for epi processes with heavily-doped substrates, where oscillators suffer most from substrate noise [55].

Since the replica biasing feedback loop adjusts the buffer currents to compensate for noise, a high bandwidth is essential for good supply rejection. The replica bias circuit is a two-pole feedback loop, with one pole at the output of the replica amplifier and one pole at the feedback to the amplifier. The maximum bandwidth of the replica loop with acceptable stability is limited by the higher-frequency pole at the amplifier input and can be adjusted by properly sizing the replica half circuit. Note that increasing the size of the replica half circuit is effective before its self capacitive loading dominates. In the 0.3-μm process technology used, the simulated closed-loop bandwidth of this loop is around 700MHz for a 1GHz oscillation frequency.

For still better supply rejection, V_{ctrl} is carefully routed so that all capacitance is only to V_{DD} and the V_{GS}-V_T stays constant even with supply noise. Any ground noise coupled to V_{ctrl} through stray capacitance would induce an error voltage that would cause a frequency error. The same care is taken in routing V_{bias}.

To suppress the effects of substrate noise, the tail NMOS current sources should have enough substrate contacts electrically shorted to the source (local ground) and close to the source diffusions. A large P+ diffusion guard ring around the VCO that is connected to a dedicated quiet ground and located inside a larger Nwell guard ring helps significantly to reduce the substrate noise in a lightly-doped bulk process [56].

The 5-stage ring oscillator in a 0.3-μm process technology has a simulated supply sensitivity as low as 0.05%/% (percentage change of the VCO frequency per percentage change of supply variation) using source-coupled stages with the replica bias circuit. This sensitivity number (0.05%/%) doubles in the absence of the replica circuit. The simulated substrate noise sensitivity is 0.1% per 100mV of substrate noise with the replica compensation circuit, and this number increases to 0.25% when the replica loop is deactivated.

Static phase errors (phase offsets) in generated clocks also degrade system performance by reducing the timing margins. One source of static phase error is the unequal spacings between the multiple phases of the ring oscillator, caused by the mismatch between the ring delay stages. This mismatch is more evident at low frequencies where V_{ctrl} and V_{bias} are low, and thus the effect of V_t mismatch on the stage delay is larger. To reduce this phase error, the VCO should be tuned properly so that the control voltages V_{ctrl} and V_{bias} are large enough[1] within the frequency range of interest. In addition, these stages should be laid out symmetrically to reduce any physical mismatches.

Phase detector, charge pump and loop filter

Figure 4.9a shows the functional diagram of the phase detector used in the PLL. Whichever edge arrives first asserts the corresponding *Up* or *Down* signals causing the charge pump to supply or extract charge. The later edge would assert the other signal.

1. The difference between the desired control voltage (V_{ctrl} or $V_{bias)}$ and V_t must be much larger (factor of 10) than the V_t offset such that the resulting buffer delay mismatch is negligible.

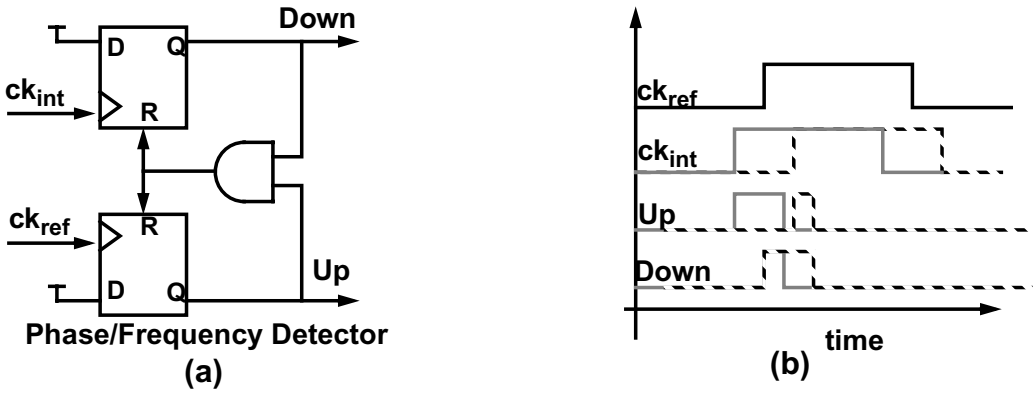

Figure 4.9: a) Phase/frequency detector schematic, b) phase detector timing diagram.

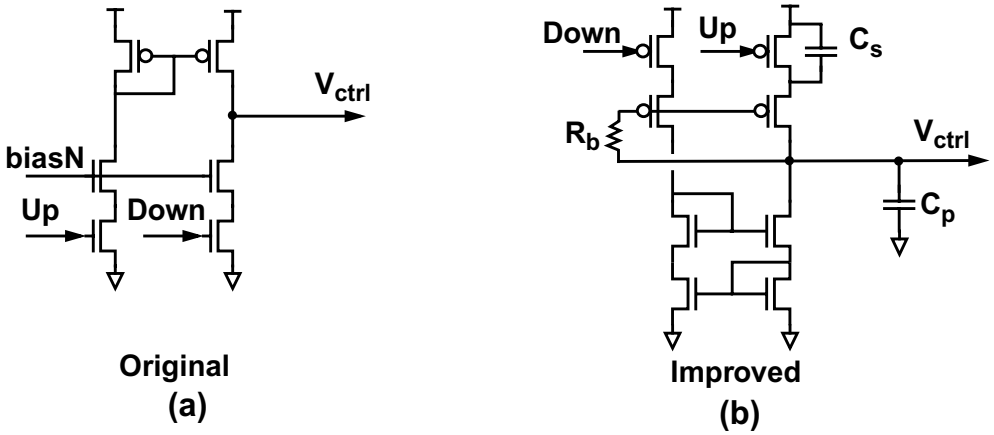

Figure 4.10: a) Original charge pump, b) improved version

After a short reset delay both *Up* and *Down* signals are deasserted by the reset signal. This action also aids acquisition since large differences in frequency would cause *Up* or *Down* to be asserted continuously. The output of the phase detector switches the current of a charge pump, shown in Figure 4.10a, that adjusts the voltages in the filter. Figure 4.9b illustrates the timing of the phase detector outputs in response to a leading (grey signal) and lagging (dashed signal) internal clock phase. If the clocks are slightly out of lock, either the *Up* or *Down* pulse widens to inject a non-zero charge. When phase-locked, the

phase detector produces equal *Up* and *Down* pulses of width equal to the reset delay. By guaranteeing a certain minimum pulse width, the detector always provides phase information to the charge pump. This guarantee avoids a common problem in some phase detectors of not providing phase information when the input difference is small. If a phase detector has such a dead band, the corresponding uncertainty in the locking could result in additional jitter. It is important to note that, in practice, there is always ripple on the control voltage *Vctrl*, even in lock, due to mismatches in signal paths for the *Up*/*Down* current pulses in any practical charge pump (Figure 4.10). As explained earlier, this ripple, which is coherent with the input clock frequency, modulates the VCO frequency and thereby produces output jitter if input clock frequency is different from the VCO frequency.

To reduce control voltage ripple, the charge pump topology is turned upside down to perform current mirroring with fast NMOS devices and thus reduce the delay mismatch between the pulses, as illustrated in Figure 4.10b. Since there is still some delay due to the NMOS mirror, a capacitor C_s is placed on the *Up* pulse path and is chosen to equalize the delay of the *Down* and *Up* pulses. In addition, to reduce any remaining output current spike caused by delay differences, a capacitor C_p is placed at the output of the charge pump. The current sources are implemented with long-channel PMOS devices and the current mirror uses a cascode structure for large output impedance, so that the *Up* and *Down* currents are equal for equal input pulse widths and different values of *Vctrl*. Otherwise, the feedback loop adjusts different widths for *Up* and *Down* pulses so that the net charge injected by charge pump is zero when in lock. Such a mismatch in the widths of the two *Up* and *Down* pulses would be another source of ripple.

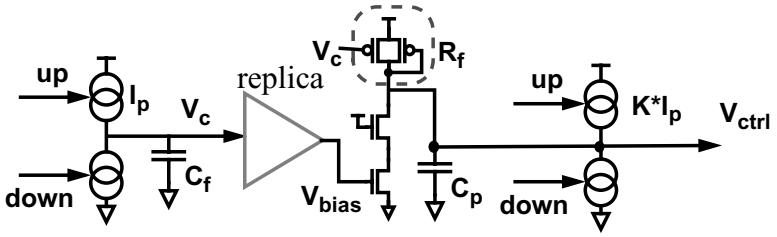

Figure 4.11: Feedforward filter implementation using replica biasing to implement the zero resistor.

The need for an extra circuit to provide the bias for the current source devices is removed in the new charge pump by self-biasing the PMOS current sources with the control voltage *Vctrl*. However, the main advantage of the self-biasing scheme is reduction of current (gain) variations in the charge pump with process and temperature, because the PLL feedback loop compensates for these variation by adjusting the loop control voltage [46]. Also, this biasing always guarantees that the PMOS current sources are in saturation region for all values of *Vctrl*, improving the output impedance of these current sources.

Since resistors in most CMOS digital technologies have large variations, a nominal loop bandwidth less than 1/20 f_{in} is often used to ensure stability. Here, however, the filter resistor is a VCO load element (Figure 4.8b) whose resistance over process and temperature is compensated by the PLL feedback. As this resistor requires a DC bias current, the original filter topology shown in Figure 4.2 cannot be used. Therefore, this design uses the feedforward method (Figure 4.3) to implement the loop zero. This design uses the feedforward structure proposed by Maneatis [46] that is implemented by integrating the replica bias circuit in Figure 4.8a with two charge pump circuits, as shown in Figure 4.11. The other advantage of this approach is that the loop zero is always placed accurately as a fixed fraction of the operating frequency, since the resistor element in the

filter is a half-buffer replica and its value is proportional to the frequency of operation [46]. Therefore, the loop phase margin does not vary considerably over the operating frequency range.

4.2 Timing Recovery

The previous section described the generation of precisely-spaced multiple clock phases. This section describes how to align the clock to the incoming data in the receiver. The task of the timing recovery circuit is to recover phase and frequency information from the input by extracting the clock from transitions in the data stream. In this work, the optimal sample point is about midway between the possible data transition times. Noise and mismatch inherent in the timing recovery circuit introduce jitter and static phase offset, degrading the timing margin. Moreover, this clock extraction is made more difficult by transmitter jitter which introduces uncertainty in the transition points that the timing recovery circuit must filter or track.

There are two main approaches for recovering timing from a serial data stream: phase-picking timing recovery by a digital PLL using oversampling as in [10] and illustrated by Figure 4.12a; and tracking phase detection by a data recovery PLL that aligns the sampling clocks to data, as shown in Figure 4.12b. In the phase-picking technique each transmitted symbol is sampled N times ($N \geq 3$) by multiple clock phases and the sample that is closest to the symbol center is selected by logic as the data. This approach allows very fast timing recovery, but suffers from large input loading (due to the large number of samplers), and phase quantization error. Furthermore, it requires complex

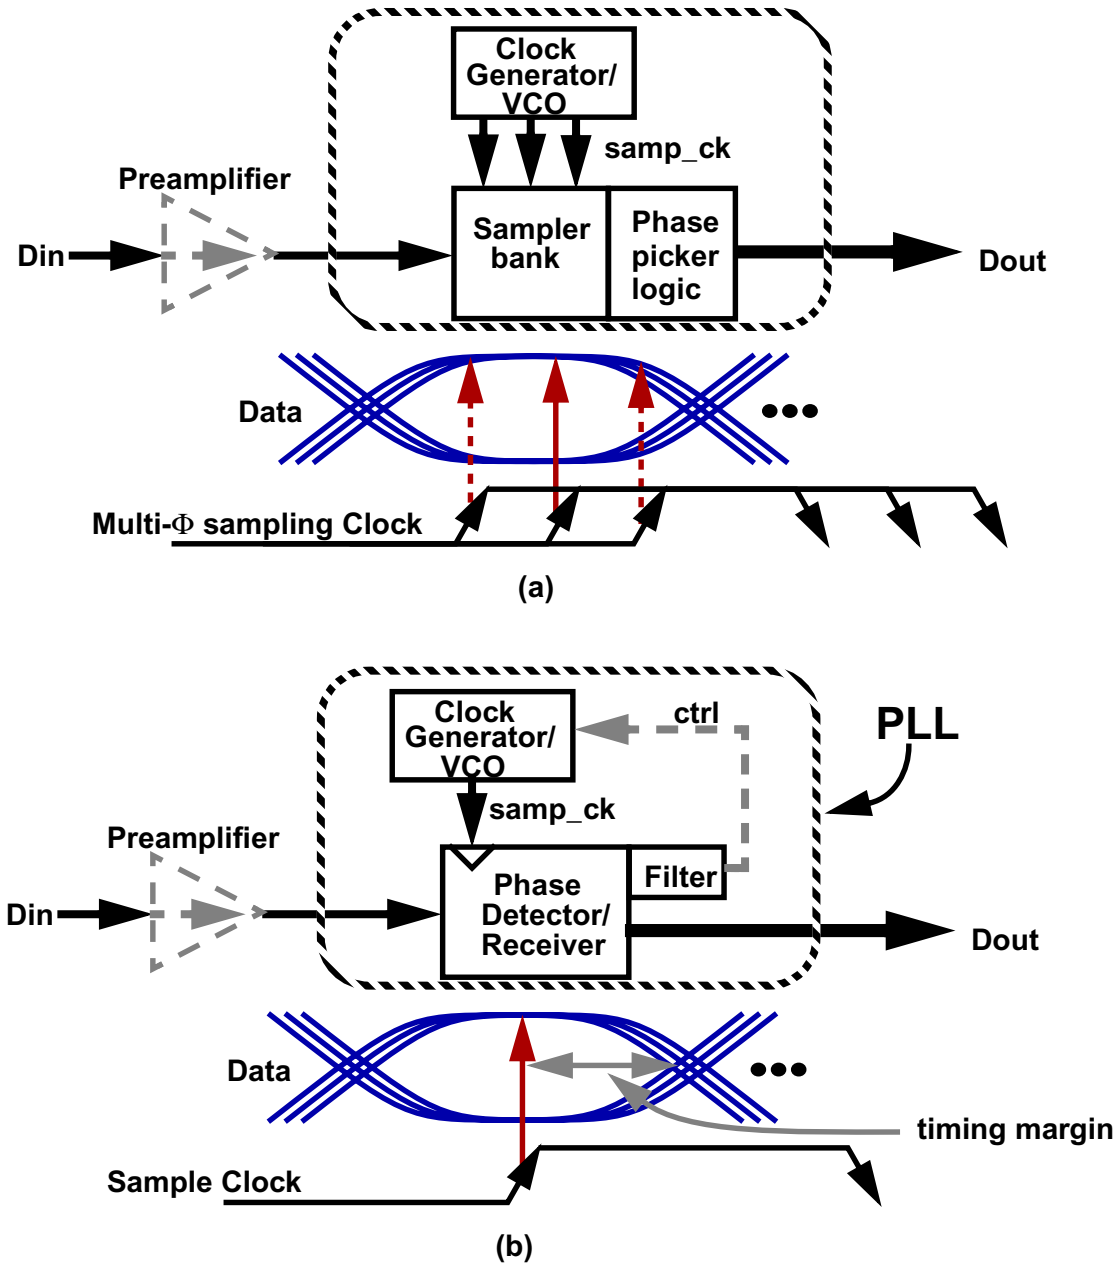

Figure 4.12: Data recovery architectures using a) 3x oversampling, b) phase-tracking PLL.

logic to process many samples at high frequency. In the tracking phase detection technique, a data phase detector measures the phase difference between the transition edge of the transmitted symbol and the sampling clock. This error value is used to align the sampling point to the symbol center. Traditional proportional tracking data PLLs offer good loop stability and bandwidth, but most suffer from a systematic phase offset and,

more important, they require a sampling clock frequency equal to the high serial data rate [45],[49].

Sampling transitions by the same mechanism as for the symbol centers reduces the systematic phase offset in data recovery. In addition, sampling and demultiplexing the fast serial symbols at the transitions and centers allows the phase comparison to be performed at a lower speed, or a given phase comparator to function at higher data rates. However, conventional sampling digital loops use bang-bang control, resulting in limited bandwidth and stability, as well as large control voltage ripple [17],[11]. It is also very important for the tracking bandwidth of the data recovery PLL to exceed the bandwidth of the transmitter PLL, because jitter in the transmitter VCO peaks at its loop bandwidth. If the bandwidth is less than the transmitter PLL, the transmitter phase noise would appear as a peak-to-peak timing error at the receiver.

This section shows the design of a novel proportional sampling phase detector for data recovery that overcomes the problems and limitations of traditional approaches. The phase detector takes advantage of a new frequency acquisition technique to lock to the input stream despite the large frequency variation of the on-chip VCO.

