PRECISON CLOCK SYNTHESIS USING DIRECT MODULATION OF FRONT-END MULTIPLEXERS/DEMULTIPLEXERS IN HIGH SPEED SERIAL LINK TRANSCEIVERS

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

Patrick Chiang
July 2007

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

(William J. Dally) Principal Adviser

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 the University Committee on Graduate Studies.

Abstract

High-speed CMOS serial links have been crucial in matching off-chip system bandwidth with the increasing on-chip demand. Conventional serial link architectures typically use multi-phase clocking structures to achieve a pin bandwidth faster than the on-die logic switching speed. However, such multi-phase clocking architectures typically use several stages of clock fanout buffering, which dissipate considerable power and suffer from significant sources of timing uncertainty, namely power supply induced jitter and process dependent static phase offset.

This thesis presents a new serial link architecture which addresses the timing uncertainties caused by power supply noise and process mismatch. First, process variation and power supply sensitivity in serial link clock buffers are examined and then shown to degrade further in future scaled CMOS processes. Next, a new architecture is presented that alleviates this problem: direct drive resonant clocking. The complementary phases of an integrated LC-VCO directly drive the final output multiplexer in the transmitter, resonating the capacitive load and eliminating clock buffers, thereby reducing power dissipation, power supply induced jitter, and static phase offset. In the receiver, a similar technique is applied, where the front-end 10GHz input sampler is directly driven by a different LC-VCO. Several side-effects of this direct drive resonant clocking technique such as increased kickback-induced jitter, reduced tuning range, and reduced bandwidth are examined. Finally, two test chips are designed and fabricated in 0.13um CMOS technology, exhibiting a 20Gb/s data rate, low power(165mW in the transmitter), and low area. The resulting power supply susceptibility is reduced by 10x and process mismatch phase error by 5x.

Acknowledgements

It is said that a person cannot fully understand and appreciate his own life until he steps away from his local surroundings, opens his eyes, and sees how things in the rest of the world are so different and yet so similar. In many ways, my graduate life at Stanford encompasses these same feelings. After leaving Stanford, I have traveled to other places and worked closely with many people from India, China, and other places in America. Never before did I understand the magnitude of my graduate school opportunities (some students described these as "chances"). I am extremely thankful that I was able to attend Stanford and participate in a place with so many "chances".

The first person I would like to thank is my advisor Professor Bill Dally. I have never met a person who can understand so many various topics with such depth, insight and ease. I still can recall some of the group meeting talks he gave–one quarter he lectured about complier optimization for code parallelization, while another Professor Dally built a CMOS multiplying delay-locked loop, and showed the SPICE simulations for it. Such breadth of knowledge and foresight are absolutely critical for academic research breakthroughs, and I can now fully appreciate the opportunity to see and learn from such experiences. More importantly, I would like to thank Professor Dally for giving me the opportunity to challenge myself and my research. Looking back, I have worked on some amazing projects with some amazing objectives. Through these experiences, I have learned to think the impossible, take risks, and not be afraid to learn from any mistakes.

Second, I would like to thank my co-advisor Mark Horowitz for his support and encouragement over the years. While Mark is well-known in academia and industry

for VLSI circuit design, I would like to thank him most for the things not described in the journal publications, but for his general concern and care for graduate students and their learning experience at Stanford. I will not forget two instances in particular– the time he wrote a 2-page detailed email in response to my first testchip publication, which included a hour-long phone conversation over the Christmas holiday; and for a talk he gave to incoming graduate students, recommending them to take time from their research in order to take advantage of the numerous resources/curriculum available in other disciplines.

Third, I'd like to thank Professor Wooley for being the third reader for my thesis, and for taking the time in graduate school to give me advice on staying in school and away from the monetary distractions that occur so often around Silicon Valley.

Stanford is the amazing place that it is not because of the tangible things that are physically located there, but because of the brilliant and smart people that work so hard to pursue cutting-edge ideas. Some people I'd like to thank are: Pauline, Anting, Rich, Frank, Bruce-k, Mattan, Ed, Beth, John, C, Patience, Renee, Sam, Alan, and likely others I have forgotten. The bay area also brought about many friends, and I'd like to thank them for basketball interludes and other non-academic distractions: Jonathan, Matt, Matt, Steve, Dan, basketball friends, Bert, Jyan, Nancy, Diana, Sam, etc.

I'd like to thank my sister Sophia and brother-in-law Andy, along with the small ones Willow and Lleyton. They have supported me for many, many years–back when I would drive my sister's Miata around Berkeley Campus, to the times I needed to borrow a bicycle or needed a place to stay while in the Bay Area. I am fortunate to have such close siblings, and I thank them for their unwavering support of me through graduate school.

Last, I would like to thank my parents Peter and Eleanor for raising me and handing me opportunities I can only appreciate the magnitude of until now. From playing soccer to learning violin to not worrying about financial constraints in order to concentrate of studies, I have experienced in my personal growth luxuries that have allowed me to pursue my dreams. While it has always been a constant struggle to find my own way while also living up to my parents standards, I now honestly admit

that my parents' goals of higher education and academic pursuits is a challenge that I find in myself my higher calling. I hope that in my lifetime, I can only pay my parents back by instilling future young students the desire and hunger to pursue their academic dreams, and hope that their future "chances" are bigger, wilder, and more successful than my own.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Overview of Problem

Continued advancement of technology scaling, as predicted by Moore's law, reduces transistor and wiring size, contributing towards the exponential increase in computation and bandwidth on a single chip. However, moving that potential computing capability off-chip does not significantly improve with scaling due to fundamental limitations in the off-chip physical properties, resulting in a growing gap between on-chip and off-chip bandwidth. It is therefore essential to investigate new architectures that use technology scaling to improve interconnect bandwidth and mitigate this on-chip/off-chip divergence.

With this increase in data bandwidth, timing precision becomes a critical limiting factor. Conventional serial links currently address the timing generation issue by creating a clean local clock with a frequency synthesizer, whose output is followed by several stages of buffering from the synthesizer to the point of use at the multiplexer/demultiplexer input. While this technique is adequate for current data rates (1 - 10Gb/s), such conventional high speed serial link architectures may not be suitable at extremely high data rates. The reason is that technology scaling degrades the process matching and power supply rejection of these clock buffers, affecting the timing precision of clock synthesis and ultimately limiting the scalability of conventional serial links for extremely high data rates (i.e. 20Gb/s).

1.2 Solution

In this thesis, a new clocking architecture is proposed that addresses these timing uncertainty considerations. This architecture, by directly driving the transmitter multiplexer and receiver demultiplexer with the clock synthesizer, eliminates the need for clock buffers, thereby eliminating crucial components that are susceptible to process mismatch and supply induced jitter. More specifically, a LC-VCO (resonator-type voltage controlled oscillator) is used in the clock synthesizer and directly drives the load by subsuming the load capacitance (in either transmitter multiplexer or receiver demultiplexer) directly in the LC resonator. This has several benefits. First, the power supply sensitivity is determined only by the LC-VCO power supply rejection, typically more than an order of magnitude better than a common ring oscillator. Secondly, the static phase mismatch of LC-VCO complementary phases is extremely low, even with the possible introduction of capacitance and threshold voltage mismatch within the oscillator. Finally, power consumption is reduced as the large transistor capacitance loading of the final multiplexer/demultiplexer stages is subsumed directly into the LC-VCO.