4.2.1 Proportional data recovery phase detector

Figure 4.13 shows the 2x oversampling receiver front-end that is part of the phase detection scheme. When the receive PLL is locked properly to the input data, half of the 10 samples represent symbol values at the center of the symbols (S_c), and half are samples at the data transition edges (S_e). The S_c samples are digitized by 2-bit flash ADCs and result

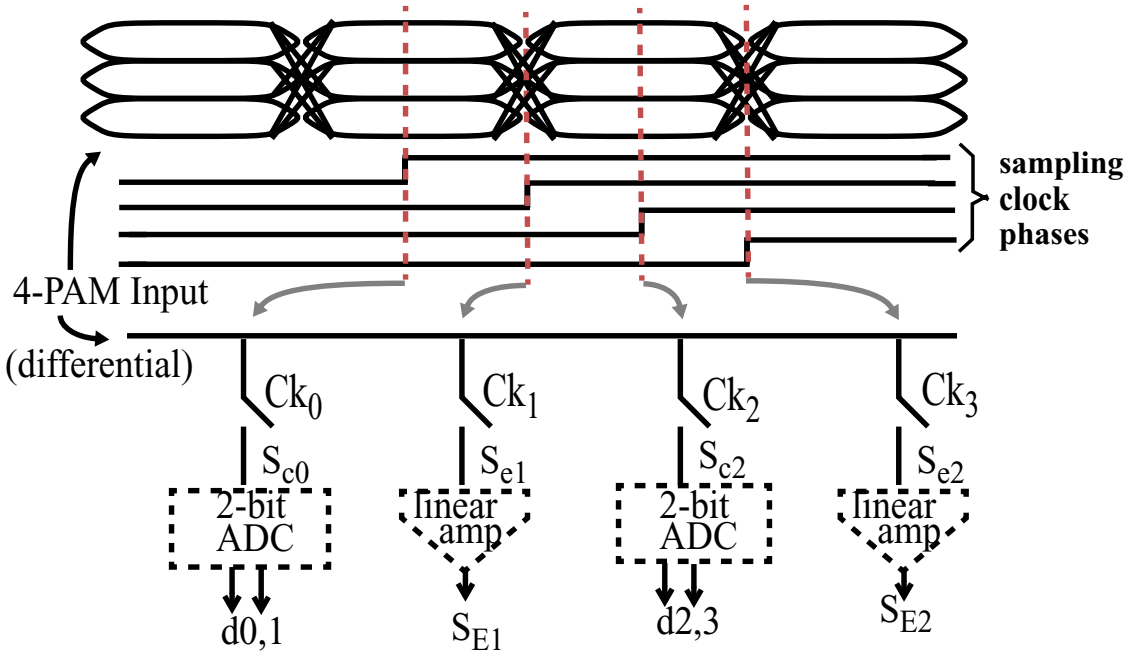

Figure 4.13: 2x oversampling receiver front-end

in the received data bits that are next resynchronized to a global clock, as explained in Chapter 3. The S_e samples are amplified linearly and kept as analog values (S_E) for use by a linear phase detector for timing recovery. The front-end equalization modules are not shown in this figure for simplicity.

Figure 4.14 illustrates the proposed phase detection method for a special case of 2-level data and lagging sampling clock. Arrows in the figure show the clock sampling points only at symbol boundaries. When the loop is not in lock and a transition occurs, the edge samples are non-zero and a monotonic function of the phase difference ($\Delta\phi$) between sampling clock edge and data zero crossing. This function can be approximated by a linear function, when the sampling edge occurs within the data transition interval and the loop is near its locking point. Thus, for edge samples within this interval of interest we have:

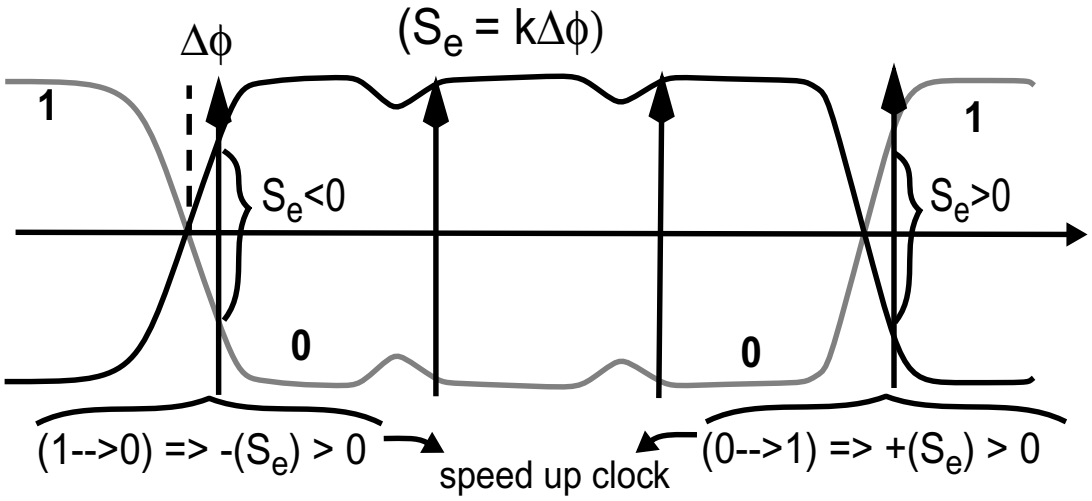

Figure 4.14: Proportional tracking phase detection method (sampling clocks lag the data)

$$S_e \approx k \cdot \Delta\Phi \,, \qquad [4.10]$$

where k is the slope of the transition edge. The *Se* values are added together with correct polarity, determined by the direction of each transition, and used to adjust the loop control voltage to correct for the phase error. Because the correction on the loop control voltage is proportional to the phase error, this method results in proportional loop control. This PLL therefore combines the advantages of both a linear and a sampling loop. Also the analog edge samples (S_e) at transitions are zero when in lock, nominally resulting in zero ripple on the loop control line. In bang-bang control, fixed-amplitude correcting pulses are always applied to the control line, resulting in ripple and, hence, timing error.

Note that in a differential 4-PAM stream, there are 3 distinct transition types (Figure 4.15). Of these 3 types, only *type1* makes a transition to the same magnitude but opposite polarity, where a zero crossing occurs exactly at the midpoint between two symbols and which therefore can be used for clock recovery. The two other types are ignored as they

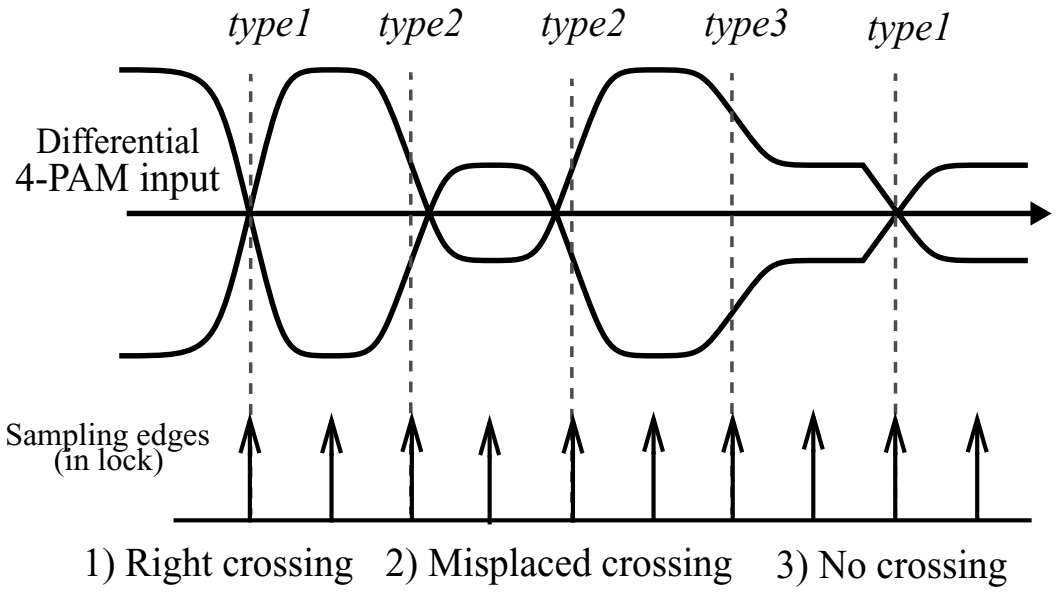

Figure 4.15: Three types of transitions in a 4-PAM symbol stream

convey wrong phase information. In every cycle (5 symbols), one *type1* transition is guaranteed by the transmitter's 4/5sym encoder illustrated in Chapter 3.

Figure 4.16 shows the block diagram of the data phase detector used here. The five amplified analog edge samples S_E, the amplified versions of S_e samples, are each fed into a switch structure and decision logic block of the phase detector. Based on the two 2-bit symbol values before and after the transition (e.g., d0d1 and d2d3), the phase detector adds the S_E values of *type1* with correct polarity to the control voltage of the loop and ignores the other two types of transitions by turning off all the switches of that stage. The add/subtract function is implemented by current summing the differential analog samples with correct polarity at the output of the phase detector (V_{ph}).

The response of the decision logic needs to be very fast, because the whole operation of input data regeneration and logic decision must be complete in less than half a clock cycle before the new S_E values are applied to the corresponding switch structure. Therefore,

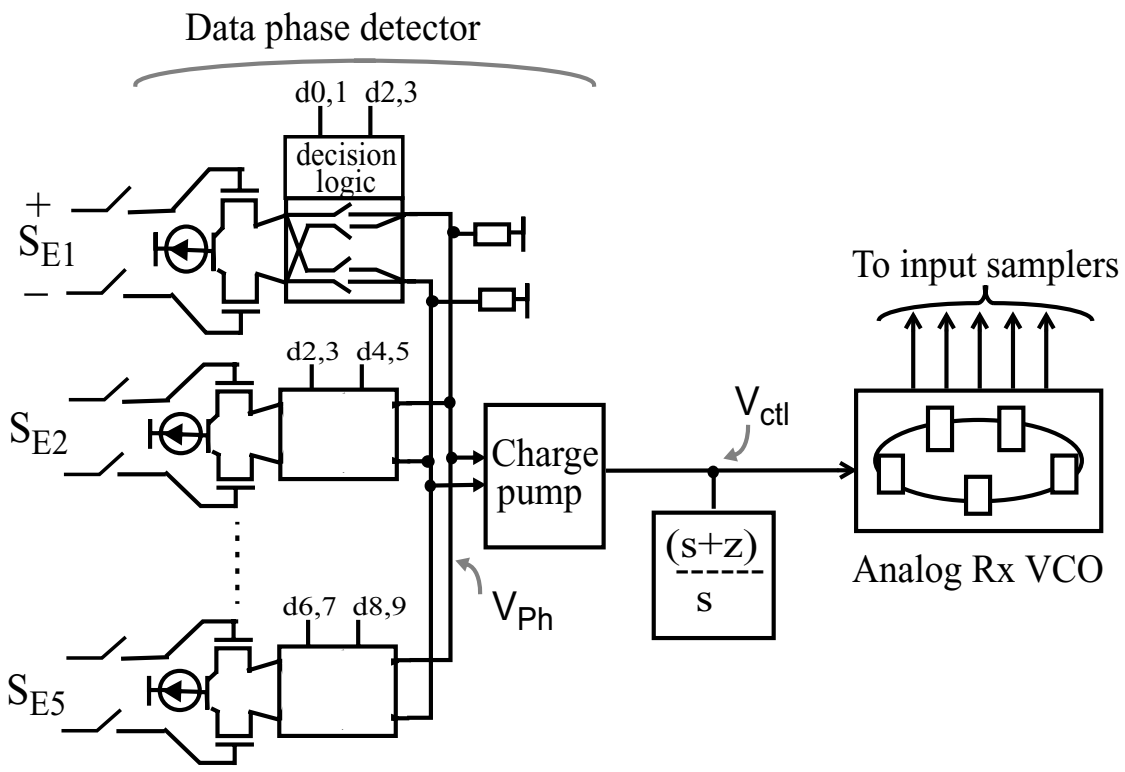

Figure 4.16: Proportional data phase detector architecture

the decision logic of Figure 4.17a is used to limit the delay to only a passgate. In order to reduce the delay from the sampler to the output of decision logic, this design uses the previous and present thermometer coded data values, $a_0a_1a_2$ and $b_0b_1b_2$ respectively, directly from input comparators before the thermometer-to-binary decoder. The two signals, *s1* and *s2*, are responsible for switching the polarity of the edge sample accordingly, and the two other signals, *s3* and *s4*, determine if any other invalid types of transitions have happened. Note that for a binary transmission only the logic for *s1* and *s2* is required.

Originally in this work, the switch structure shown in Figure 4.17b was used together with the decision logic to process each edge sample S_E, where the top switches change the polarity of the input, and the tail switches pass or block the input signal. For better speed

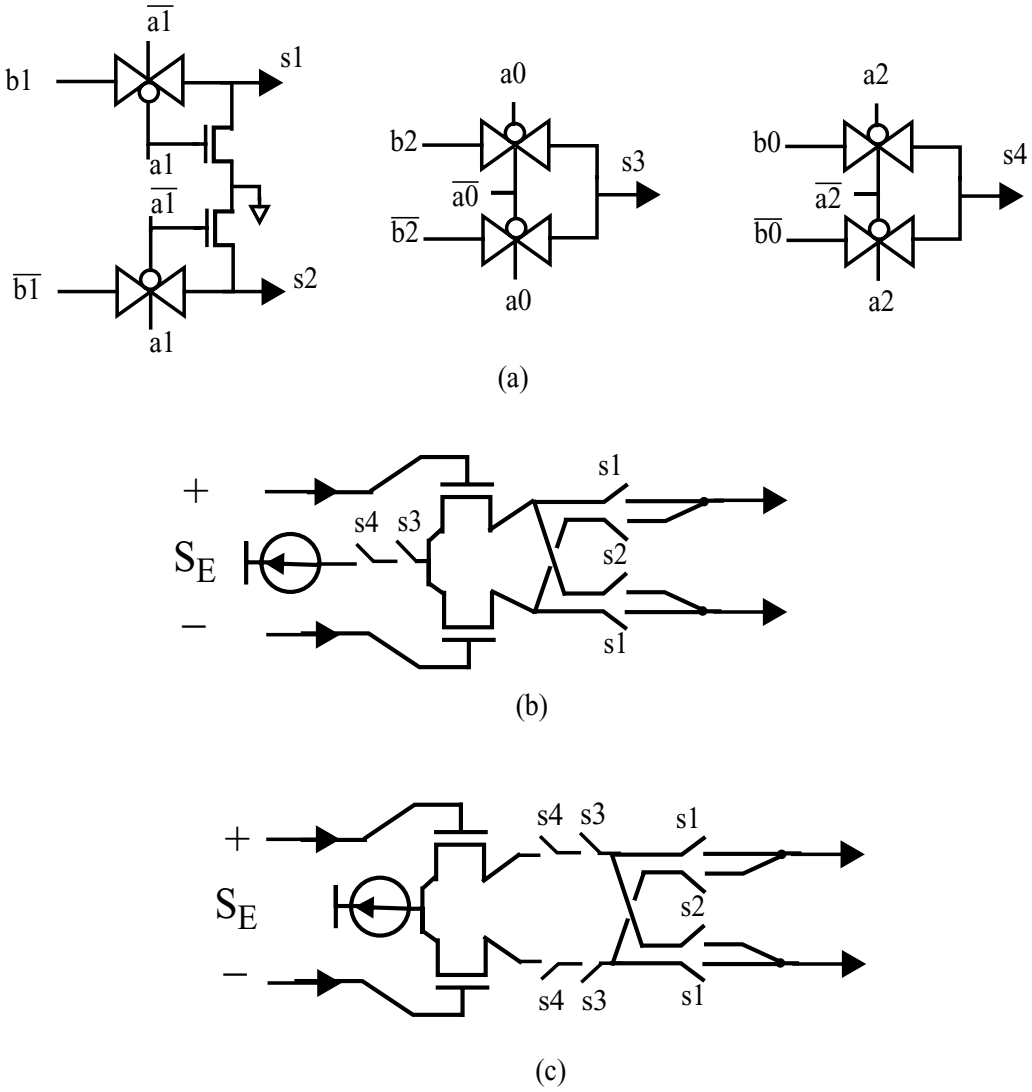

Figure 4.17: a) Phase detector decision logic, b) original switch structure, c) improved switch structure

performance all the switches are built by NMOS devices. The problem with this design is that when the edge sample S_E is invalid and either one of the tail switches is off, the toggling of switch *s3* pulls charge from the differential stage tail and turns it on randomly. The random charge is differentially injected to the output of the phase detector based on the invalid input voltage and results in noise on the output. This noise on the phase detector output translates into noise on the PLL control voltage and therefore jitter in the VCO. Also

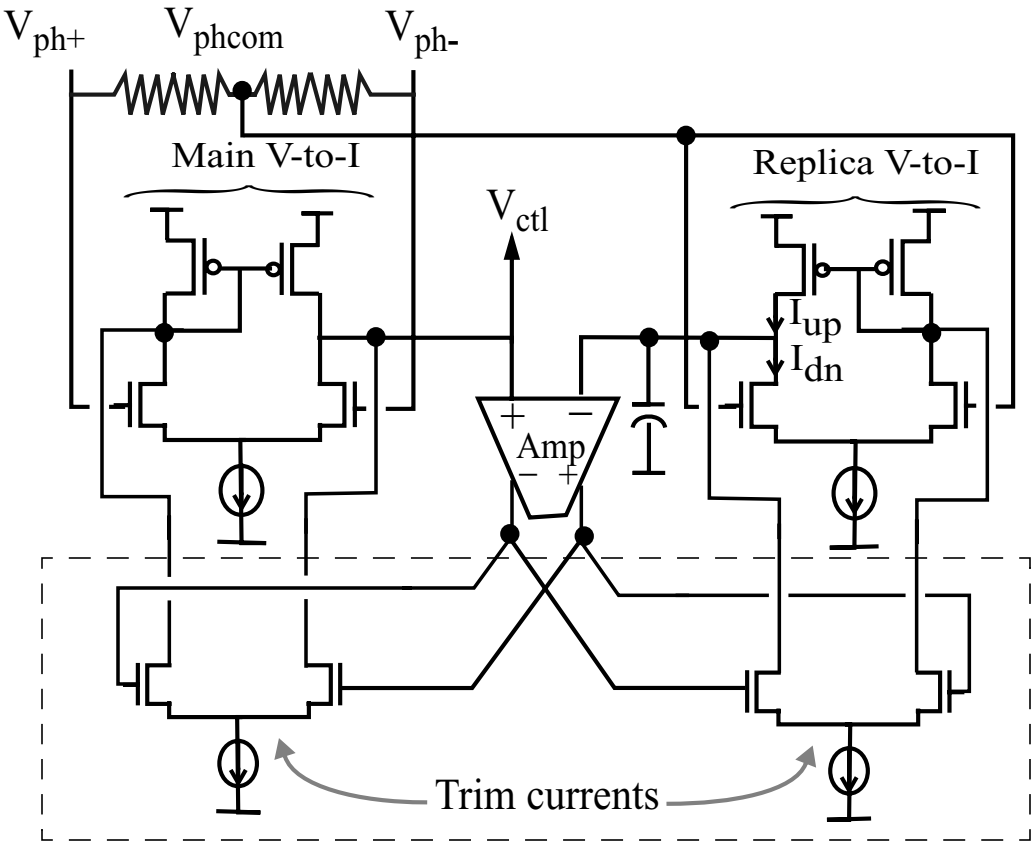

Figure 4.18: Charge pump offset calibration replica circuit

for certain input 4-PAM patterns, this problem is severe enough to force the loop to lose lock. To avoid this problem the structure in Figure 4.17c is used, where the tail switches are moved up to the top. Thus, when there is an invalid edge sample input, at least one of the two switches *s3* or *s4* is off and disconnects the input from output. The charge injection due to the switches in this design is also common-mode and is therefore largely canceled by the following differential charge pump stage.

The charge pump converts the differential signal V_{ph} to a proportional current using a differential voltage to current (V-to-I) converter, as shown in Figure 4.18. The voltage offsets in the charge pump, phase detector, and receiver front-end stages (equalizer and linear amplifier) directly translate into a phase offset between the sampling clocks and input

data. Random offset due to transistor mismatches is reduced by increasing the device sizes and careful layout. The systematic offset of the charge pump (V-to-I) is cancelled using an offset calibration loop that forces the charge pump to inject zero net charge (current) into the loop filter when V_{ph}'s differential voltage is zero, as shown in Figure 4.18. The calibration circuit has an exact replica of the main V-to-I converter, whose inputs are tied together and set equal to V_{ph}'s common mode voltage (V_{phcom}). The replica V-to-I converter has a capacitor at its output that acts as a charge integrator. Thus, the source (I_{up}) and sink (I_{dn}) currents should be exactly equal to avoid charging the replica circuit output to either of the supply rails. The two I_{up} and I_{dn} currents are forced equal by a differential comparator and a current trimming circuit combination (Figure 4.18) that compares the replica output to the loop control voltage (V_{ctl}) and makes the replica output equal to V_{ctl} by trimming I_{up} and I_{dn}. A replica of the trim currents is applied to the main V-to-I converter. As a result, the loop charge pump generates equal I_{up} and I_{dn} currents when its differential inputs are equal, i.e. when V_{ph}=0, or $V_{ph+} = V_{ph-} = V_{phcom}$.

To make loop parameters (gain, bandwidth, phase margin) track process variations and frequency of operation, the feedforward filter design by Maneatis [54] is used. As there is ideally no ripple on the loop control voltage when in lock, the loop filter does not require a third order pole capacitor to damp the control voltage ripple. Therefore, the loop theoretically has only two poles and one zero, and its gain-bandwidth product is therefore limited only by Equation (4.7) due to its discrete-time nature. However, the capacitive loading of the VCO stages on the loop control line introduces a third order pole that can make the loop unstable for very large gains. The loop gain, and consequently the bandwidth, increases with the number of useful (*type1*) transitions per cycle, and the slew rate of input data signal (k in 4.10). Using the 4/5sym encoder, the density of *type1*

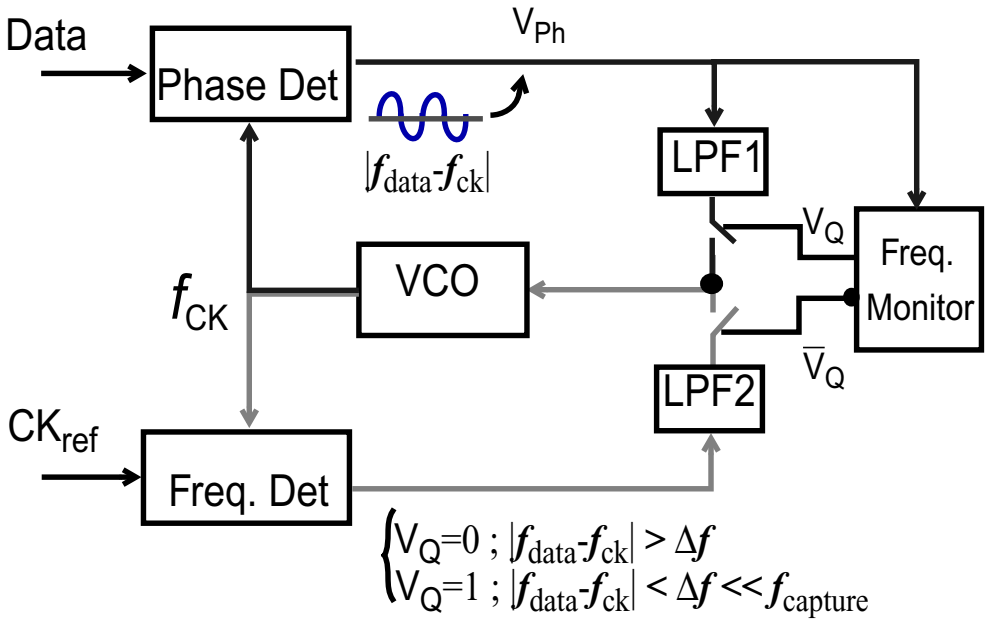

Figure 4.19: Frequency acquisition loop for data phase detector

transitions varies from a minimum of one to a maximum of 5 per clock cycle. The input slew rate (k) is determined by the signal amplitude and transition time, which is limited by channel bandwidth. Hence, the loop parameters are chosen carefully to guarantee a loop bandwidth of >20MHz and a phase margin >45° under the worst operating conditions (lowest and highest loop gains). The loop is optimized for a random data sequence with an average *type1* transition density of two per cycle, a differential input amplitude of 1V (500mV single-ended), and a risetime (10% to 90%) of 200ps. In this condition, the loop has a bandwidth of 35MHz (BW/f_{ref}> 0.07) and phase margin of 50°.

4.2.2 Frequency acquisition

As the phase detector has a limited frequency capture range, a frequency acquisition aid is employed at startup to help acquire lock to a local reference crystal clock, which typically has a maximum frequency error of +/-200ppm (Figure 4.19). As long as this frequency

error is smaller than the frequency capture range of the data recovery phase detector, phase lock to data can be performed successfully. Several approaches have been proposed in the past for frequency acquisition in data recovery, e.g. [57] and [58]. However, almost all of them require certain parts of the detector to run at the high-speed serial data rate (e.g. 5Gsym/s in this work), which is not practical for any of the present CMOS process technologies. The following paragraph describes the operation of a frequency acquisition loop that does not require any circuit to run faster than the external reference clock (e.g. 500MHz).

If the Rx-VCO frequency is initially different from that of the incoming data, cycle slipping occurs. During cycle slipping, sweeping of the clock phase causes the phase detector output (V_{Ph}) to oscillate between *Early* and *Late* signals. The frequency of this oscillation (sweep speed) is equal to the frequency difference between the receive sampling clocks (VCO frequency) and the incoming data. The frequency acquisition loop is activated at circuit startup, and a frequency monitor circuit keeps this loop activated as long as the frequency difference is large. When this difference gets smaller than the capture range of the PLL, the frequency monitor circuit activates the data recovery loop and deactivates the frequency acquisition aid. To guarantee successful phase lock by the data recovery phase detector over all corners, the frequency monitor does not deactivate the frequency acquisition loop until the frequency difference between data and VCO is below one-fourth the capture range of the phase detector.

The frequency acquisition loop (Figure 4.19) is identical to the PLL architecture used for clock generation, as explained in Section 4.1, where the frequency detector is in fact the phase/frequency detector of Figure 4.9 with the single-ended charge pump of Figure 4.10. Using a phase/frequency detector, instead of just a frequency detector, is necessary for the

frequency acquisition loop to guarantee zero frequency offset[1] between the VCO and reference clock, whose nominal frequency is equal to that of the input data stream.

Figure 4.20a shows the top view of the frequency monitor circuit. If there is a considerable frequency difference, the oscillations at the phase detector output (V_{Ph}) cause the edge detector to produce pulses that continuously discharge C_{one} and keep the one-shot circuit output (V_{one}) at zero. Once the VCO frequency is close enough to the incoming data frequency (within the data PLL capture range), the pulse rate of the edge-detector decreases such that C_{one} can charge high enough to switch V_{one} to one. At the rising edge of V_{one}, V_Q, which is reset to zero at start-up, is asserted and hands loop control over to the data phase detector. The maximum frequency difference that causes the frequency monitor to switch control between the two loops is determined by the capacitor C_{one}, pull-up current I, and the threshold of the following comparator V_{th}, as shown in the following equation

$$\Delta f = \frac{I}{C_{one} \cdot V_{th}}. \qquad [4.11]$$

The current source I is implemented in a manner similar to the charge pump current source in Figure 4.10b, using PMOS devices biased with a copy of the loop control voltage V_{ctrl} to generate a constant current over the process corners. In addition, by self biasing the current source with V_{ctrl}, at lower frequencies of operation (smaller V_{ctrl}) that the capture range of the phase detector reduces [63], the maximum frequency difference Δf, which triggers the frequency monitor to turn off the acquisition loop, reduces proportionally[2].

The edge detector is designed to have hysteresis (Figure 4.20b), using positive feedback

1. As phase is the integral of frequency, a phase-locked loop may have a phase error but zero frequency error. A frequency-locked loop may demonstrate frequency errors due to mismatches in the frequency detector.

2. Smaller V_{ctrl} results in smaller pull-up current I, and thus smaller Δf.

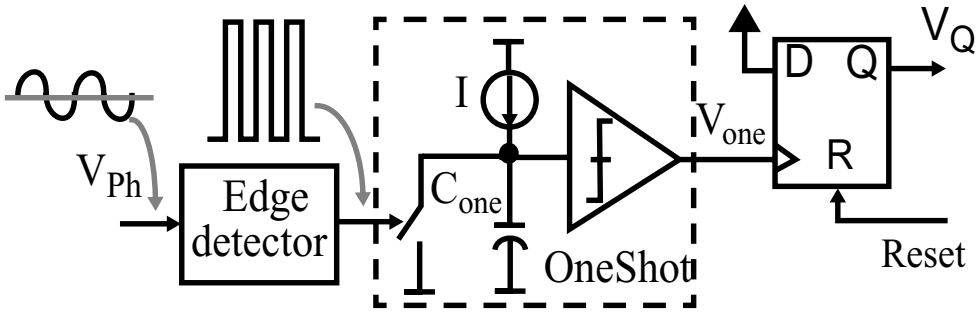

(a)

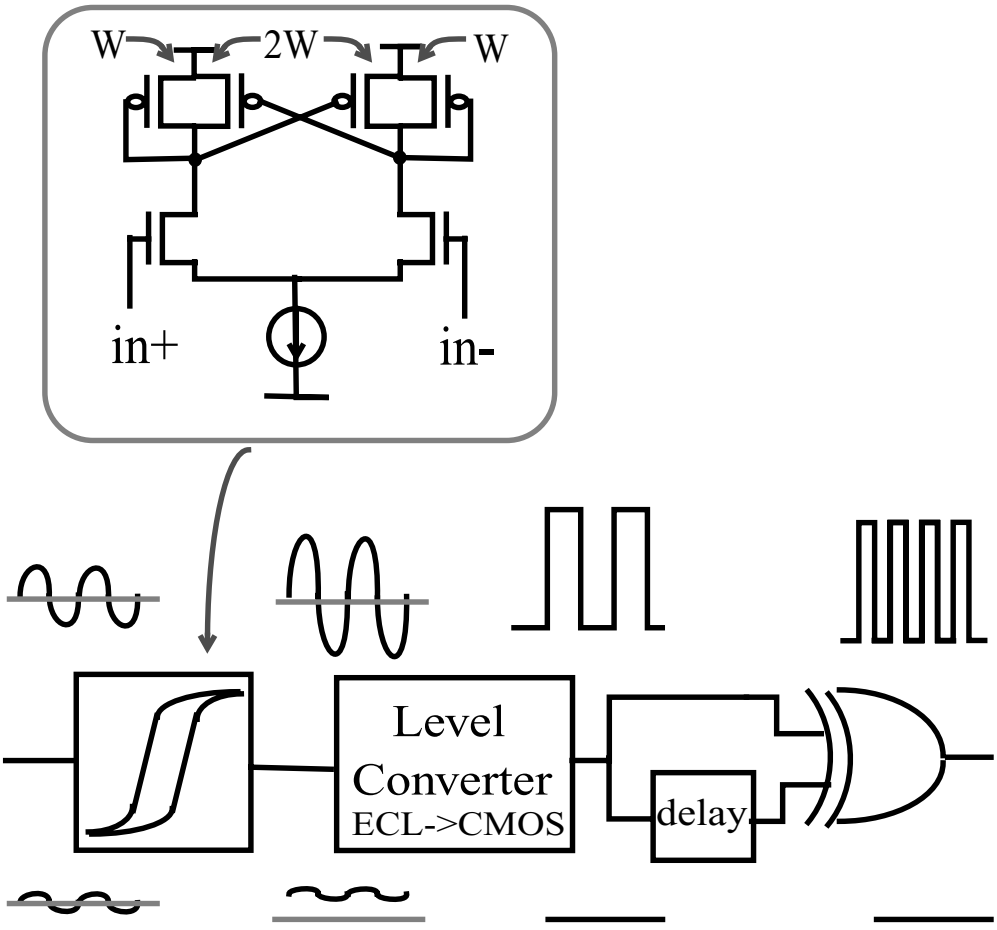

(b)

Figure 4.20: Frequency monitor: a) top view, b) edge detector

in its first stage amplifier. Thus, it reacts only to oscillation amplitudes larger than a certain threshold level, which helps prevent erroneous transitions due to noise.

4.3 Summary

This chapter addressed two issues: generating the multiple clock phases for the transmitter and receiver, and recovering timing information from the input stream to ensure a good timing relationship between the data and sampling clock. Both affect the timing margin and performance of the system.

Ten evenly spaced clock phases are generated by tapping from a 5-stage differential ring VCO for the 5:1 multiplexing, 1:5 demultiplexing, and timing recovery. To reduce the jitter in these clocks, the VCO delay elements are built using differential stages with linearized PMOS load resistors, and tail current sources. A replica bias circuit increases the output impedance of the tail current source to make the stage more insensitive to supply variations. Because of mismatch between the delay stages, there is always a phase error between adjacent phases. This phase error can be minimized by careful layout and tuning of the stage elements for the target frequency range. Using a single-ended charge pump results in control voltage ripple that degrades the phase spacing further. To minimize the ripple, an improved version of the charge pump is used in addition to a third-order capacitance.