Two 20Gb/s 0.13um CMOS prototype chips were built to validate this concept. The first prototype illustrates the concept of a LC-VCO directly resonating the transmitter multiplexer, eliminating the need for clock buffers, and reducing power supply induced jitter and static phase mismatch. This first prototype also implements a data multiplexing structure that uses an on-chip two-tap FIR filter to compensate for bandwidth reduction and latch hysteresis in the 10Gb/s, 4:1 preamp multiplexer. The second prototype implements an improved version of this direct drive resonator technique – the new transmitter 2:1 output multiplexer topology eliminates the susceptibility of data dependent kickback noise. This second transmitter prototype also implements a two-tap FIR filter that is used to compensate for intersymbol interference and channel loss, since there is significant frequency dependent attenuation expected at the Nyquist frequency of 10GHz. This second prototype also introduces our direct LC-VCO method in a demultiplexing receiver – two complementary phases

are used for the front end 1:2 demultiplexing track-and-hold circuit. An architecture for demultiplexing the receive data by multiple 1:2 demultiplexing before actual quantization to the digital domain is introduced (called downsampling demultiplexing). This technique relaxes power and design difficulty in building a high frequency slicer array. Eventually, the eight 2.5Gb/s downsampled analog signals are quantized by offset compensated slicers. Finally, the receiver implements a baud rate clock and data recovery architecture, using an on-die BER measurement system to find the optimum clock placement. This architecture reduces system overhead/complexity and power as compared to a 2x oversampled receiver approach.

1.3 Organization

Chapter 2 discusses the background for understanding the importance of minimizing the timing uncertainty issue, especially for serial link data rates greater than 10Gb/s. This section will describe two major causes for timing uncertainty, static phase offset and power supply induced jitter in both the oscillator as well as the clock buffers. After understanding the causes for timing uncertainty, this chapter will describe how conventional serial link architectures are susceptible to both static phase offset and supply-induced jitter, making these architectures not amenable to future very high date rate serial links.

Chapter 3 presents the clocking solution proposed in this thesis: resonating the output multiplexer in the transmitter (or input demultiplexer in the receiver) directly from the clock source itself using the LC resonator of the clock synthesis. This chapter describes the power supply sensitivity of the LC-resonator, the reduction in static phase offset, and reduced power consumption due to resonance of the parasitic capacitance. This chapter also discusses the intrinsic static phase offset mismatch due to capacitance and threshold mismatch in the voltage controlled oscillator, by using HSPICE simulation. Degradation in LC resonator quality factor is analyzed, and jitter enhancement due to data-dependent jitter kickback is evaluated.

Chapter 4 describes the first generation prototype of a 20Gb/s transmitter that uses the direct drive LC method as described in this thesis. This architecture achieves

20Gb/s by creating eight 2.5Gb/s low speed data streams, 4:1 multiplexing these streams into two 10Gb/s data streams, retiming these data streams using an analog retiming latch, and then 2:1 multiplexing the 10Gb/s data streams into one 20Gb/s data stream. The 4:1 10Gb/s multiplexer structure is described along with its bandwidth limitations. The 10Ghz analog latch is also described, and its negative hysteresis effect is also described. To compensate for these bandwidth limitations, a two-tap FIR pre-emphasis filter is built for the 10Gb/s data streams, compensating for both the bandwidth limitations and the analog latch hysteresis. Finally, in the 20Gb/s 2:1 selector, data dependent kickback noise is analyzed and shown to be negligible.

Chapter 5 illustrates the second generation transmitter that is modified from the previous architecture. This transmitter consists of a 2:1 first stage multiplexer achieving a 5Gb/s data rate, a 2:1 CML multiplexer to achieve 10Gb/s data, and a final 2:1 multiplexer to achieve 20Gb/s output data. No explicit retimer/latching is involved, necessitating timing precision. The 10Gb/s 2:1 multiplexer is described with its bandwidth and process variation/sensitivity discussed. A revised 2:1 20Gb/s output multiplexer is described that removes any of the residual data dependent jitter that was predicted using the first generation architecture.

Chapter 6 describes the architecture and circuits used for a 20Gb/s demultiplexer receiver. First, the receiver demultiplexes the incoming data stream 1:2, from 20Gb/s to 10Gb/s, with the sampler clock coming directly from the LC-VCO. As 10Gb/s is still too fast for slicing the input, the 10Gb/s data streams are again analog 1:2 demultiplexed into 5Gb/s, at which time they are converted from analog to digital. A bit error rate backend counts bit errors, and thus, the optimum clock placement for demultiplexing can be determined by advancing the sampling clock and then measuring the BER bathtub curve. The effect of sinusoidal clocking and common mode bias are analyzed for the receiver input bandwidth.

Chapter 7 presents the experimental results of the two fabricated test chips. The first 0.13um CMOS test chip illustrates the 20Gb/s transmitter data rate using the LC-VCO to directly modulate the output driver. The second 0.13um CMOS test chip shows a revised 20Gb/s transmitter and a 20Gb/s receiver.

Chapter 8 summarizes the results, suggests future work, and concludes this thesis.

Chapter 2

Background

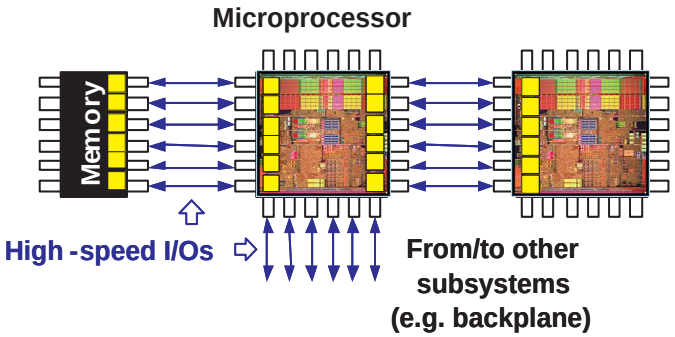

Figure 2.1: Serial Link Interconnections

2.1 Overview

High speed serial links (I/Os) are critical circuit blocks on microchips that allow separate chips to communicate information between each other. A simple example of using a serial link can be see in a microprocessor like the IBM Cell processor [44] in Figure 2.1. The processor above needs to store and load data from an off-chip memory or communicate with another Cell processor; serial links allow this data to be transmitted and received, connecting these different chips. The performance of serial links is a critical component to improved computer system performance, and

in many designs, the bottleneck for these systems. The reason is that relative to the on-die bandwidth that has improved significantly with technology, the off-chip bandwidth (i.e. serial link speeds) has not improved at the same rate. Many previous works have described this phenomenon. [45, 46, 47]

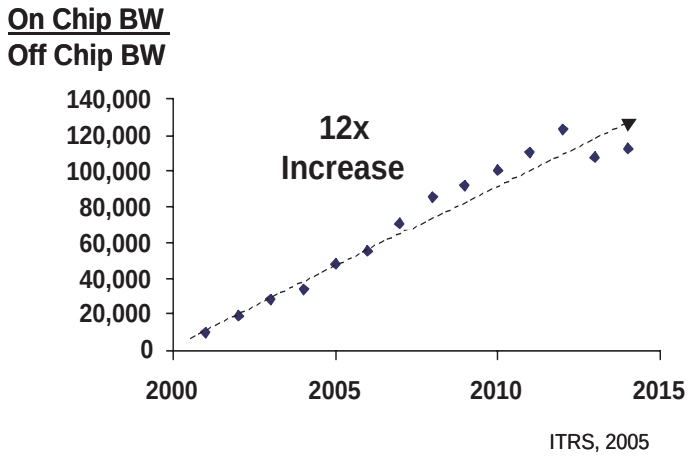

Figure 2.2: Disparity Between On-chip and Off-chip Bandwidth

Figure 2.2 shows a predicted ratio of the maximum possible bandwidth between the on and off chip bandwidth–it can be seen that bandwidth disparity is increasing over time, using data extrapolated from the ITRS roadmap [1]. This phenomenon of the disparity between on-chip/off-chip bandwidth is dictated by the difference in manufacturing technologies between the two bandwidths. [1] On-chip bandwidth of a microprocessor increases exponentially due to Moore's law scaling. For example, deeper submicron CMOS technology helps improve transistor switching speed, increases the number of metal wiring layers, reduces the pitch between parallel wires, and increases the physical length and width of a microprocessor die; all of these improvements in technology improve the possible on-chip bisection bandwidth. However, the maximum possible off-chip bandwidth is limited to the number of I/O pads