The clock recovery takes advantage of a new proportional sampling phase detector that recovers timing from a very fast input stream. This phase detector combines the advantages of traditional bang-bang detectors (integrated with demultiplexer to operate on high-speed data) and proportional ones (good loop bandwidth and stability). In addition, control voltage ripple is eliminated in this approach, at least in principle. One of the main problems with this approach is processing analog edge samples through multiple analog stages, resulting in phase error in recovered clocks due to voltage offset in those stages.

The recovery loop uses a frequency acquisition aid to lock the VCO frequency to a reference clock that is just slightly away from the target frequency, guaranteeing that data

phase recovery does not fail due to the limited frequency capture range of the phase detector. The frequency acquisition aid has a frequency monitor circuit that uses cycle slipping information to measure the frequency difference between the VCO and input stream, and thus determine when loop control should be transferred to the main loop.

Chapter 5

Measurement Results

This section describes measurements performed on the high-speed link including the transmission medium, package, and the transceiver chip. The target transmission medium for this project is a 10-meter twinax cable (or a pair of coax cables) plus the board and package traces and connectors. Channel characterization has to be performed in advance to provide the required information for receiver and transmitter filter design. This work uses a new approach to model high-frequency loss (e.g. skin effect loss) in the the transmission media more accurately as is discussed in Section 5.1.

Section 5.2 describes the board design and chip packaging for the transceiver test setup, and techniques to improve the high-speed signal integrity. Finally, the experimental results from the two test chips in 0.4-μm and 0.3-μm CMOS technology from LSI Logic are discussed in Section 5.3.

5.1 Cable Modeling

One of the main limiting characteristics of the copper cables at high frequencies is the skin effect resistance that causes frequency-dependent attenuation, as described in Chapter 2. This work uses the cable impulse response in time domain that is obtained by time domain transmissometry (TDT) measurement in order to characterize channel transient response. The PCB traces are not included in these measurements, because in the test setup of this work, the PCB traces (~1inch) are considerably shorter than the cable (10 meter) used to introduce any low-pass filtering in the channel.

Assuming that an ideal TDT device, which generates an ideal unit step signal, is used to excite the transmission line, the output waveform of the line is the channel step response. The channel impulse response can be obtained by taking the derivative of the step response. Having the cable impulse response, one can calculate the transient response of the channel to any input waveform by convolving the input waveform with the impulse response in the time domain. These steps are shown in Figure 5.1. However, the actual TDT device used in the lab showed a finite rise time of ~50ps, thus resulting in an inaccurate step response. Hence, a simple digital signal processing step is used to solve the limited slew rate problem of input step. Since the TDT device's rising transition shows a constant slope, the derivative of TDT step waveform is an ideal 50-ps wide pulse (Figure 5.2a). One can form an ideal step waveform by normalizing the amplitude of these 50-ps pulses and putting an infinite number of them together with 50-ps spacings, as shown in Figure 5.2a. Considering the cable as a linear time invariant (LTI) system, one can perform

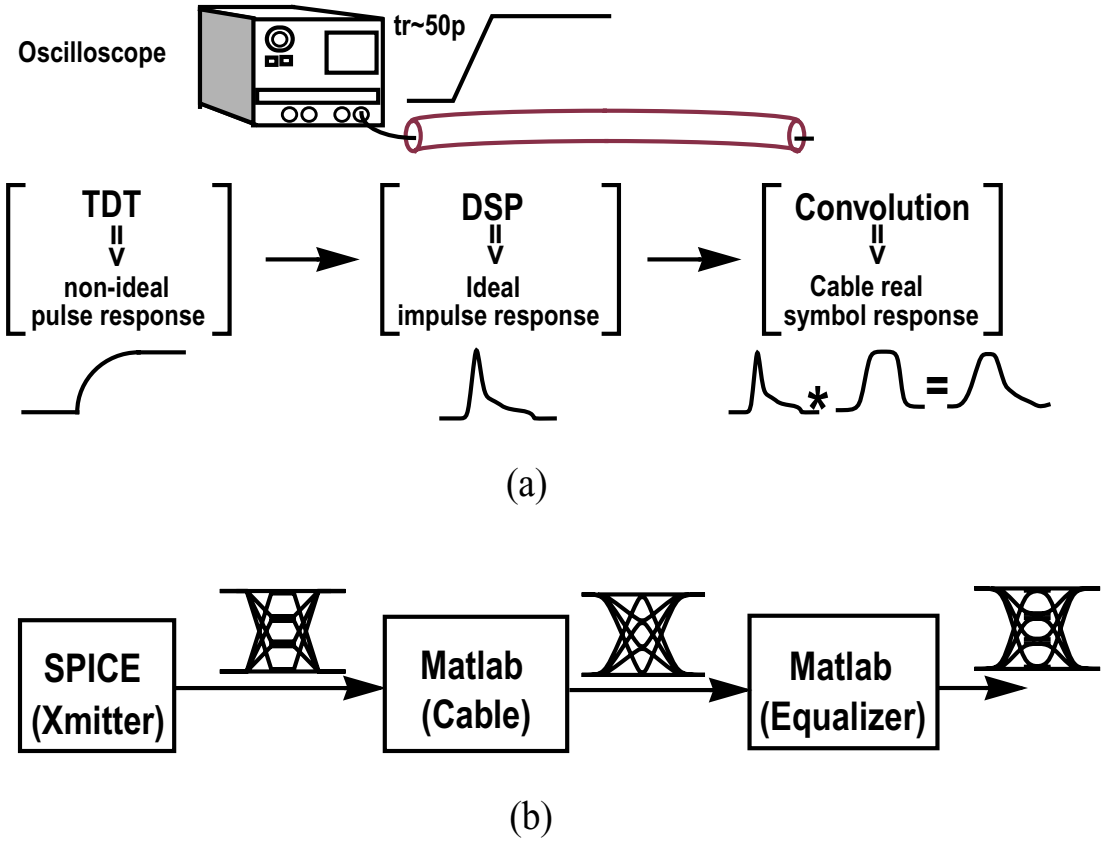

Figure 5.1: a) Steps to obtain the cable impulse response to an input waveform, b) simulation steps to obtain the transmitted eye diagram at the receiver

the same steps, in Figure 5.2a, on the cable TDT step response to obtain the cable response to an ideal input step, as shown in Figure 5.2b.

Figure 5.3 shows the impulse response of a 10-meter PE142LL coaxial cable, used for link measurement purposes, for two cases of the impulse response. The wider impulse response is simply the derivative of the cable response to a non-ideal TDT step signal, and the narrower impulse response is obtained by the steps illustrated in Figure 5.2.

Figure 5.4 shows the frequency response of this cable that is obtained from the Fourier transform of the cable narrow impulse response in Figure 5.3. The -3dB frequency

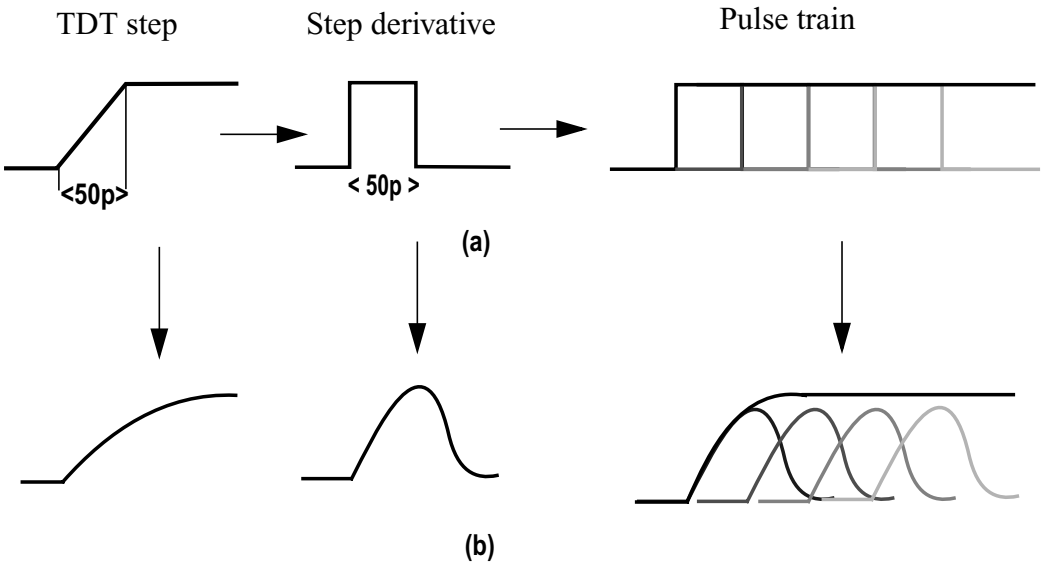

Figure 5.2: a) Steps to construct an ideal step from the slew-limited step signal generated by TDT device, b) cable response to each input signal

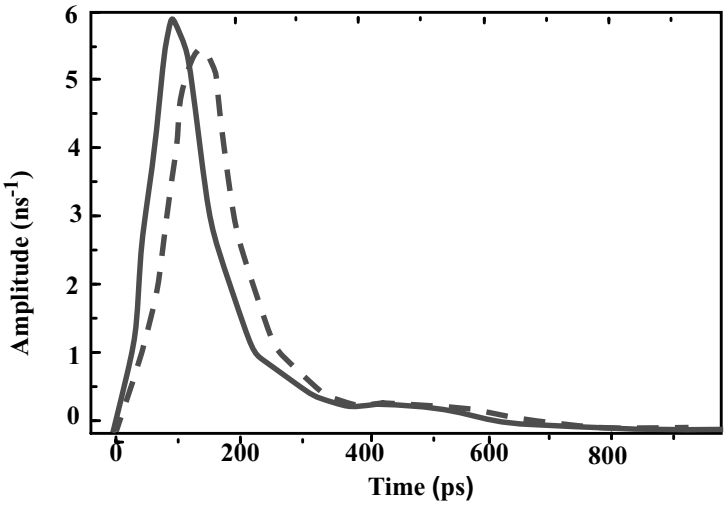

Figure 5.3: The 10-m cable impulse response from non-ideal TDT step (dotted) and an ideal step signal (solid).

of the cable is measured to be ~900MHz, which is in good agreement with the measurements from a network analyzer.

The impulse response of Figure 5.3 is used to find the transmitted 4-PAM eye diagram after the 10-meter cable using MATLAB simulations. As it is not possible to plot

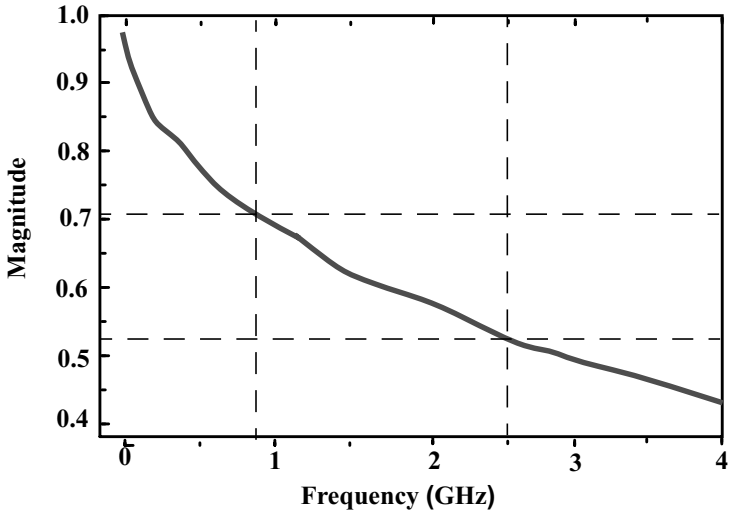

Figure 5.4: The 10-m cable frequency response

a continuous equalized waveform at the receiver due to the discrete-time FIR equalizer, MATLAB is used to simulate the effect of the receiver 1-tap filter on the eye diagram.

Figure 5.5 shows the HSPICE simulated 10-Gbps 4-PAM eye diagrams after the 10-meter PE142LL coaxial cable, which is the HSPICE generated waveform from the transmitter convolved with the channel impulse response. The first plot (Figure 5.5a) clearly indicates that without channel equalization, zero eye opening is achieved, while second plot in Figure 5.5b shows the improvement due to transmitter pre-emphasis. Figure 5.5c shows the eye diagrams after the cable with transmitter pre-emphasis and receiver equalization. Based on the simulation result, the receiver 1-tap equalizer mortifies (increases) the 4-PAM eye area (height x width) by 40% at 10Gbps. The measured results, however, show that the receiver equalizer is not as effective at very high speeds as suggested by simulations. Possible factors degrading the equalization performance are discussed in the following sections.

The actual test setup uses a pair of 10-meter PE142LL coaxial cables, instead of a single twinax cable, for the differential serial transmission. The main reason for this choice is the availability of high-quality radio-frequency SMA connectors for coaxial cables in the market, as opposed to twinax cable connectors that are hard to find or not even made for the frequency range of interest in this work[1]. However, the lengths of the two coaxial cables need to be carefully matched so that the skew between the two differential signals in the cable pair is not more than 10% of the minimum symbol period of 200ps (5Gsym/s). Therefore, the two 10-meter cables are cut and matched manually within 0.5cm, so that the time skew between them does not exceed 20ps (200ps x 10%). The actual differential TDT measurement on these two cables shows that their delay match is about 10ps.

5.2 Package and Board (Traces and Connectors)

To minimize the loss of PCB traces for high-speed serial data and achieve test results closest to the simulations with the measured cable impulse response (Section 5.1), high-quality synthetic board materials[2] are used. The trace width on the board is designed to maintain a 50-Ω environment and to prevent reflections from corrupting the signal. With a dielectric constant near 4, the trace width for 50Ω needs to be roughly 2x the dielectric thickness. A four-layer board is used for a thinner dielectric (8mils instead of 62.5mils) and hence better routing density. Mismatch between board trace impedance and

1. The best twinax connector found at the time of experiments was made by GORE for 1000BASE-C Gigabit Ethernet (1.25Gbps).
2. RO4003 supplied by Rogers Corp.

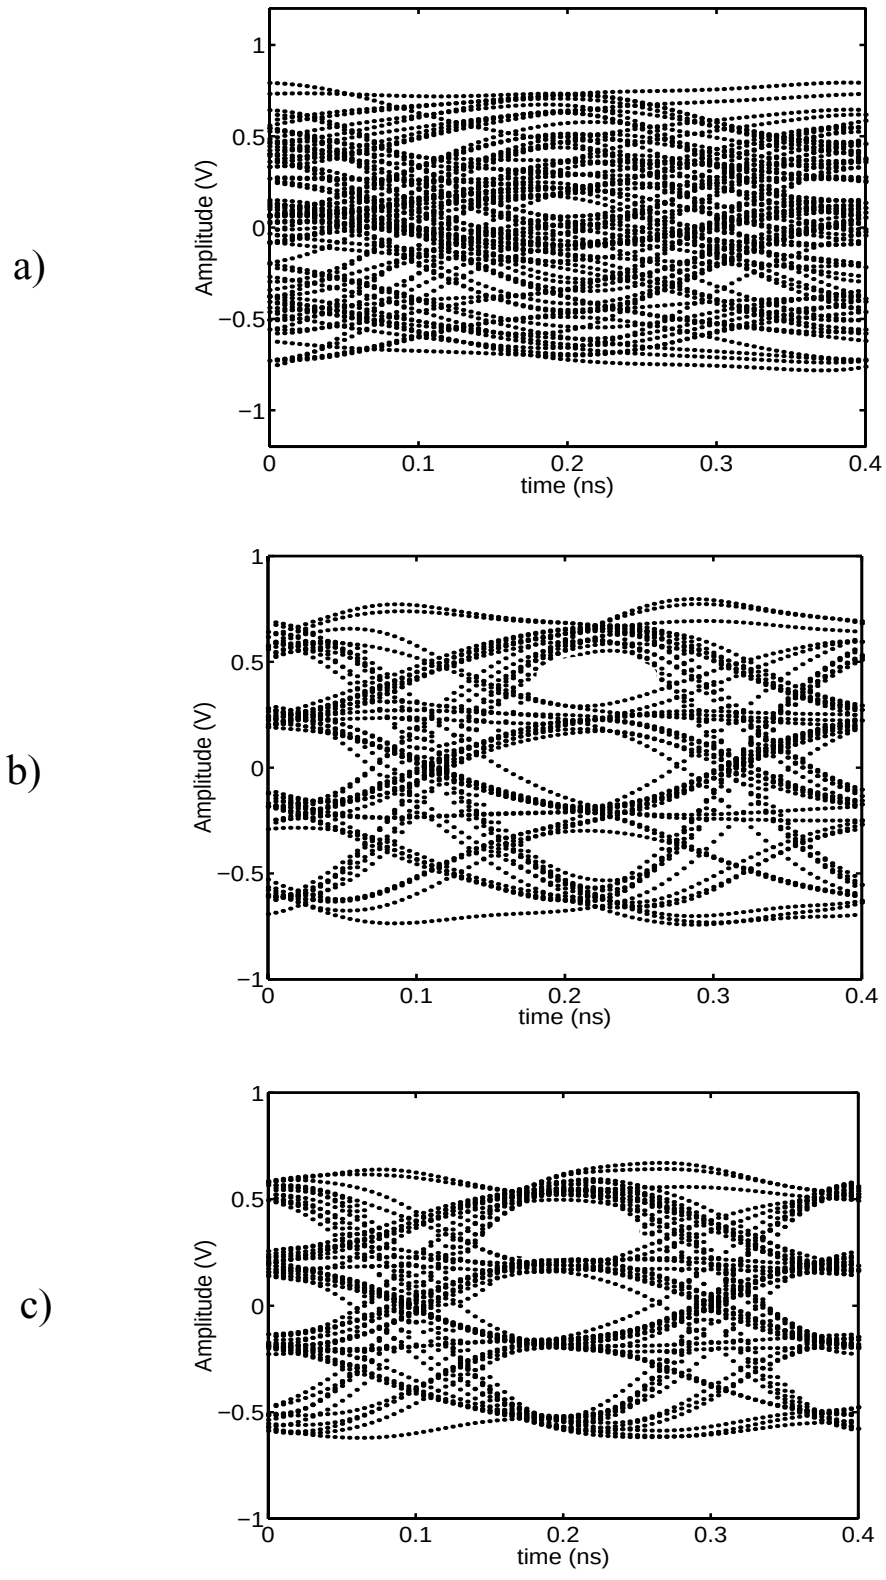

Figure 5.5: 10-Gbps 4-PAM eye diagram after the 10-m cable a) with no pre-emphasis, b) with transmitter pre-emphasis, c) with receiver equalization.

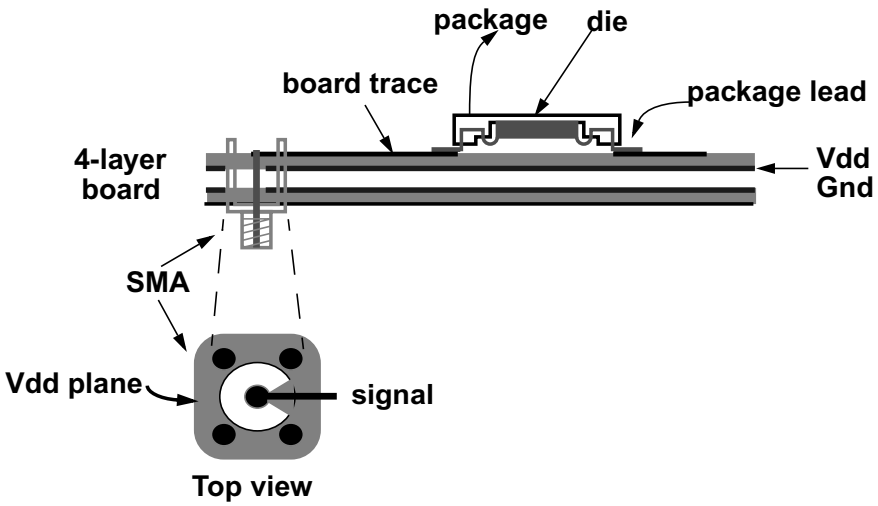

Figure 5.6: High-speed serial signal path from die package to SMA connector

the coaxial cable can cause reflections. Fortunately, the impedance can be well controlled by the PCB fabrication process to result in less than 1-% impedance mismatch.

To ensure good signal integrity for the high-speed serial traces on the board, the design shown in Figure 5.6 is used. To minimize the parasitic effects of the package leads, they are cut short enough to be soldered to the 50-Ω trace. In addition, to reduce the signal path discontinuity at the connection point of the board trace and SMA connector, the SMA connector is installed at the back of the board so that its center conductor end is closest to the board trace. Mounting the SMA connector on the other side causes the center conductor acts as stub (discontinuity) in the signal path. The center conductor is cut flush with the traces to eliminate any additional discontinuity. The design of the SMA launch is shown in Figure 5.6. The internal planes are cleared near the center conductor to avoid perturbations in the characteristic impedance arising from capacitive coupling and loading.

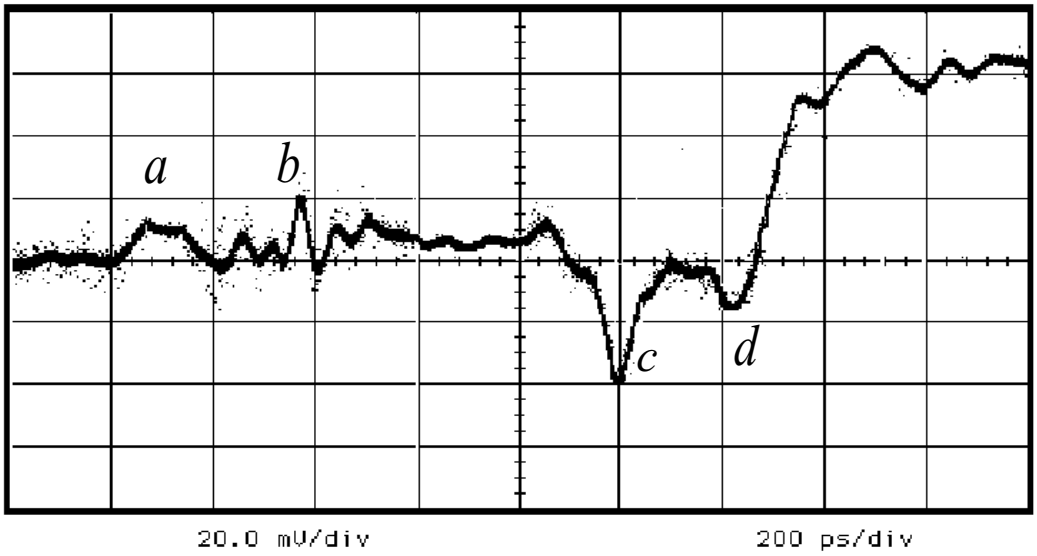

Figure 5.7: Differential TDR of the serial receiver input path.

Figure 5.7 shows the time domain reflectometry (TDR) plot of the differential high-speed path shown in Figure 5.6, for the receiver side of the transceiver chip, where the amplitude of TDR step is 1.0V. The peak in point *a* is the connection section between the SMA connector and the cable that has a slightly larger impedance than the cable and trace. The SMA connection to the board is only evident as a small ripple at point *b*, due to the design shown in Figure 5.6. The remaining signal between points *b* and *c* is flat and corresponds to the controlled impedance on the board. The big dip at point *c* is due to the package pin capacitance, estimated around 0.4pF from the area of the dip. The relatively constant impedance section between *c* and *d* is the controlled impedance trace of the package. The small dip at point *d* is due to the bondwires of the receiver inputs that introduce a small capacitive effect. The small discontinuity on the bondwires is due to the parallel differential structure that implements a transmission line as described in Chapter

2. Finally, the signal step after point *d* (bondwires) is a result of the unmatched on-chip termination resistors, corresponding to ~60Ω. The reason for this high termination impedance at receiver inputs is explained in the next section.

On-chip termination is clearly preferable to off-chip termination with discrete resistors. Off-chip termination always results in an unterminated stub composed of the package parasitics and the internal circuitry, and this stub generally introduces unacceptable reflection back into the signal line [12]. In this design, to mitigate the effect of the undesired signal reflections, the high-speed serial line is doubly terminated at both transmitter and receiver side. In addition, the on-chip resistor helps damp any ringing at the die pads, caused by any parasitic LC due to package and bondwires.

To reduce supply noise, the board's power supplies are tightly decoupled with capacitors. Each supply is bypassed with a network of chip capacitors ranging from small values to large values. Because the small values have lower series inductance, a wide range of noise frequencies is bypassed by using a range of these capacitors. As the mutual impedance between the on-board supplies is very small and the chip does not have much digital circuitry to generate considerable digital supply noise, the digital and analog power planes are shared on the board. However, the two supplies are still kept separate on chip due to the larger impedance of supply connections from the die to the board, including bondwires and package traces.

The package used in this design is a 52-pin ceramic quad-flat package (CQFP)[1]. This package has controlled impedance traces from the pins to the bonding pads using

1. Supplied by Vitesse Semiconductor Inc. and typically used for their 2.5-Gb/s SONET transceivers.

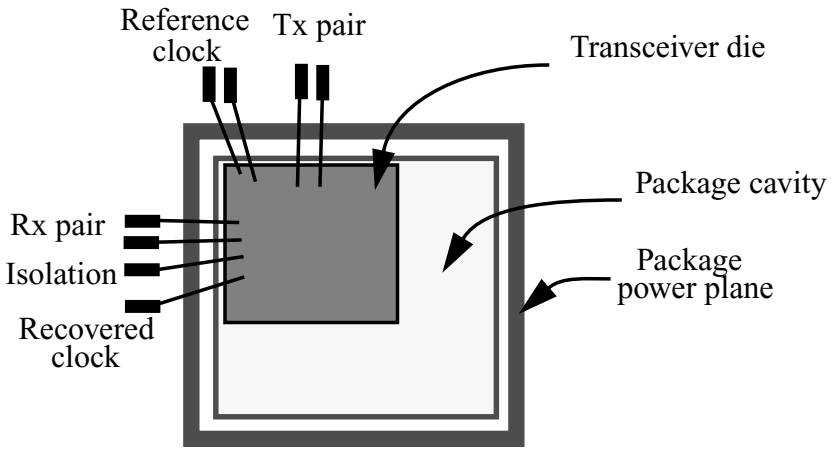

Figure 5.8: Bondwire configuration for the high-speed signal

internal supply planes to eliminate the package inductance, as verified by Figure 5.6. However, the impedance mismatch due to long bondwires can still disturb the signal, even using the parallel differential bondwire technique described in Chapter 3. Bondwire lengths are kept below roughly 2mm by placing the die as close to the package edge as possible. In order for both the transmitter and the receiver to be close to the package edge, the two blocks are placed on adjacent edges instead of opposing edges, as shown in Figure 5.8. Having both the transmitter and receiver on one side was avoided, as it results in electromagnetic interference from the high-frequency transmitter bondwires coupled to that of the sensitive receiver inputs. The latter effect is also called near end crosstalk (NEXT) that can cause an error even with differential signals because the coupling may be stronger to one of the two inputs. The 500-MHz reference clock is also brought to the die on the transmitter side to avoid coupling to the inputs bondwires.

The recovered clock from the incoming serial data is bonded out of the chip to observe receiver phase locking (Figure 5.8). Due to some last minute changes, the pad for

this output clock was unfortunately placed on the receiver side, and caused some performance degradation in the receiver measurements even with isolation bondwires/ pads between the recovered clock and the receiver inputs.

5.3 Silicon Test Results

The 4-PAM serial pre-emphasis transmitter and the receiver together with the proposed linear clock recovery architecture were implemented in silicon. Initially, the 4-PAM pre-emphasis transmitter was implemented to verify the simulation results for successful 4-level data transmission at 10Gbps both in the actual silicon and over the 10-meter copper cable. The results from this transmitter test chip were crucial in proving the feasibility of this signaling method. The second test chip was the full transceiver, and included an improved version of first transmitter chip.

Figure 5.9 shows the complete top level transceiver block diagram showing the main chip I/O pads, where each block is described in detail in Chapter 3 and 4. Both the transmitter and receiver have a fixed 20-bit data register for verifying and debugging the functionality of the transceiver by selecting a user programmable sequence and monitoring the transmitted waveform and recovered data at receiver. To facilitate the BER measurements, a 2^7-1 encoder on the transmit side and its matching decoder on the receive side are implemented on chip. The PRBS sequence generated in the transmitter is 4/5sym encoded for clock recovery to test the full link with random data similar to real data. In addition, to measure the clock generation sensitivity to voltage noise on the analog supplies, a large NMOS device is connected between the supplies that pulls the two on-

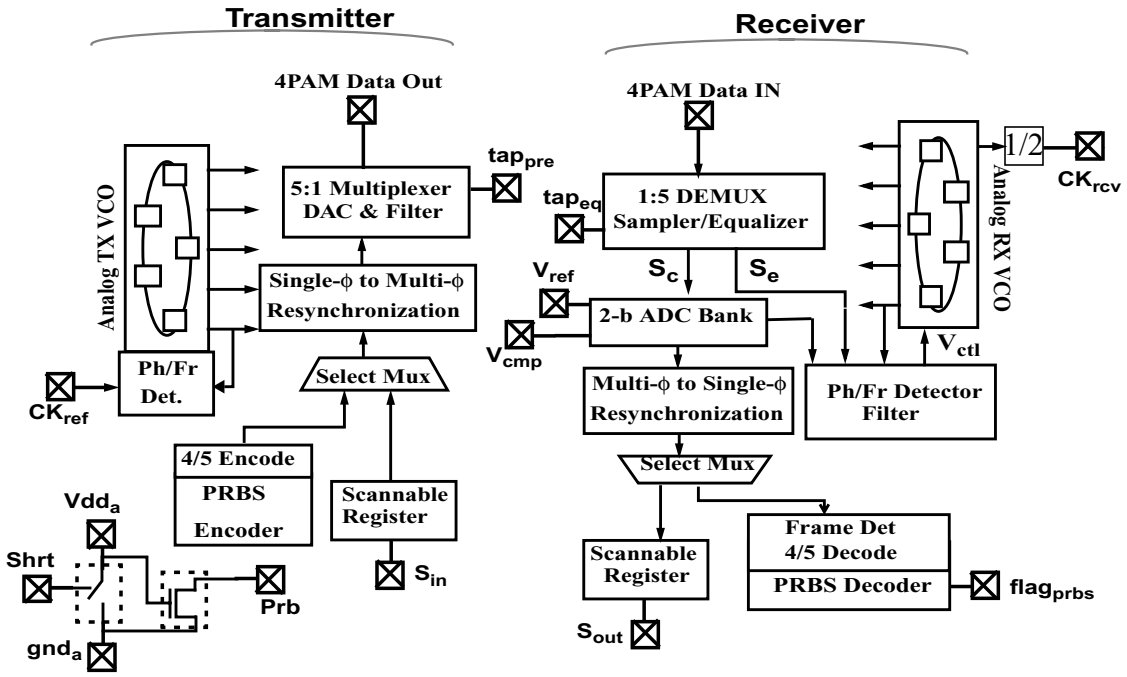

Figure 5.9: Complete transceiver block diagram

chip analog supplies together when the off-chip signal *Shrt* is turned high. In order to probe the supply variation on chip, an NMOS open-drain amplifier device is tied to supply pins as shown in Figure 5.9. A low off-chip resistor (~100Ω) is placed at the NMOS amplifier output (*Prb* pad) to convert the current into voltage. This simple amplifier design was chosen due to its high bandwidth, so that the analog supply noise at large frequencies is not filtered out.

The following two sections show the measurement results from the two test chips. The measurements from the first test chip, the 4-PAM 0.4-μm transmitter, are illustrated in Section 5.3.1, followed by the results of the second test chip showing the performance of the complete 0.3-μm transceiver with the improved transmitter circuitry.