[1]Graph is extrapolated by calculating the possible bisection bandwidth on a CMOS die vs. the bandwidth coming off of a die. The on-die bandwidth is calculated by the speed of on-die wires, governed by the width and spacing of wires, along with the number of wiring layers and maximum die size. The off-chip bandwidth is calculated by the expected increase in the number of pads per die, along with the link speed of each I/O pad.

that can be physically placed on the surface of a microprocessor or the corresponding ball grid array package as well as the interconnection speed of each physical pad. As the number of chip I/O pads does not increase very rapidly with technology improvements, and the bandwidth/pin speed is limited by the physical parasitics of the interconnect channel, this disparity in bandwidth ratio increases over time. Therefore, new developments in improving the data rate of serial link are crucial in mitigating the gap between these two different bandwidths.

It should be noted that while serial link data speed is an important metric, two other factors below are also crucial towards future microprocessor performance–these two other factors have not been addressed directly by this particular thesis work. The first factor is the importance of the link power dissipation. As the number of these links placed on a microprocessor increases, the power dissipation of these links becomes a significant fraction of the total dissipated power. For example, a serial link running at 2.5Gb/s in 0.13um process may dissipate about 100mW [50]. For a microprocessor incorporating 400 such links, the total I/O power is 40W, or about 40 percent of the total power on a 100W processor. For future link development, the metric of power dissipated per gigabit of bandwidth needs to continue to improve, and there has been much research work in this area. [50, 51, 52, 48, 49].

The second factor that is crucial toward link design is in serial link channel equalization methods. Due to channel losses (i.e. reflections, via stubs, frequency-dependent attenuation), it is difficult to signal at high frequencies through this lossy off-chip interconnect [45]. Due to the high-frequency signal attenuation, the channel essentially acts as a low-pass filter and therefore significant intersymbol interference (ISI) reduces the signal-to-noise ratio at the receiver. State-of-the-art serial links have incorporated many types of equalization topologies to combat this high-frequency signal loss, such as transmit pre-emphasis [50], receiver equalization [39], multi-level signaling [37], decision feedback equalization [38], duobinary signaling [36], among others. As the channel attenuation issue is a serious problem, many researchers have addressed this issue and therefore it is not address in this thesis.

2.2 Timing Uncertainty

As data rates increase from 1Gb/s to 10Gb/s and beyond, the importance of timing precision in link performance becomes a dominant factor in the ability to scale off-chip with on-chip bandwidth.

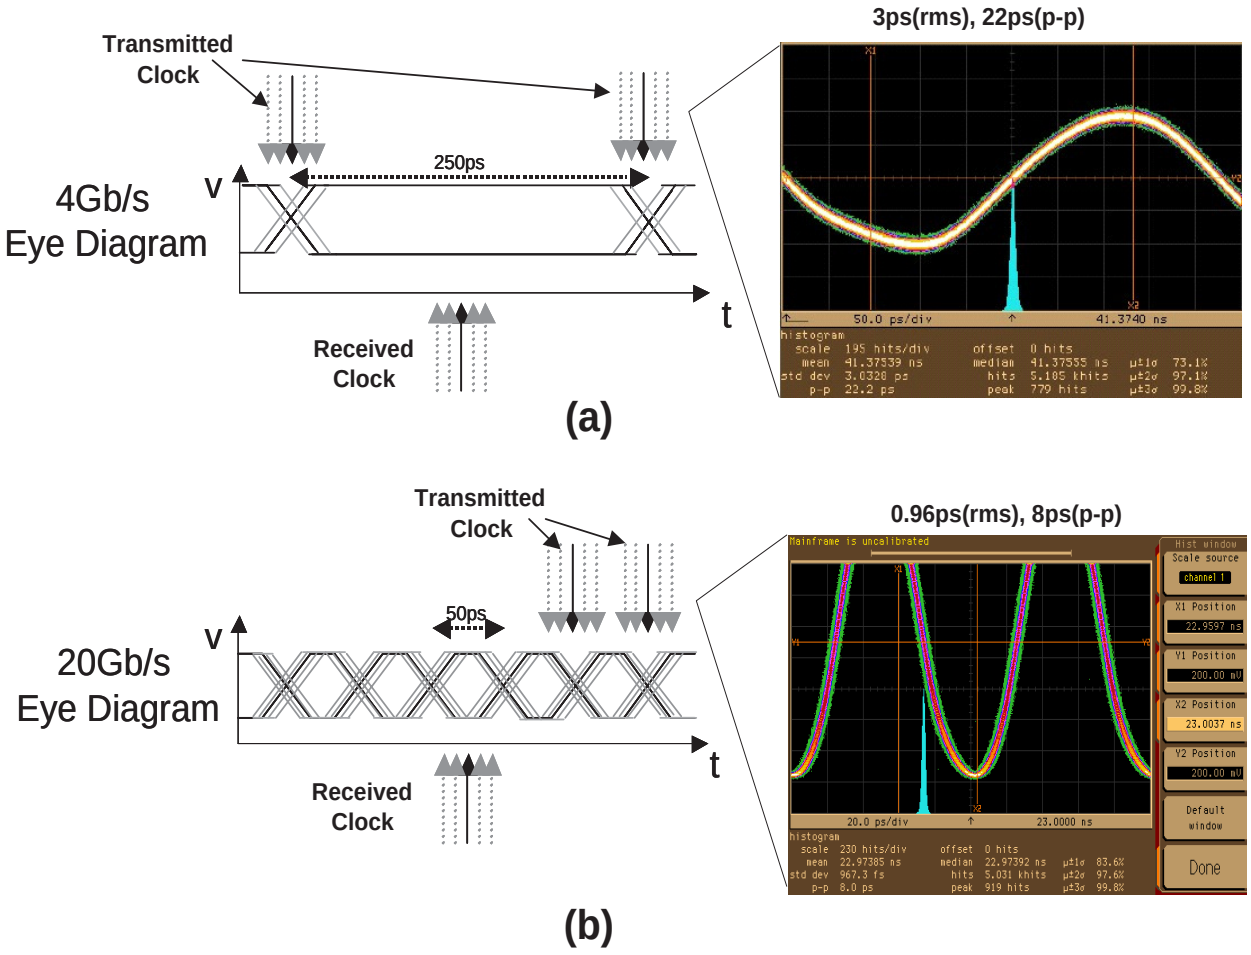

Figure 2.3: Importance of Timing Precision at Higher Data Rates

The simple example in Figure 2.3(a) illustrates the importance of timing uncertainty at a data rate of 20Gb/s. Take as an example a previously designed serial link in 0.25um CMOS running at 4Gb/s (250ps per bit) [50]. It illustrates an internal synthesized clock jitter in the transmitter of 3.0ps(RMS), 22.2ps(pk-pk). This synthesized clock is used to create both transmitter data edges as well as the receiver

sampling clock. Assuming the receiver sampling clock uses the synthesized internal clock similar to that in the transmitter, to first order the clock uncertainty and eye diagram looks similar to the Figure 2.3(a). This results in a relatively large timing margin of 84 percent, since the clock uncertainty is still a small fraction of the entire bit time.

Consider now a 20Gb/s serial link data rate (50ps per bit), Figure 2.3b where the data rate has increased by five times – now there exists very little timing margin. Indeed, if we used a synthesized clock with jitter performance similar to the lower data rate (4Gb/s) for clocking and retiming our new 20Gb/s data stream, we now exhibit a timing margin of only 20 percent, resulting in a significant increase in bit error rate. However, if the 20Gb/s serial link uses a synthesized clock developed in this thesis [40, 41], the synthesized clock jitter and uncertainty is reduced to 0.97ps(RMS), 8ps(pk-pk). This corresponds to a timing margin of 60 percent, now a significant improvement.

One major assumption made in this example is the conjecture that technology scaling doesn't decrease jitter absolutely (in regards to picoseconds) when migrating a conventional architecture to a finer line length CMOS technology for faster data rate [8]. It will be shown in the following sections how transistor scaling doesn't necessarily decrease synthesized clock jitter. Therefore, it will be necessary to design a new circuit architecture as the current serial link circuitry at 4Gb/s does not scale well to 20Gb/s serial links and beyond.

The rest of the chapter is organized as follows: Section 2.3 discusses how synthesized clock timing jitter arises in high speed serial links. Section 2.5 describes the phenomenon of static phase offset, which arises from DC process and layout mismatches causing a systematic DC mismatch and frequency spur. Section 2.6 describes the high frequency jitter component, which is caused by high-frequency power supply fluctuations. After the mechanisms for clock jitter are understood, Section 2.8 gives an example of a conventional serial link architecture and the timing uncertainty that arises from such an architecture.

2.3 Serial Link Circuit Block Diagram

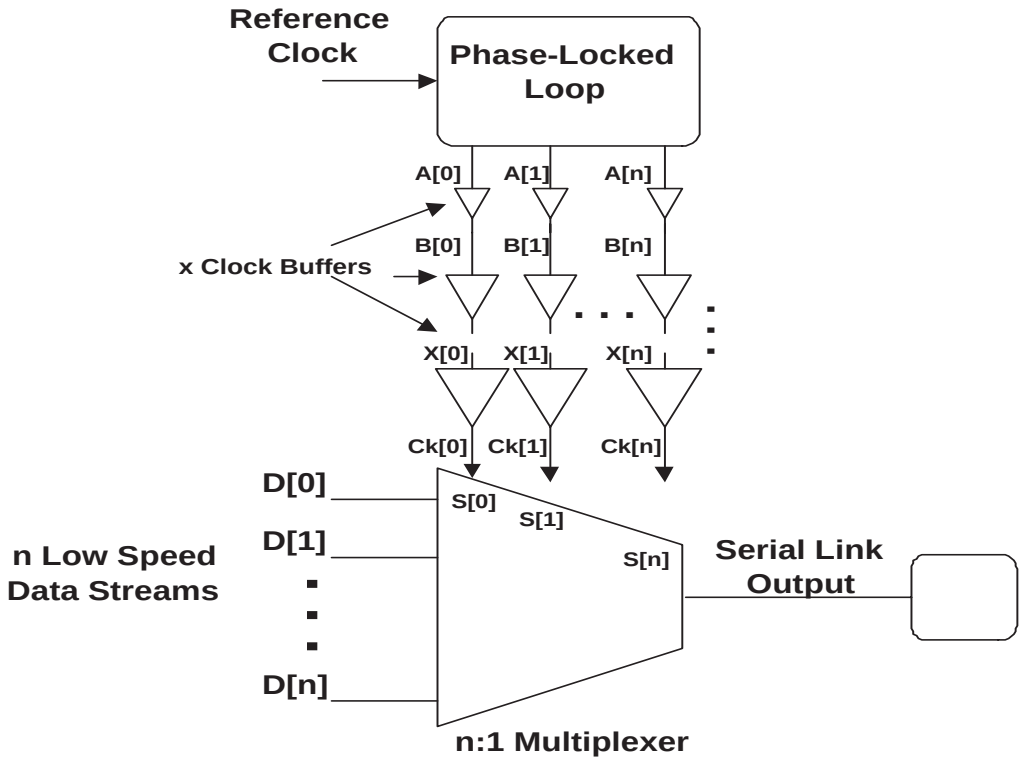

Figure 2.4: Top Level Diagram of Clock Synthesis and Data Multiplexing in a High Speed Serial Link

Figure 2.4 shows a high level view of the design and implementation of a high speed serial link. Multiple data streams (n) at a low data rate are up-multiplexed to a higher data rate using multiple clock phases (n) synthesized from a phase locked loop (PLL). The PLL essentially creates high frequency, time-interleaved clocks A[n:0] on-die from a low frequency, off-chip reference clock. These synthesized clock phases are then passed through several clock buffers before these phases are finally used at the multiplexer input to clock the (n) data streams.

Notice that the number of clock phases needed is the same as the upconversion multiplexing ratio. For example, in Figure 2.5(a) of a 2:1 transmitter multiplexing output stage, two phases of complementary clock are needed to multiplex two data streams to the serial link output. This is in contrast with Figure 2.5(b)), where a 4:1

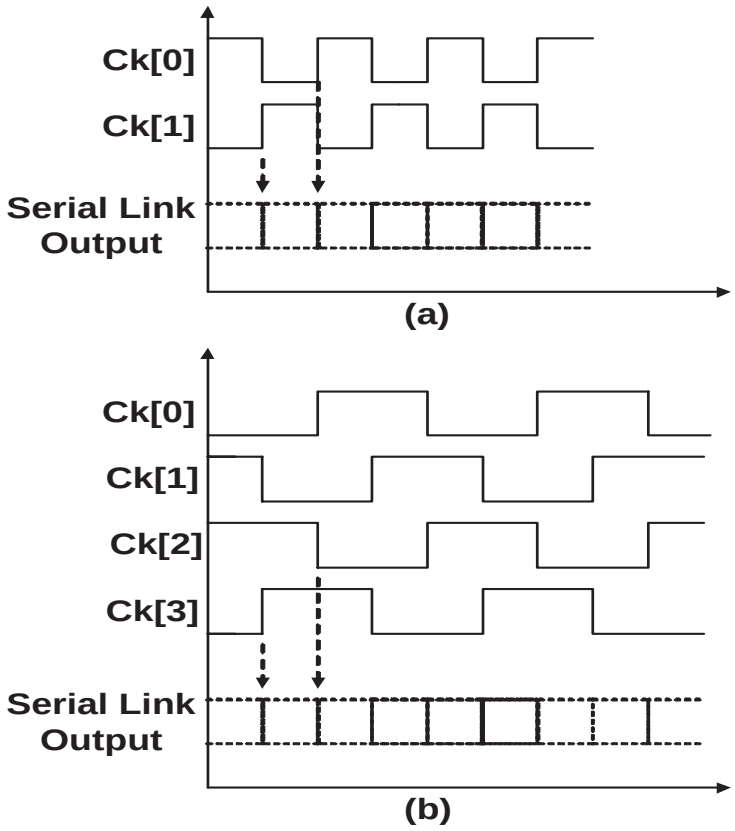

Figure 2.5: a) 2:1 Transmitter Multiplexing, b) 4:1 Transmitter Multiplexing

transmitter multiplexing requires four phases of time-interleaved clocks to multiplex four input data streams to the final serial link output speed. As will be shown in the following sections, while higher multiplexing ratios (i.e. 4:1) improve design ease since both the clocks and the data elements are running at slower clock frequencies compared to low multiplexing ratios (i.e. 2:1), higher data multiplexing also introduces timing uncertainty as a result of the difficulty in generating precise multiple clock phases and synchronizing them with the multiple data streams.