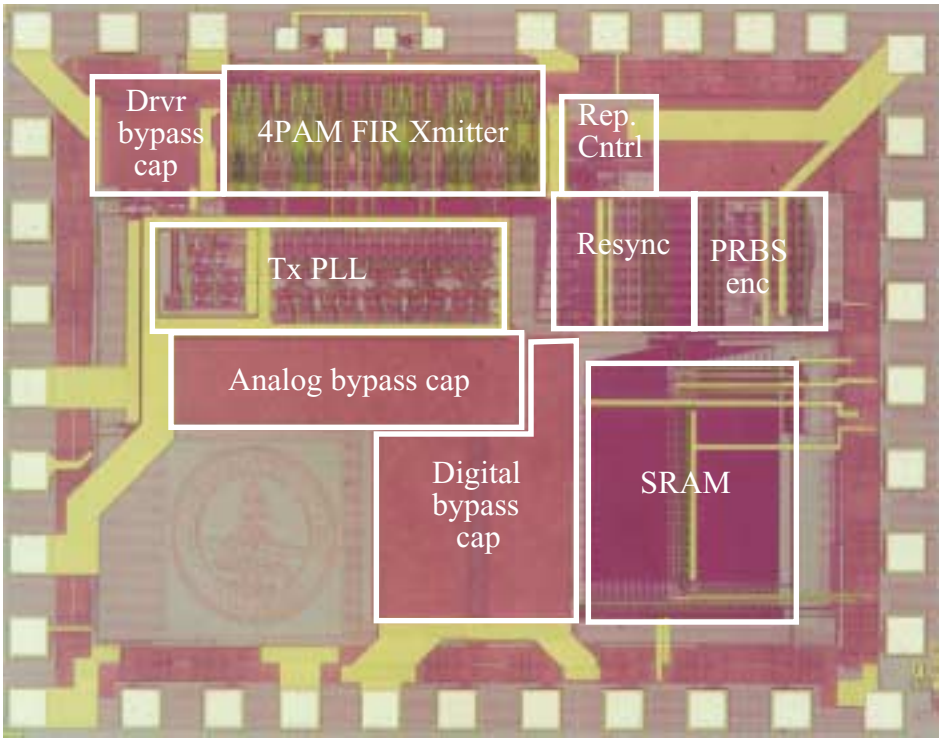

Figure 5.10: Transmitter die micrograph

5.3.1 Transmitter test chip

The transmitter chip was implemented in a 0.4-µm CMOS process technology offered through LSI Logic. The die photo of the 2mm x 1.5mm transmitter chip is shown in Figure 5.10. The 4-PAM pre-emphasis driver occupies an area of 0.8mm x 0.3mm. Three independent supply inputs are used on the chip for the analog clock generation block, digital functional blocks, and the 4-PAM pre-emphasis driver, each with separate on-chip bypass capacitors. The separate digital supply helps prevent noise of the digital section from entering other stages and affecting the output jitter. In addition, a separate supply for the 4-PAM driver helps isolate the sensitive analog stages in the PLL from the

supply noise due to driver large stages buffering the random data. The individual supply for the 4-PAM driver also allows adjusting the amplitude of the output swing.

The size of the output pads is reduced to 50µm x 50µm to keep pad capacitance to a minimum. To guarantee less than 10% amplitude loss due to output RC filtering, the total output capacitance at the 25-Ω I/O (for a doubly terminated 50-Ω line) should not exceed 3.6pF for a 5Gsym/s signal. The 5:1 multiplexing transmitter has a total capacitance of 1.4pF, 400fF of which is due to the pad and metal interconnects.

The transmitter has an output jitter of 19ps (p-p) and 3ps (rms) for a fixed 10-symbol long (20-bit long) sequence. However, the jitter increases to 32ps (p-p) and 8ps (rms) for the long PRBS sequence, due to the random supply noise generated by the large 4-PAM driver stages. Note that the high ratio of the rms jitter to the peak-peak jitter (32/8=4) in the latter case indicates that the jitter distribution is not Gaussian and is affected by another noise source. Mismatches in the VCO stages and output driver paths increase the total phase error, effectively forming a smaller data eye. All the clock buffers in the output multiplexing driver were designed with nominally equal fanout. However, due to underestimating the wire capacitances, one of the clock buffers experienced a different loading, causing misplacement in one of the output phases (transition edges) by ~40ps. The misplaced phase results in a wider eye opening (240ps) on one side and a narrower one (160ps) on the other side of that transition edge. This problem is fixed in the transmitter implementation of the next test chip by designing a completely symmetrical signal path for all the driver clock buffers, so that the remaining output phase errors are mainly due to device mismatches and fixed-pattern noise in the supply and control voltages of the PLL.

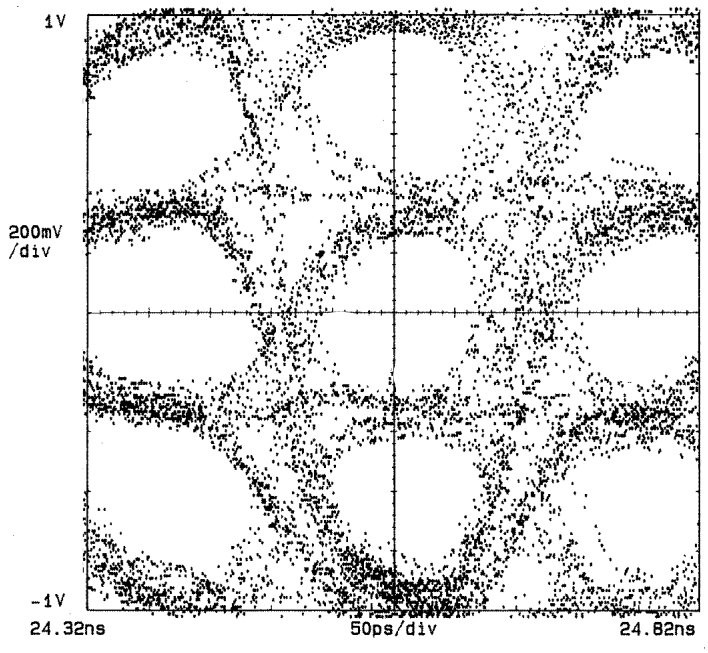

Figure 5.11: Differential data eye at 10Gbps over ~0.3-m wire

The transmitter achieves a symbol rate of 5Gsym/s (10Gb/s) with a differential eye-height of >390mV and eye-width of >110ps at a total differential swing of ~2.0V (1.0 single-ended) over board traces and 0.25 meters of cable, as shown in Figure 5.11. Note that the eye opening measurements exclude the effects of the unbalanced loading in the driver clock path, as explained earlier. After 10 meters of PE142LL coaxial cable the 4-PAM eye diagram collapses. However, by tuning the pre-emphasis filter taps a 4-PAM eye diagram with an eye-height of >200mV and eye-width of >90ps is achieved, shown in Figure 5.12. As symbols without pre-emphasis after the 10-m cable show a zero eye opening at 5Gsym/s, it is clear that ISI mitigation is essential.

The transmitter test chip has three controllable pre-emphasis tap weights that allow channel equalization for different cable types and lengths. Although the simulation results show that all three filter taps are effective in channel ISI reduction, the measurements

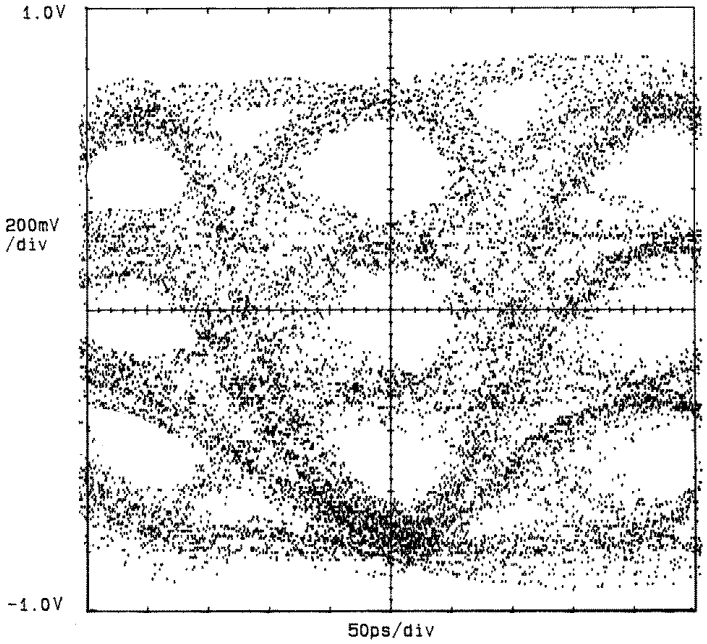

Figure 5.12: Differential data eye at 10Gbps over 10-m cable

performed on the actual transmitter chip and cable clarify that the contribution by the third filter tap is smaller than the driver output amplitude noise and improves the eye opening insignificantly. Therefore, only two taps are used by the pre-emphasis filter in the second transmitter implementation to reduce complexity and output loading.

The transmitter chip dissipates a total power of 1.45W at full speed operation (10Gbps) and maximum output amplitude of 1.2V single ended (2.4V differential). The analog section of the transmitter, consisting of the input clock amplifier, source-coupled multi-phase clock generation PLL, source coupled buffers, and low-swing to high-swing converters, consumer the largest share of power, at 0.65W. The transmitter 5:1 multiplexing driver, including the differential 4-PAM drivers, the pre-emphasis filter modules, and all the clock and data pre-buffers, dissipate a total of 0.50 Watts. The

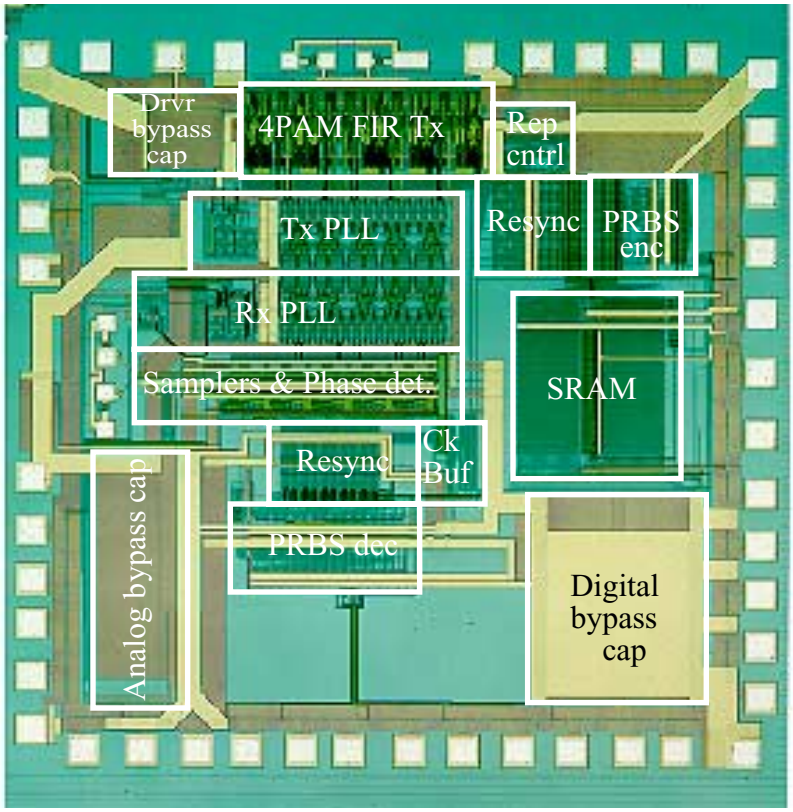

Figure 5.13: Full Transceiver chip micrograph

remaining 0.30W of power is consumed in digital functional blocks such as the resynchronization logic, PRBS encoder, and 4/5sym coder.

5.3.2 Transceiver test chip

The die photo of the 2mm x 2mm ($4mm^2$) transceiver chip is shown in Figure 5.13. The transceiver chip was implemented in a 0.30-µm CMOS process technology offered by LSI Logic, with a maximum suggested supply voltage of 3.0V. The data recovery phase detector and all the input samplers occupy a total area of 0.8mm x 0.15mm ($0.12mm^2$). The net active area of the full transceiver, consisting of transmitter and receiver clock generation PLLs, transmitter pre-emphasis driver, receiver data samplers, clock recovery,

and the resynchronization blocks, is less than 1.0mm^2. Larger bypass supply capacitors are used in this chip compared to the transmitter chip to suppress supply noise further.

As this process offers well-controlled non-salicided polysilicon interconnect, this layer is used to implement the 50-Ω on-chip termination resistors. These have no diffusion loading, unlike the large PMOS termination resistors used in the previous transmitter design. The poly resistor in this process has 10% variation over the process corners, therefore, small PMOS devices are used in parallel with the poly resistors to trim the total resistor value close to 50Ω. Unfortunately, the foundry omitted the salicide-block mask during fabrication, which caused all poly resistors to be very low resistance (~8Ω). Thus, all the termination poly resistors were cut by a YAG laser, and the trimming PMOS overdriven to act as the 50-Ω termination resistors. However, the lowest resistance value achieved by the PMOS termination devices is ~60Ω. This effect, which is also shown earlier in the TDR plot of Figure 5.7, degrades the signal integrity of the high-speed path.

The transmitter in this chip achieves a symbol rate of 5Gsym/s (10Gbps) with an eye opening of >200mV and >90ps, which is similar to the performance obtained by the previous transmitter chip, however at a lower power dissipation and jitter. The output driver, together with its two filter tap drivers and all the clock and data buffers, consumes a total power of 340mW at the full output amplitude of 1.2V and 5Gsym/s. The complete clock generation block consumes 530mW to generate the 10 clock phases at 1GHz, corresponding to 5Gsym/s. The output signal also shows better jitter performance of 11ps (peak-peak) and 2ps (rms) for a 10-symbol long data pattern, and 25ps (peak-peak) and 6ps (rms) for the PRBS sequence. The unbalanced data path in the driver of the original transmitter is fixed in this test chip, and thus no large phase offset is observed in the output

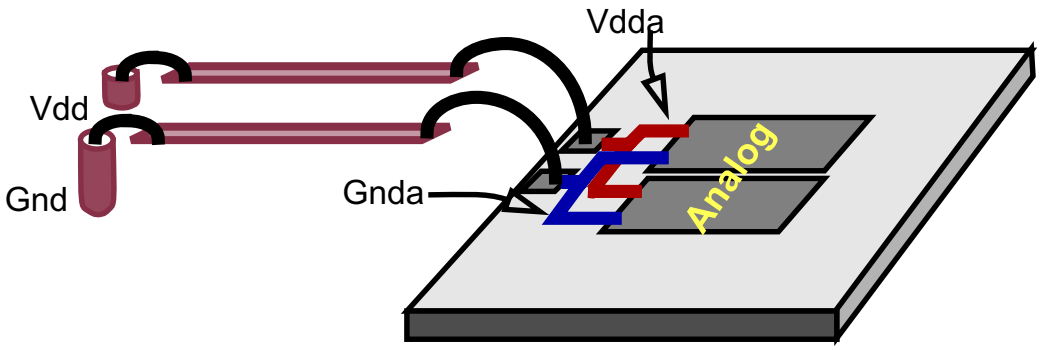

Figure 5.14: Only one pin per analog supply (Vdd & Gnd) is available in the transceiver test chip

eye diagram. At 5Gsym/s a maximum phase spacing error of 13ps (~7% of the 200-ps symbol width) was measured at transmitter output. Therefore, the 200-ps symbol period is degraded by a maximum of 38ps due to jitter and phase spacing errors.

This chip similar to the transmitter chip has three separate supply inputs on chip for the analog blocks, digital blocks, and the 4-PAM driver. The 52-pin CQFP package has only two separate supply planes that are allocated to the digital and 4-PAM driver supplies because these blocks consist of many noisy stages and require a lower impedance path to the off-chip supply. Therefore, the dedicated connections for the analog power and ground has to be made through the other package pins, and due to the pin limitation in the transceiver, the analog supply (Vdd and Gnd) pins were limited to only two (Figure 5.14). Unfortunately, the total series resistance adds up to ~1.7Ω causing a considerable drop for the on-chip analog supplies at large current consumption. For example, at 5Gsym/s (10Gbps) the total analog current draw is ~360mA, with corresponding voltage drop of 0.57V. As a result, for the transmitter to operate at 10Gbps, the off-chip supply has to be

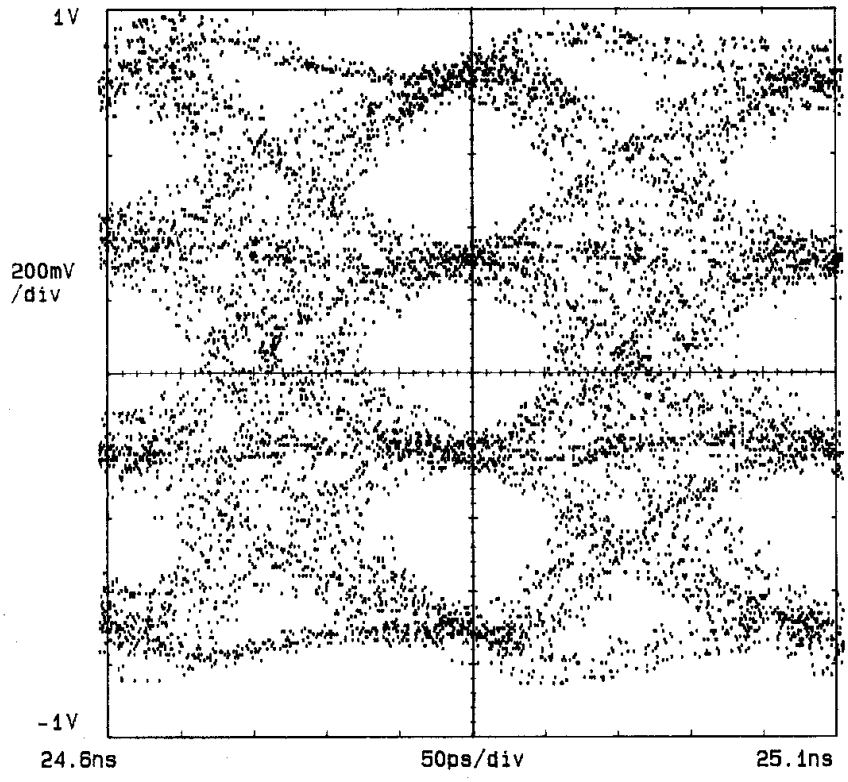

Figure 5.15: Differential data eye at 8Gbps over 10-m cable

increased to at least 3.4V. At 3.0-V supply, the transmitter operates at a maximum speed of 8.4Gbps. The clock recovery block is limited to a maximum of 8Gbps at 3.0-V, which limits the operation speed of the receiver, and thus the complete transceiver loop, to 8Gbps. Increasing the off-chip supply voltage to 3.3V helps the full link to run close to 9Gbps. However, an off-chip supply higher than 3.3V does not help increase the receiver speed, because the input serial stream is still referenced to the off-chip supply while the larger drop in the on-chip analog supply causes the input pre-amplifiers (Chapter 3) to fall out of their gain region. Hence, the measurements for the transceiver chip are mostly done with a supply of 3.0V at a maximum speed of 8Gbps.

Figure 5.15 depicts the differential eye diagram generated by the transceiver test chip at 8Gbps (4Gsym/s) with a maximum single-ended amplitude of 1.2V after the 10-meter cable. As the eye diagram shows, an eye opening of 350mV and 120ps is achieved over the cable with transmitter pre-emphasis, while without pre-emphasis this 8-Gbps eye opening reduces to 60mV and 50ps. To convert this 4-PAM signal into binary bits, the receiver reference voltage for the 2-bit ADCs should first be adjusted. The receiver has a built-in comparator that generates a ONE, at pad V_{cmp} (Figure 5.9), when all receiver ADC outputs toggle. Using this comparator, the user can apply a differential DC voltage equal to the reference voltage level of the 4-PAM data, obtained from the eye diagram, to the receiver inputs and vary the off-chip tap for the ADCs' reference voltage until the comparator output, V_{cmp}, turns high. Due to device mismatches in the ADC differential samplers, all the ADC outputs do not toggle at the same reference voltage.

The same receiver comparator is used to measure the maximum input referred offset of the ADC stages. In this test, the ADC reference voltage is set to zero and a DC voltage is applied to the receiver input that increases from zero differential value. The largest input differential voltage, in either the positive or negative direction, that causes all the ADCs to toggle shows the largest input offset voltage over all ADC samplers. Using this technique, a maximum input referred offset voltage of 50mV is measured over all the input samplers.

The BER measurements are performed using the test setup shown in Figure 5.16. The PRBS encoder in the transmitter generates an 8-Gbps pseudorandom sequence that is sent over the line. The receiver detects the serial signal from the line and, after proper framing, sends it to the PRBS decoder. Whenever there is a bit error in the received sequence, the

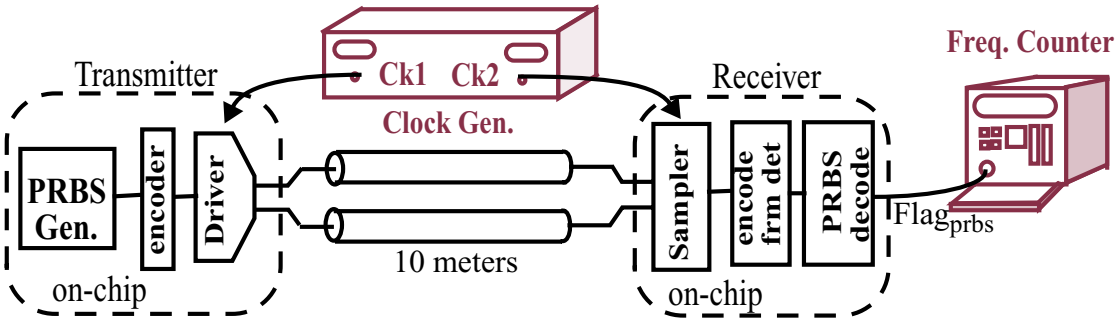

Figure 5.16: Test setup for BER measurements

PRBS decoder generates an error pulse. The number of these pulses per second is the system BER. The valid data window is measured by connecting the receive and transmit PLLs to two clock sources, as shown in Figure 5.16, and varying the delay of one clock source versus the other until a rapid increase occurs in BER. The receiver successfully detects an 8-Gbps 4-PAM data stream after 10 meters with a 3-V supply. At 8Gbps over 10 meters, the receiver had a BER of 10^{-7} for a time window of 50ps and minimum eye height of 250mV. At 6Gbps, the BER decreased to 10^{-15} for a window of 150ps and minimum eye height of 150mV.

Receiver equalization helps reduce the required transmitter pre-emphasis for the 10-m cable, effectively allowing the use of longer cables for the link. The receiver equalizer is adjusted manually, as is the transmitter pre-emphasis filter. However, as opposed to the transmitter output, the equalized waveform in the receiver cannot be viewed and used to set the optimized tap weight value. Therefore, the equalizer tap is adjusted to minimize the measured BER. At 6Gbps, the equalizer helps increase the 150-ps valid sampling window to 180ps, while it did not improve the valid window noticeably at 8Gbps.

The reasons for the considerable improvement in the link performance at 6Gbps compared to 8Gbps include lower on-chip analog voltage drop at 6Gbps, which results in better performance (e.g. larger gain) of the receiver front-end stages including the equalizer, and less spurious interference due to coupling from high-frequency adjacent bondwires (Figure 5.8), and smaller signal reflections from imperfect terminations (Figure 5.7).

The receiver data-recovery PLL requires that the input symbols have a minimum peak-to-peak swing of 800mV differential (400mV swing on each line) to acquire lock, and 600mV differential swing to maintain lock. This PLL has a capture range of >20MHz in the worst case, which is for a symbol stream with one transition per clock cycle (5 symbols). The frequency acquisition circuit switches the loop control to the data phase detector when there is less than 100kHz frequency difference between the transmitter and receiver reference clocks. The receive PLL has a jitter of 28ps (p-p) and 4ps (rms), when locked to the transmitted 20-bit long data sequence with a jitter of 11ps (p-p). The decision logic of the data phase detector injects undesired charge onto the VCO control line, causing error in the sampling clock phases and data detection. In case of a long random sequence, such as the 2^7-1 PRBS sequence, these random charge injections cause the clock recovery loop to lose lock. The source of this problem is detected and fixed as explained earlier in Chapter 4 (Figure 4.17), and proved to be functioning correctly by simulation.

Since the receiver clock recovery loop uses a linear phase detector that operates on the analog edge sampled values, any voltage offset in the phase detector front-end samplers and preamplifiers translates into phase error in the recovered clock. Assuming that the phase detector samplers and preamplifiers have a maximum voltage offset close to

that of the ADCs samplers and amplifiers, one should also expect a voltage offset of 50mV for the phase detector. At 8Gbps (4Gsym/s), considering the worst case maximum transition time of 250ps and a typical eye height of 300mV, the 50mV input offset voltage translates into a maximum phase offset of

$$\Delta\phi_{err} = 250ps \cdot \frac{50mV}{300mV} = 41ps. \quad [5.1]$$

However, the actual phase error in the phase detector is expected to be less than 41ps, since the phase detector uses the summed outputs of five of these sampler and amplifier stages that statistically result in a smaller final offset voltage.

Table 1 summarizes the performance of the transceiver test chip.

Table 1 Performance Summary

Transmitter performance	
Maximum transmitter rate	10Gb/s @3.3V, 8Gbps @ 3V
Output jitter @ 8Gbps	11ps (p-p), 2ps (rms) (short length)
Output jitter @ 8Gbps	25ps (p-p), 6ps (rms) (PRBS)
Supply sensitivity	0.35ps/mV
Max. eye opening @8Gbps	350mV, 110ps (10-m cable)
Max. eye opening @10Gbps	200mV, 90ps (10-m cable)
Receiver performance	
Maximum receive rate	9Gbps @3.3V, 8Gbps @3V
Data PLL jitter @ 8Gbps	28ps (p-p), 4ps (rms)
Data PLL capture range	>20MHz
Min. swing to capture lock	±400 mV
Min. swing to maintain lock	±300 mV
Data PLL dynamics	BW ~35MHz, Ph.m. ~50°

Table 1 Performance Summary

Power dissipation @ 8Gbps, 3V	
Output driver	220mW
Analog (2 PLLs)	750mW
Input samplers and logic	130mW
Total	1100mW

5.4 Summary

Test results show that equalizing the channel is crucial in transmitting a 4-PAM 5Gsym/s signal over the 10-meter coaxial cable pair with an acceptable eye opening. In addition, the techniques to improve the signal integrity of the high-speed path, such as a parallel differential bondwire structure, are verified to be effective.

The link speed is limited to 8Gbps at 3.0V mainly due to the large on-chip analog supply drop at high frequencies that corrupts the operation of the receiver input amplifiers and equalizer. The link performance is further degraded because of the high on-chip termination resistor (~60) and EMI interference from high-speed bondwires close to the receiver sensitive inputs.

The experiments proved that the proposed linear clock recovery architecture operated successfully up to 4.5Gsym/s (9Gbps) at 3.3V with a minimum frequency capture range of 20MHz, while it requires considerably lower power and area compared to other clock recovery architectures running at this data rate. The performance of this new clock recovery design is sensative to voltage offset in the phase detector stages that shift the recovered sampling clocks by a fixed phase error. Fortunately, using the techniques discussed earlier in Chapter 2, voltage offset of these differential stages can be digitally cancelled on chip to improve the clock recovery performance.

Chapter 6

Conclusion

The major factors that limit the performance of high-speed links are identified to be on-chip clock frequency, timing accuracy, and most importantly channel bandwidth. This dissertation showed that the performance of a multi-Gbps transmission link can be significantly improved by taking advantage of subtle circuit techniques together with simple communication techniques that can be implemented in modest CMOS technologies (e.g. 0.4-μm and 0.3-μm CMOS).

The simplest approach to achieve higher bit rates at lower on-chip clock frequencies is to use multiplexing and demultiplexing at the transmitter and receiver respectively. The demultiplexing at the receiver is performed by interleaving multiple samplers. The maximum bandwidth of the sampler is one of the primary limitations on receiver performance. This bandwidth is limited by the sampling aperture and the sampling uncertainty of the sampling switch. The multiplexing at the transmitter is also performed by interleaving parallel drivers using overlapping clock phases. One of the limitations on the transmitter bandwidth is the maximum rate that each driver can switch

the full output current. The large capacitance of the high fan-in receiver and high fan-out transmitter also can limit the maximum operating signal frequency. However, the multiplexing and demultiplexing can be performed at the chip I/O pads and merged together with the ESD structures, so that the low off-chip impedance (typically 50Ω-25Ω) provides a small RC time constant, enabling higher off-chip rates. To track technology scaling, additional complexity is required to handle the increased device mismatch and on-chip noise levels. Although jitter is expected to scale with technology, more attention must be paid to the clocking architecture, in order to minimize phase noise and phase offset of the on-chip clock and maintain high bit rates.

A 4-PAM modulation scheme reduces the signal bandwidth by a factor of two (doubles the symbol period), relaxing the on-chip clock frequency by the same amount. Therefore, the maximum on-chip data rate in a certain technology is simply doubled by moving from traditional binary 2-PAM to 4-PAM modulation. The fundamental limitation of data rate scaling stems from the transmission channel, rather than the device process technology. Frequency-dependent channel attenuation limits the bandwidth by introducing intersymbol interference. The lower symbol rates of the multi-level method relaxes multi-Gbps data transmission over the low-pass copper cables used as target transmission medium in this work, by concentrating the signal power spectrum in the lower frequency regions of the low-pass channel.

Although multilevel data transmission does alleviate the channel low-pass effects, mainly ISI, it cannot cancel them completely. Thus, line equalization must be used. As long as sufficient signal power exits, the simplest technique to cancel the line ISI and extend the data rate beyond the cable bandwidth is transmitter pre-emphasis. This work

proposes a method to implement a multi-tap FIR transmit filter at a very low-complexity, and therefore low power and area overhead. Similar but slightly more complex is the receive side equalization, which requires a much smaller amount of power compared to transmit pre-emphasis. This work also introduces a new approach to implement a FIR equalizer at the samplers' input by exploiting the parallel architecture of the high-speed receiver. The two transmit and receive filters compensate for the low-pass channel response and deliver an ISI-free signal over the 10-meter cable.

Using the above techniques, this work achieved a data rate of over 8Gbps on a pair of 10-meter copper cables with a -3dB bandwidth of <1GHz in 0.3-μm CMOS technology.

Future work

Since the scaling of CMOS provides such high bandwidths, while the off-chip electrical interconnects pose the main speed bottlenecks, one possible application is toward optical interconnects where the bandwidth of the medium is extremely high. For example, the 10-Gbps backplanes of future OC192 routers are planned to be optical. Applying techniques such as multi-level pulse modulation and signal equalization to the optical medium with linear optical drivers and receivers, one can increase data transmission rate-distance product further. However, electronic components to drive and receive the optical signals have more stringent criteria than electrical links. The receivers require resolution on the order of several millivolts and the drivers often require driving voltages greater than 3V. The design of these circuits while scaling continues can potentially introduce a difficult challenge.

One interesting area of circuit design is dealing with the transistor mismatches. Because transistor mismatches are static, offset-correction schemes, such as an open-loop correction at the start-up period of the circuit, can be applied. Input-offset voltage can be corrected using techniques commonly applied in ADCs [94]. However, the challenge is to perform the offset corrections with very little performance overhead for the receiver and at very short cycle times if performed continuously, which makes dynamic, closed-loop techniques very difficult to apply. Therefore, the work done in [33] and [95] performs receiver offset cancellation by measuring and saving the input offset value digitally at the circuit startup and using a DAC to subtract the digital to analog converted value, which represents the offset, from the input signal. The accuracy of multiple-phase clock generation is also degraded by device mismatches that result in delay mismatches in the oscillator and buffer stages. Such static phase offsets can be adjusted by programmable interpolators using digital control [83], or by adjustable delay stages using analog servo controls. If these techniques can be applied, the scaling of on-chip circuit performance can be expected to continue.

Based on Shannon's theory [2], the limit on the data rate of information that can be transmitted per Hz of channel depends entirely on the signal to noise ratio. In applications where the transmission medium is an impedance-terminated copper cable, the main source of random noise that can be eliminated or reduced by careful design techniques (e.g. supply, crosstalk) is the thermal noise of the terminating resistors (50Ω). Even for channel bandwidths up to 20GHz, this noise is very small compared to the typical signal amplitudes used (>100mV). Therefore, the maximum channel capacity in these cables can be quite high if optimum communications schemes are used. As CMOS process

technology improves, implementing more complex communication techniques to utilize the optimum channel capacity becomes more practical. For example a method that has been adopted recently by hard disk manufacturers is partial response maximum likelihood (PRML) detection [6]. The PRML method uses basis waveforms that extend a couple of bit times, and performs sequence detection using the Viterbi algorithm [7] at the receiver. The water-filling technique, as used in digital subscriber lines (DSL) [9], also increases channel capacity by allocating the better-behaved channel frequency bands to transmit a higher data throughput [3]. Should CMOS technology scale to perform such complex digital signal processes fast enough for multi-Gbps communication, these techniques will solve the problem of long low-pass electrical interconnects for a while, obviating the need for expensive optical interconnects.

Bibliography

[1] Jan W. Bergmans, *Digital Baseband Transmission and Recording*. Kluwer Academic publishers, 1996.

[2] C. E. Shannon, “Communication in the presence of noise,” *Proceedings of IRE*, vol. 37, pp. 10-21 , January 1949.

[3] E. A. Lee and D. G. Messerschmitt, *Digital Communication*. Kluwer Academic publishers, 1994.

[4] Bruce A. Carlson, *Communication systems*. McGraw-Hill, 1968.

[5] J. G. Proakis, *Digital Communications*. McGraw-Hill publications, 1995.

[6] D. Fisher *et al*., “PRML detection boosts hard-disk drive capacity,” *IEEE Spectrum*, pp. 157-169, November 1996.

[7] G. D. Forney, “The Viterbi algorithm,” *Proceedings of the IEEE*, vol. 61, no. 3, Sep. 1973.

[8] Roger E. Ziemer and Roger Peterson , *Introducotion to Digital Communication*. Prentice Hall Publicaions, 1996.

[9] Walter Chen, *DSL: Simulation Techniques and Standards* . Macmillan Technical Publishing, 1998.

[10] C.K.K. Yang *et al*., "A 0.6-µm CMOS 4-Gbps Transceiver with Data Recovery using Oversampling," *IEEE Symposium on VLSI Circuits*, Digest of Technical Papers, pp. 71-72, June 1997.

[11] John Poulton *et al.*, "A Tracking Clock Recovery Receiver for 4Gb/s Signalling," *Hot Interconnects Symposium*, Digest of Technical Papers, pp. 157-169, August 1997, Stanford.