2.4 Origination of Timing Uncertainty

Typically, since the PLL's synthesized clock is multiplied up from the low-frequency reference clock, there is accumulated jitter in the phase locked loop, making the PLL a

large source of clock jitter. Design of low jitter PLLs has been researched heavily in the literature [10]. There has been much work on power-supply insensitive PLLs [21, 50], differential Current Mode Logic(CML) PLLs [18], low phase noise ring oscillators [16], and on-die LC oscillators [17]. In general, the feedback of a phase locked loop rejects any noise accumulated in the VCO (voltage controlled oscillator) based upon the bandwidth of the PLL feedback loop. A larger loop bandwidth results in the output PLL noise tracking the input reference closely, while a smaller loop bandwidth implies that the PLL does not track reference clock noise. Choice of loop bandwidth depends on the location of the noisiest clock reference–either the on-die voltage controlled oscillator or the off-chip reference clock. A ring oscillator VCO typically has a large jitter sensitivity to power supply noise, such that the loop bandwidth will be large (in the MHz range) such that the PLL output noise will track the relatively clean input reference clock. A LC-VCO uses a passive inductor and capacitor to form its oscillation frequency, and therefore is more immune to power supply noise. Hence, for some high frequency LC oscillator synthesizers, the reference clock may be noisier than the LC-VCO, and the PLL loop bandwidth may be much lower (in the 100kHz range or lower).

Typically after the synthesized clock phases have been generated from a low jitter PLL, they are then sent through multiple stages of clock buffering. For example, in Figure 2.6, one of the multiple PLL phases is sent through three stages of buffering before being used at the transmitter multiplexer. These buffers are used for several reasons. One is that these buffers allow the PLL to be placed at some distance away from the point of use, so that there is isolation of the noisy digital circuitry from the noise-sensitive PLL. Secondly, the PLL is optimized to create precise clock phases with low timing uncertainty, but their fanout or ability to drive large capacitance is limited. Hence, the buffers allow for increased fanout or drive strength, both in driving long metal trace capacitance as well as transmitter/receiver multiplexer input capacitance. Additionally, the clock buffers can allow the sharing of one PLL with multiple links, where the clock buffering provides the clock distribution from one centrally located PLL.

Each clock buffer is typically an inverter, or some variation of an inverter topology

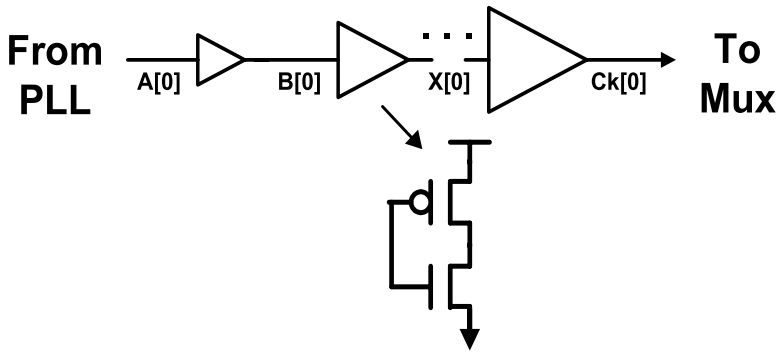

Figure 2.6: Clock Buffering from PLL Output to Transmitter Multiplexer Clock Input

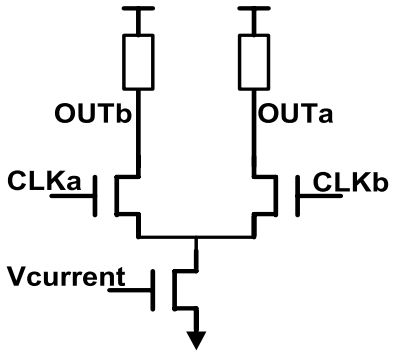

Figure 2.7: Current Mode Logic (CML) Clock Buffer

[21]. It will be described in the following sections why such topologies can be the major culprits in creating larger timing uncertainty.

It should be mentioned that another type of clock buffer that is often used is the Current Mode Logic (CML) buffer, seen in Figure 2.7. The load for the differential pair normally consists of some type of low-impedance, such as a diode connected PMOS [18], a simple resistor, or a variation such as a resistor with peaking inductor in series or parallel with the inductor. This type of buffer is typically used when the data is so fast that conventional digital inverter topologies do not act as effective clock buffers. From a jitter sensitivity standpoint, power supply rejection is quite good due to the differential nature of the CML. [2] Due to their resistive loading and

[2]Static phase offset, however, may be exacerbated with CML buffers, since mismatches may exist not only in transistors, but also in the wiring capacitance, via resistance, and load resistors, causing duty cycle distortion.

constant tail current bias, power dissipation is a major disadvantage for these types of buffers. The next sections will consider both inverter-based as well as CML-based clock buffering techniques and the advantages/ disadvantages of both.

While jitter arising from phase locked loops has been one of the most heavily researched area, jitter from the post-PLL clock buffers has not been investigated thoroughly [21, 13]. While at first glance it may seem that clock buffers contribute minimal clock jitter since the clock is not multiplied as in a PLL (and therefore the jitter is not accumulated), such buffers contribute two significant sources of synthesized clock jitter: static phase offset and power supply induced jitter. These two sources of jitter will now be explained further; it will be seen that such post-PLL clock buffer jitter can significantly reduce serial link performance, even if the PLL itself has superior performance. In some sense, creating a precision phase locked loop with optimum low phase noise/jitter characteristics may be overshadowed by the jitter introduced by the post-PLL clock buffers.

2.5 Static Phase Offset

Static phase errors result as systematic DC phase shifts in the arrival of a clock edge, from where it was suppose to cross to where it actually switches. Take Figure 2.8 for example, which is slightly revised from Figure 2.5(b). Due to duty cycle error in the clock[2] waveform its falling clock edge is off-phase by 20 percent from nominal, essentially causing the duty cycle to be distorted. As this is a DC type of error, every clock period will exhibit this same 20 percent distorted clock cycle. If both rising and falling edges of this asymmetric clock are used to multiplex data to the high speed I/O output, then the transmitted data waveform will also exhibit the same asymmetry, and therefore reduced timing margin.