[12] William J. Dally and John W. Poulton, *Digital System Engineering*. Cambridge University Press, 1998.

[13] S. Kim, B. Lee, and D. Neikirk, "Compact Equivalent Circuit Models for the Skin Effect," *http://weewave.mer.uteaxs.edu/*, Univ. Texas at Austin 1996.

[14] P. J. Black and T. Meng, "A 1-Gbps, four-state, sliding block viterbi decoder," *IEEE JSSC*, vol. 32, no. 6, 1994.

[15] W. J. Dally and John Poulton, "Transmitter equalization for 4-Gb/s signaling," *Hot Interconnects Symposium,* Digest of Technical Papers, pp.29-39, Aug. 1996.

[16] R. Farjad-Rad *et al*., "An Equalization Scheme for 10Gb/s 4-PAM Signalling over Long Cables" *IEEE Mixed Signal Conference*, Digest of technical papers, Cancun, Mexico, July 1997.

[17] A. Fiedler *et al*., "A 1.0625 Gbps transceiver with 2x-oversampling and transmit signal pre-emphasis," *IEEE International Solids-State Circuits Conference*, Digest of Technical Papers, pp. 238-9, Feb. 1997

[18] A.X. Widmer *et al*., "Single-chip 4*500-MBd CMOS transceiver", *IEEE Journal of Solid-State Circuits*, vol.31, no.12, pp. 2004-14, Dec. 1996

[19] J. Proakis and M. Salehi, *Communication Systems Engineering.* Prentice Hall, New Jersey, 1994

[20] L. Thon, "540MHz 21mW MDFE equalizer and detector in 0.25-µm CMOS," *IEEE International Solids-State Circuits Conference*, Digest of Technical Papers, Feb. 1998.

[21] D. Xu, Y. Song, and G. T. Uehara, "A 200MHz 9-tap analog equalizer for magnetic read channels in 0.6-µm CMOS," *IEEE International Solids-State Circuits Conference*, Digest of Technical Papers, Feb. 1996.

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

[23] Thomas H. Lee, *The Design of CMOS Radio-Frequency Integrated Circuits.* Cambridge University Press, 1998.

[24] H.O. Johansson, J. Yuan, and C. Svensson, "A 4-GSamp/s Line-Receiver in 0.8µm CMOS," *Symposium on VLSI Circuits*, Digest of Technical Papers, pp. 116-7, June 1996.

[25] J. Crols and M. S. J. Steyaert, "A 1.5 GHz highly linear CMOS downconversion mixer," *IEEE J. Solid-State Circuits,* vol. 30, pp. 736–742, July 1995.

[26] J. Sylvan, "Using sample and hold circuits in multichannel environments," *Power conversion and Intell. Motion,* no. 15, pp. 12–15, Apr. 1989.

[27] M. Sauerwald, "Effects of aperture time and jitter in a sampled data system," National Semiconductor Corporation, 1994, Application Note AD-03.

[28] B. Razavi, *Principles of Data Conversion System Design.* Piscataway, NJ: IEEE Press, 1995.

[29] D. F. Hoeschele, *Analog-to-Digital and Digital-to-Analog Conversion Techniques.* New York, NY: Wiley, 1994.

[30] C. Schiller and P. Byrne, “A 4-GHz 8-b ADC system,” *IEEE J. Solid-State Circuits,* vol. 26, pp. 1781–1789, Dec. 1991.

[31] H.O. Johansson and C. Svensson, “Time resolution of NMOS sampling switches used on low-swing signals,” *IEEE Journal of Solid-State Circuits*, pp. 237-245, Feb. 1998.

[32] C.K.K. Yang, *A multi-Gbps transceiver in CMOS technology.* Ph.D. dissertation, Stanford University, 1998.

[33] B. Ellersick *et. al*., “GAD: A 12-GS/s CMOS 4-bit A/D converter for an equalized multi-level link,” *Symposium on VLSI Circuits*, Digest of Technical Papers, June 1999.

[34] P. R. Gray and R. G. Meyer, *Analysis and design of analog integrated circuits.* John Wiley & Sons, 1984.

[35] Y. P. Tsividis, *Operation and modeling of the MOS transistors.* McGraw-Hill Electrical Engineering series, 1988.

[36] Pelgrom, “Matching Properties of MOS Transistors,” *IEEE Journal of Solid State Circuits*, vol. 24, no. 10, pp. 1433-39*,* Oct. 1989

[37] C.L. Portmann *et. al*., “Metastability in CMOS library elements in reduced supply and technology scaled applications”, *IEEE Journal of Solid State Circuits,* Jan. 1995, vol. 30, no. 1, pp. 39-46

[38] T. Mizuno, “Experimental study of threshold voltage fluctuation due to statistical variation of channel dopant number in MOSFET,” *IEEE Transactions on Electron Devices*, vol.41, no.11, pp. 2216-21, Nov. 1994

[39] W.T. Lynch and H.J. Boll, “Optimization of the Latching Pulse for Dynamic Flip-Flop Sensors,” *IEEE Journal of Solid-State Circuits*, vol.9, no.2, pp. 49-55, Apr. 1974

[40] C.K.K. Yang and M. Horowitz, "A 0.8-μm CMOS 2.5-Gbps oversampling receiver for data recovery," *IEEE International Solids-State Circuits Conference*, Digest of Technical Papers, Feb. 1996.

[41] B. Kim *et al*., "A 30MHz hybrid analog/digital clock recovery circuit in 2-mm CMOS," *IEEE Journal of Solid-State Circuits*, pp. 1385-1394, Dec. 1990.

[42] K. Lee *et al.*, "1.04 GBd low EMI digital video interface system using small swing serial link technique," *IEEE Journal of Solid-State Circuits*, pp. 816-823, May 1998.

[43] B. Razavi, *Monolothic phase-locked loops and clock recovery circuits*. IEEE Press, 1996.

[44] S. Y. Sun, " An analog PLL-based clock and data recovery circuit with high input jitter tolerance," *IEEE Journal of Solid-State Circuits*, pp. 325-330, Apr. 1989.

[45] C. Hogge, " A self correcting clock recovery circuit," *IEEE Journal of Lightwave technology*, pp. 1312-1314*,* Dec. 1985.

[46] John G. Maneatis, "Low-Jitter Process-Independent DLL and PLL Based on Self-Biased Techniques," *IEEE Journal of Solid-State Circuits*, pp. 1723-1732, Nov. 1996.

[47] R. Farjad-rad *et al*, "A 0.3-μm CMOS 8-Gbps 4-PAM serial link transceiver," *IEEE Journal of Solid State Circuits*, May 2000.

[48] F. M. Gardner, "Charge-pump phase locked loops," *IEEE Trans. Comm.*, pp. 1849-1858, Nov. 1980.

[49] T.H. Lee, J. Bulzachelli, "A 155MHz Clock Recovery Delay-and Phase-Locked Loop", *IEEE Journal of Solid-State Circuits*, vol.27, no.12, pp. 1736, Dec. 1992.

[50] D. Mijoskovic *et al.*, " Cell-based fully integrated CMOS synthesizers," *IEEE Journal of Solid-State Circuits*, pp. 271-279, March 1994.

[51] H. Ransijn and P. O'connor, "A PLL-based 2.5Gbps GaAs clock and data regeneration IC," *IEEE Journal of Solid-State Circuits*, pp. 1337-1344, Oct. 1991.

[52] A. Hajimiri and T.H. Lee, "A General Theory of Phase Noise in Electrical Oscillators," *IEEE Journal of Solid-State Circuits*, pp. 179-194, Feb. 1998.

[53] B. Kim *et. al*, “PLL/DLL system noise analysis for low-jitter clock synthesizer design,” *Proc. of ISCAS*, June 1994.

[54] J.G. Maneatis *et al.*, “Low-jitter process-independent DLL and PLL based on self-biased techniques,” *IEEE Journal of Solid-State Circuits*, vol.31, no.11, p. 1723-32, Nov. 1996.

[55] D.K. Su *et al.*, “Experimental results and modeling techniques for substrate noise in mixed-signal integrated circuits,” *IEEE Journal of Solid-State Circuits*, Nov. 1996.

[56] N. Verghese and D. Allstot, *Simulation techniques and solutions for mixed-signal coupling in integratedcircuits*. Kluwer Academic Publishers, 1995.

[57] A. Pottbacker and U. Langmann, “An 8-GHZ silicon bipolar clock-recovery and data-regenerator,” *IEEE Journal of Solid-State Circuits*, p. 1572-76, Dec. 1994.

[58] L. M. Devito, “ A versatile clock recovery architecture and monolithic implementation,” *Monolithic phase-locked loops and clock recovery circuits,* pp. 405-421, IEEE Press, 1996.

[59] M.Y. Hsiao and K.Y. Sih, “Serial-to-Parallel Transformation of Linear-Feedback Shift-Register Circuits,” *IEEE Transaction on Electronic Computers*, pp. 738-40, Dec. 1968.

[60] A. Lempel, W.L. Eastman, “High Speed Generation of Maximal Length Sequences,” *IEEE Transaction on Computers*, vol. C-20, no.2, pp 227-9, Feb. 1971.

[61] L. Andersson *et al.*, “Silicon Bipolar Chipset for SONET/SDH 10 Gb/s Fiber-Optic Communication Links,” *IEEE Journal of Solid-State Circuits*, vol.30, no.3, pp. 210-218, Mar. 1995.

[62] 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*, vol.8, no.1, pp. 33-44*,* Jan. 1989.

[63] R. Best, *Phase-Locked Loops*. 3rd ed, McGraw Hill, 1997.

[64] K. Chen, et.al., "Predicting CMOS Speed with Gate Oxide and Voltage Scaling and Interconnect Loading Effects," *IEEE Transactions on Electron Device*, vol.44, no.11, pp. 1951-7, Nov. 1997.

[65] R. Co, J.H. Mulligan Jr., "Optimization of Phase-Locked Loop Performance in Data Recovery Systems," *IEEE Journal of Solid-State Circuits*, vol.29, no.9, pp. 1022-1034, Sep. 1994.

[66] M. Dolle, "A Dynamic Line-Termination Circuit for Multireceiver Nets," *IEEE Journal of Solid-State Circuits*, vol.28 no.12, pp. 1370, Dec. 1996.

[67] K.S. 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 p. 1995-2003, Dec. 1996.

[68] C.A. Eldering *et al.*, "Digital Burst Mode Clock Recovery Technique for Fiber-Optic Systems," *IEEE Journal of Solid-State Circuits*, vol.29, no.2, pp. 271, Feb. 1994.

[69] F. Gardner, *Phase Lock Techniques.* John Wiley, 1979.

[70] C. Gray *et al.*, "A Sampling Technique and Its CMOS Implementation with 1-Gb/s Bandwidth and 25ps Resolution," *IEEE Journal of Solid State Circuits*, vol.29, no.3, pp. 340-9, Mar. 1994.

[71] M.J. Izzard and C.G. Thisell, "Low Power Analog Control Clock Synchronizer for 3Gb/s Data Retiming in ATM Applications," *TI ACS Report*, Nov. 1994.

[72] D.A. Johns, D. Essig, "Integrated Circuits for Data Trasmission Over Twisted-Pair Channels," *IEEE Journal of Solid-State Circuits*, vol.32, no.3, pp. 398-406, Mar. 1997.

[73] M. Kurisu, et.al, K. Yoshida, "2.8-Gb/s 176-mW Byte-Interleaved and 3.0-Gb/s 118-mW Bit-Interleaved 8:1 Multiplexers with a 0.15-µm CMOS Technology," *IEEE Journal of Solid-State Circuits*, vol.31, no.12, pp. 2024-2029, Dec. 1996.

[74] Z. Lao, U. Langmann, "Design of a Low-Power 10Gb/s Si Bipolar 1:16 Demultiplexer IC," *IEEE Journal of Solid-State Circuits*, vol.31, no.1, pp. 128-131, Jan. 1996.

[75] S.J. Lee, B. Kim, K. Lee, “A Novel High Speed Ring Oscillator for Multi-Phase Clock Generation Using Negative Skewed Delay Scheme,” *IEEE Journal of Solid-State Circuits*, vol.32, no.2, pp. 289-91, Feb. 1997.

[76] R. Matick, *Transmission lines for digital and communication networks.* 3rd edition, IEEE Press, 1997.

[77] J.E. McNamara, *Technical Aspects of Data Communication*. 2nd ed. Digital Press, Bedford, MA, 1982.

[78] J.A. McNeill, “Jitter in Ring Oscillators,” *IEEE Journal of Solid-State Circuits*, vol.32, no.6, pp. 870-879, Jun. 1997.

[79] R. Mooney *et al*., “A 900 Mb/s bidirectional signaling scheme,” *IEEE Journal of Solid-State Circuits,* vol.30, no.12, p. 1538-43, Dec. 1995.

[80] T. Sekiguchi *et al.*, “Low-Noise, High-Speed Data Transmission Using a Ringing-Cancelling Output Buffer,” *IEEE Journal of Solid-State Circuits*, Nov. 1995, vol.30, no.12, pp. 1569-74

[81] S. Sidiropoulos *et al*., “A CMOS 500 Mbps/pin synchronous point to point link interface,” *IEEE Symposium on VLSI Circuits.* Digest of Technical Papers*,* pp. 43-4, June 1996.

[82] S. Sidiropoulos *et al*., “A 700-Mb/s/pin CMOS signaling interface using current integrating receivers,” *IEEE Journal of Solid-State Circuits*, vol.32, no.5, pp. 681-90*,* May 1997.

[83] S. Sidiropoulos, M.A. Horowitz, “A Semidigital Dual Delay-Locked Loop,” *IEEE Journal of Solid-State Circuits*, vol.32, no.11, pp. 1683-1692, Nov. 1997.

[84] B.S. Song *et al*., “NRZ timing recovery technique for band limited channels,” *IEEE International Solid-State Circuits Conference.* Digest of Technical Papers, p. 194-5, June 1996.

[85] M. Galles *et al*., “Spider: a high-speed network interconnect,” *IEEE Micro*, vol.17, no.1, p. 34-9, Jan.-Feb. 1997.

[86] J. Kuskin *et al* , "The Stanford FLASH Multiprocessor," *Proceedings of the 21st International Symposium on Computer Architecture*, Chicago, IL, pp. 302-313

[87] N. Kushiyama *et al*., "A 500-megabyte/s data-rate 4.5 M DRAM," *IEEE Journal of Solid-State Circuits*, vol.28, no.4, p. 490-8, Apr 1994.

[88] R.C. Foss *et al*., "Memory- fast interfaced for DRAMs," *IEEE Spectrum,* vol.29, no.10, pp. 54-7, Oct 1992.

[89] K.Y. Chang *et al*., "A 2-Gb/s Asymmetric Link for High-Bandwidth Packet Switches," *Hot Interconnects V Symposium Record*, Digest of technical papers, pp. 171-179, Aug. 1997.

[90] T.H. Hu *et al*., "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, p. 1314-20, Dec. 1993.

[91] R.G. Swartz *et al*., "A Burst mode, Packet Receiver with Precision Reset and Automatic Dark Level Compensation for Optical Bus Communications," *IEEE Journal of Solid-State Circuits*, vol.29, no.2, pp. 325, Feb. 1994.

[92] T. Takahashi, et. al. "A CMOS Gate Array with 600 Mb/s Simultaneous Bidirectional I/O Circuits," *IEEE Journal of Solid State Circuits*, vol. 30, no. 12. pp. 1544-46, Dec. 1995.

[93] A.X. Widmer, P.A. Franaszek, "A DC-Balanced Partitioned-Block 8B/10B Transmission Code," *IBM Journal R&D*, vol.27, pp. 440-451*,* Sep 1983.

[94] B. Razavi, B.A. Wooley, "Design Techniques for High-Speed, High-Resolution Comparators," *IEEE Journal of Solid-State Circuits*, vol.27, no.12, pp. 1916-26, Dec. 1992.

[95] M. E. Lee, et al. "A 90-mW 4-Gbps Equalized I/O Circuit with Input Offset Cancellation," *ISSCC Digest of Tech. Papers*, pp. 252-3, Feb. 2000.

[96] K. Iravani, et al. "Clock and Data Recovery for 1.25Gbps Ethernet Transceiver in 0.35-μm CMOS," *CICC Digest of Tech. Papers*, pp. 261-4, May 1998.

[97] M. Fukaishi, et al. "A 4.25Gbps CMOS Fibre Channel Transceiver with Asynchronous Tree-Type Demultiplexer and Frequency Conversion," *IEEE Journal of Solid-State Circuits*, vol.33, no.12, pp. 2139-47, Dec. 1998.

[98] S. Sidiropoulos *et al*., “Adaptive Bandwidth DLLs and PLLs using regulated Supply CMOS Buffers,” *VLSI Symposium Digest of Technical Papers*, pp. 124-7, June 2000.