Variation in the rise and fall time of the inverter buffer is highly dependent on the current through each device. The current of a MOS device in the velocity saturated regime is seen in Figure 2.9. Notice that the current is highly dependent on the threshold voltage. From Pelgrom's model [33], the one sigma mismatch of the

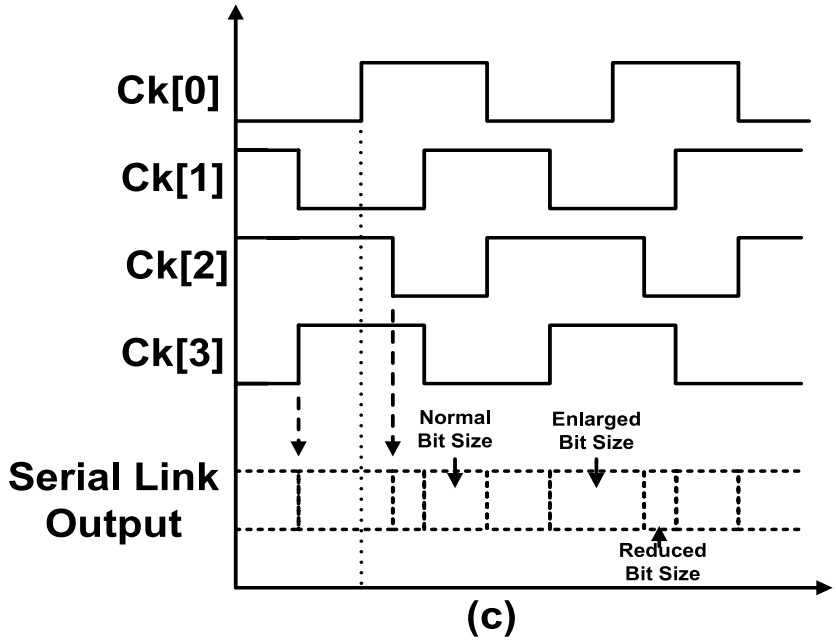

Figure 2.8: 4:1 Transmitter Multiplexing with Static Phase Offset

Transistor Current

$$I_D = \frac{\mu_N C_{OX}}{2[1+\mu_N C_{OX}\frac{W}{L}R_{SX}(V_{GS}-V_T)]}\frac{W}{L}(V_{GS}-V_T)^2$$

Threshold Mismatch

$$\sigma_{\Delta V_T} = \frac{A_{VT}}{\sqrt{WL}}$$

Figure 2.9: Inverter Equations

threshold voltage of a MOS is inversely proportional to the square root of the product of a device's width and length. This variation in the threshold voltage will hence cause variability in the transistor current, resulting in variation in the rise and fall time of the buffer, causing systematic static phase offset. This transistor mismatch problem worsens in future CMOS processes, as other non-ideal effects such as gate length variation will also cause increased static phase offset. [1]

The problem of clock buffer static phase offset is worsened if several clock buffers are used to increase the fanout from the PLL output to the actual location the clock is used for multiplexing. Figure 2.10 shows a plot of the increase in static phase offset as

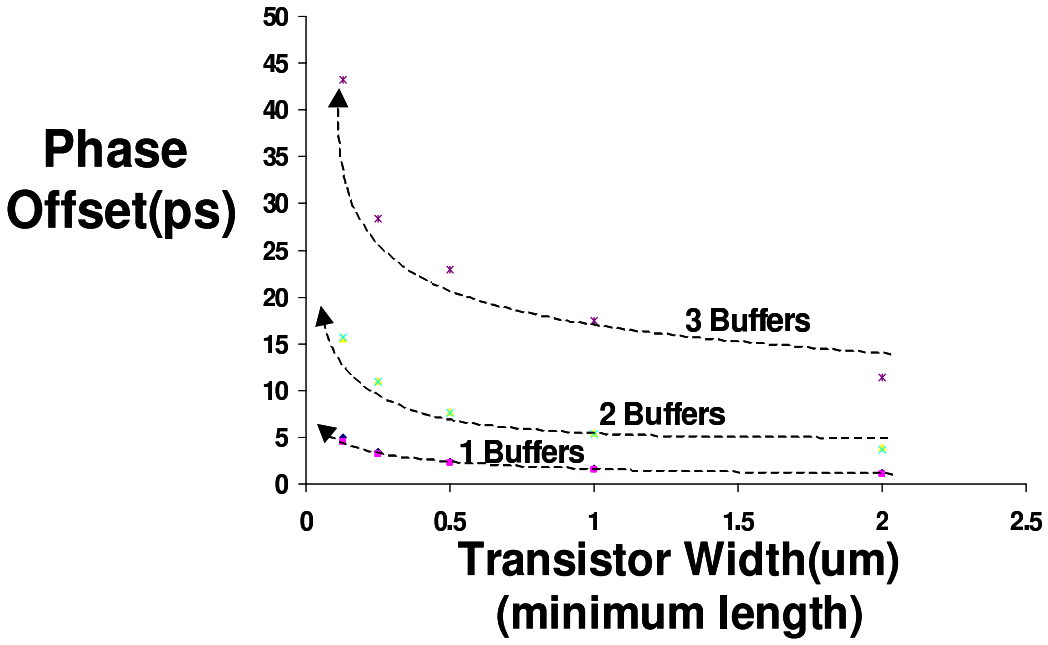

Figure 2.10: Static Phase Offset for Increasing Numbers of Clock Buffers

the number of clock buffers is increased from one to three. These have been simulated with minimum length devices in a 0.13um standard CMOS process, with the worst case delay mismatch causing the worst case static phase offset. From Pelgrom's model, we can increase either the length or width of the devices to improve matching but at the expense of higher power dissipation, which is not an ideal tradeoff.

Exacerbating static phase error is capacitive mismatch. Since the delay through a device depends on both the current dissipation as well as the capacitance of a particular node, the delay can be modulated by capacitive mismatches – for example, between nodes B[0], B[1], through B[n] in Figure 2.4 [15]. As the multiplexer ratio increases, the wiring and layout of the multiple clock phases becomes extremely difficult to route symmetrically.[3] As the number of clock phases used for multiplexing increases, it becomes increasingly difficult to match the layout and capacitance of this clock wiring from the PLL output all the way to the multiplexer input.

The previous analysis pertains only to inverter based buffering and not to CML

[3]One experimental measurement of this result is in Figure 7.6 in Chapter7: Experimental Results, where eight time-interleaved 2.5GHz clock phases are measured.

type buffering. For the most part, differential CML buffers do not suffer from threshold voltage mismatch since the input clock is differential in nature and the rise/fall times does not depend on the exact values of the threshold voltage, but on the ability of the differential pair to switch current when the zero crossing occurs. Simulations show that if the threshold voltage mismatch in the differential pair is 20mV and the differential signal swing is 1V, the static phase offset that results is 1/50 of a bit cycle, or 0.5ps. Simulations in SPICE shows that over a range of differential pair offset mismatch, the worst case phase offset is still less than 1.2ps. Still, other process mismatches may arise in the form of resistive or capacitive loading mismatch, such as in the load of a CML buffer – any differential RC variation between the two loads will cause a timing uncertainty. This is largely mitigated because the load resistance is normally very small (less than 100 Ohms), such that the width and size of these load resistors use large silicon area, mitigating layout variations. Secondly, the delay of the CML gate is based on the delay between the two output nodes, and thus any single-ended mismatches on one edge (i.e. the rising edge) will again occur on the other edge (i.e. the falling edge), resulting in little differential delay mismatch.

2.5.1 Phase Offset Calibration

Since static phase offset is a DC effect with a very small time variance, this type of timing uncertainty can be calibrated away with feedback loops. There have been many possible designs that have incorporated DC mismatch compensation [51, 14]. A more recent example was shown by Lee, Weinlader, and Yang [5], where the residual quantization phase error was around 8ps. Such methods show promise in reducing phase offsets to several picoseconds or less. However, these phase correction architectures are usually limited by the measurement circuitry as well as the phase corrector resolution. Such limitations may also be worsened by process variations and mismatches that degrade in deep submicron CMOS.

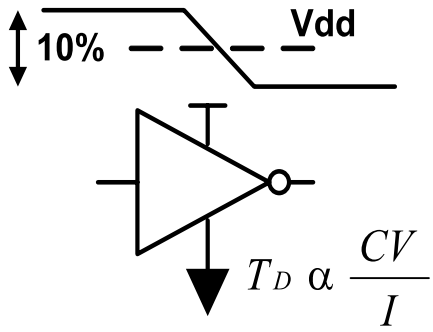

Figure 2.11: Timing Variation due to Inverter Power Supply Fluctuation

2.6 Power Supply Sensitivity

Not only do clock buffers between the PLL and the transmitter multiplexer cause potential DC mismatch problems resulting in static phase offset, but such clock buffers also present susceptibility to power supply noise. Power supply noise exists when large digital noise transients cause the power supply voltage to fluctuate [32].

Figure 2.11 shows a simple inverter where the propagation delay is to first order, linearly proportional to the supply voltage and inversely proportional to the current through the NMOS/PMOS transistors. The transistor current is dependent on the gate over-drive voltage (Vgs - Vt), which continues to decrease with CMOS scaling. While the supply voltage scaling has leveled off in deep submicron CMOS, the transistor threshold voltage has increased in order to reduce overall transistor leakage. This reduction in future gate-overdrive is an important concern due to increased sensitivity to power supply variation.

Figure 2.12 plots the power supply sensitivity across a few CMOS generations for an inverter with a 10 percent change in the power supply. Note that this is not a dynamic variation in the power supply but a 10 percent reduction in supply voltage, resulting in the subsequent measured phase change. The Y-axis plots the jitter (ps) as a proportion to the FO4 delay (ps) of the particular technology.[4]

[4]It should be noted that while the proportional jitter may be getting worse, the absolute jitter in ps for a FO4 has decreased. For example, faster signal transition (rise/fall) times results in less susceptibility for power supply variation to affect the buffer delay.

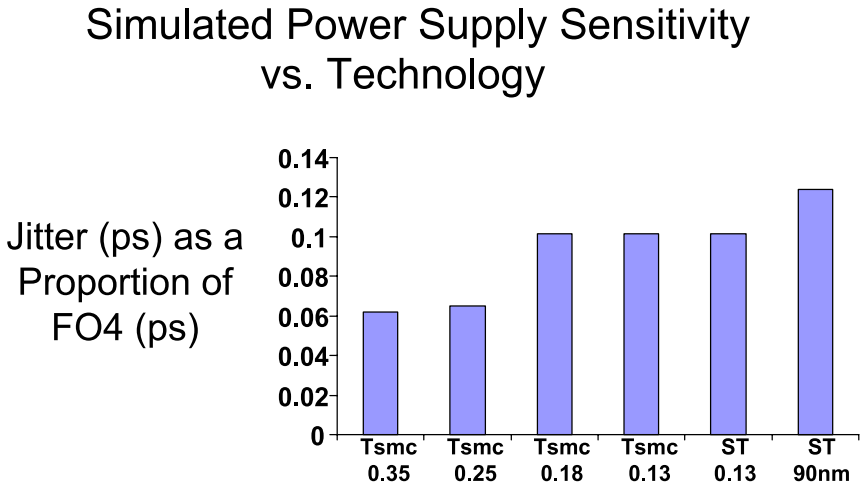

Figure 2.12: Power Supply Jitter Sensitivity Versus Process Scaling

2.6.1 Supply Regulation for Clock Buffers

One possible solution to the power supply sensitivity issue is to use supply-regulated clock buffers [50, 21]. Essentially, as these buffers provide a high power supply rejection ratio (PSRR), the regulated supply remains relatively clean rejecting much of the high-frequency noise. Supply regulation also allows the use of reducing the power supply voltage in different locations based upon required performance and clock frequency, thereby optimizing power consumption. However, there exist several difficulties in building such regulators. One is that with power supply scaling, the headroom for linear regulators decreases, making them difficult to build in future technologies. This usually results in building the regulator with thick oxide devices and the higher supply voltage (i.e. 1.8V or 2.5V), thereby reducing the PSRR and reducing the regulator efficiency. Second, if the regulated supply is less than the typical digital supply voltage, the clock buffer voltage swing will be less than full rail, resulting in negative implications on the speed/bandwidth of the transmitter multiplexing, receiver demultiplexing stages.

2.6.2 CML Buffer Power Supply Sensitivity

While the inverter is quite susceptible to power supply variation, the CML buffer does have very good supply rejection. Due to its differential nature, its delay depends on the current through the tail device, not on the absolute value of the power supply. As long as the compression in the power supply does not vary the current through the tail device and does not saturate the NMOS differential pair, the delay through the differential pair does not change significantly. One paper [24] suggests that substrate noise of both the tail node and the differential pair may change the current modulation of the device–for example, in a ring oscillator with many stages of CML buffers. However, in SPICE simulation, when used as a clock buffer, both substrate noise and supply noise effects contribute less than 1.5 ps of dynamic phase variation.

2.7 Serial Link Transmitter Example

As previously mentioned, process mismatch and power supply sensitivity can introduce timing uncertainty in the clock distribution network that follows clock synthesis. This issue can be elucidated with an example of the clocking network of a conventional serial link. Figure 2.13 shows an example of a realistic, conventional serial link transmitter, either at 4Gb/s (0.25um CMOS)[2] or at 16Gb/s (90nm CMOS)[8]. First, a multi-phase PLL creates four time-interleaved clock phases. Each phase passes through: 1) local inverter for increasing fanout; 2) low-to-high converter, since the PLL usually does not swing full rail-to-rail; 3) local buffering to again increase fanout to the rail-to-rail regime; 4) larger buffering to drive a large capacitance Cline, which is about 80fF of clock loading that drives the 100um of metal wiring distance from the VCO to the transmitter; 5) dynamic NAND structure for data multiplexing; 6) local buffer to finally drive the multiplexer; 7) input multiplexer that finally drives the final output stage. Along this whole path, transistor mismatch as well as capacitance mismatch between the four parallel clock signals can result in static phase offset and reduction in eye width. As well, each buffer also introduces susceptibility to power supply variation. Interpolator structures used to adjust DC static phase

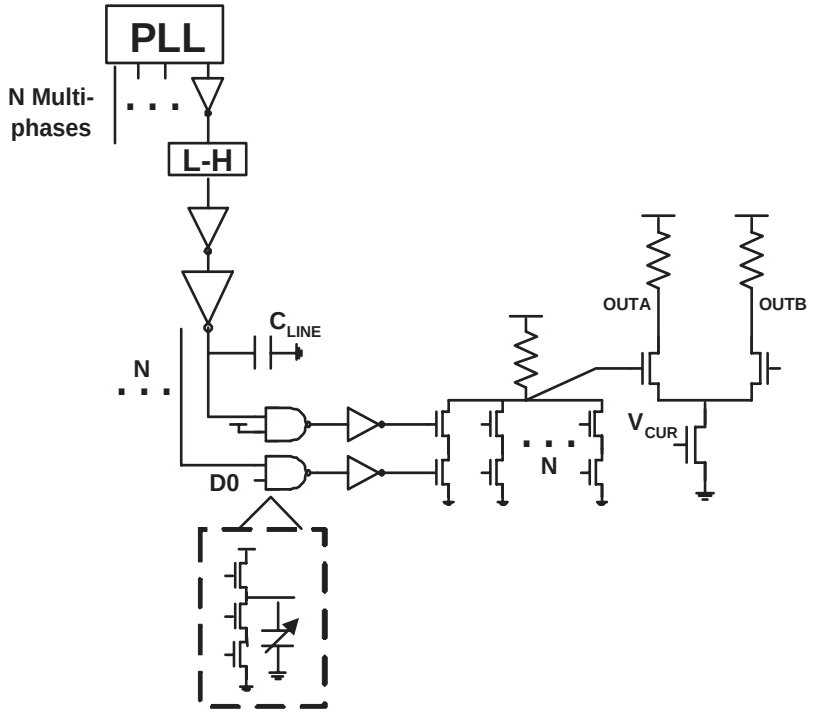

Figure 2.13: Example of Clock Network for a Conventional 4:1 Transmitter Multiplexer

offset, such as a current summing tri-state buffer [12], also introduce more sources of power supply sensitivity. As can be seen, the issue of timing uncertainty – both the static phase offset and the power supply induced jitter – are difficult and critical problems that need to be addressed as serial links approach data rates 20Gb/s and above.

Chapter 3

New Architecture for Serial Links

3.1 Overview

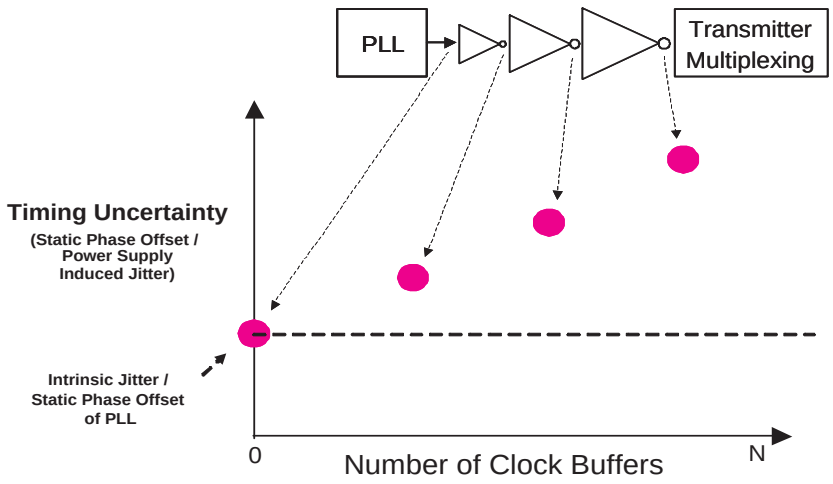

Figure 3.1: Diagram of the Increase in Clock Buffers and its Correlation with Increased Timing Uncertainty

Chapter 2 showed how clock buffering in conventional, multi-phase serial link architectures increases timing uncertainty. Figure 3.1 shows a qualitative graph, plotting the number of post-PLL clock buffers used versus the increase in timing uncertainty. If the clock is taken directly from the phase locked loop, we will find the minimum static phase offset and power supply induced jitter–the residual uncertainty

of the synthesized PLL clock. As this clock is passed through several stages of buffering, each buffer stage introduces an increase in timing uncertainty, both from static mismatches as well as high frequency jitter from power supply noise.

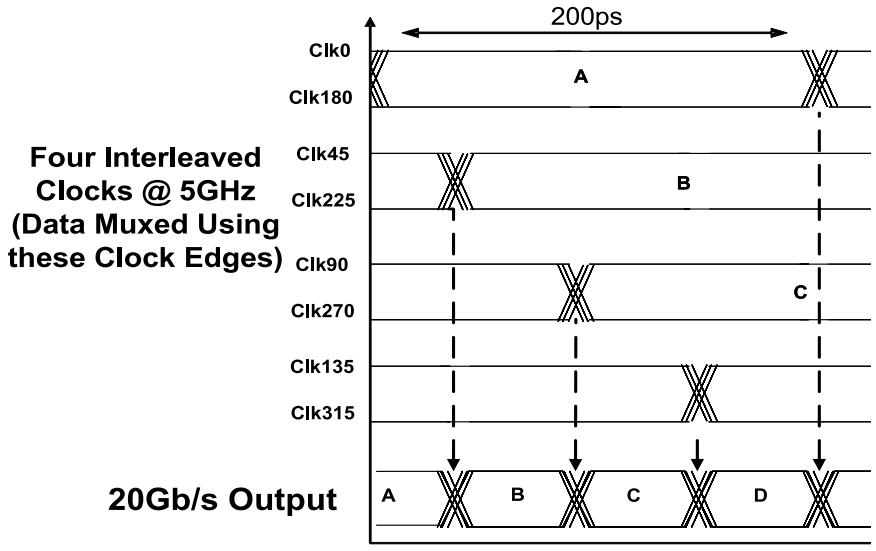

Figure 3.2: Timing Waveforms of a Conventional Low to High Speed Transmitter Multiplexing

Ultimately, if these "dirty" clocks are used to do the final transmitter multiplexing, as in Figure 3.2, the timing uncertainty of these clocks will be passed to the output pin and thus cause timing uncertainty in the up-multiplexed output data.

If the timing uncertainty increases after passing the clock through noisy buffers, it would seem that the critical bottleneck are the clock buffers. A solution, then, is to directly drive the transmitter multiplexer from the PLL as in Figure 3.3.[1] Figure 3.4 shows how such a design might be achieved. First, two 10Gb/s data streams are created at 10Gbs but offset in phase by 50ps. This is typically done using clocks created from noisy clock buffer trees, resulting in timing uncertainty of these 10Gb/s data edges. However, we then create a two phase, complementary 10GHz LC oscillator clock waveform that directly drives the final 2:1 output multiplexer with no digital or CML buffers between the oscillator and the final multiplexer. As will be seen

[1]Note that the figures display clock buffering that drives the transmitter multiplexer. A similar clock distribution also occurs for a receiver demultiplexer datapath.

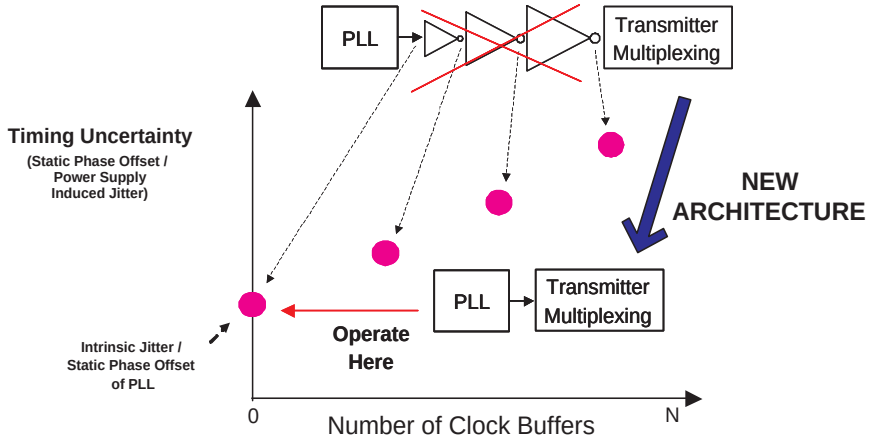

Figure 3.3: Diagram of the New Architecture to Directly Drive the Output Multiplexer with the PLL

in this chapter, because the LC clock has many advantages in regards to low jitter performance, using the LC clock as the selector for the final 2:1 output multiplexer results in a clean output as compared to the conventional architecture. A number of advantages (and possible disadvantages) arise when the VCO/PLL directly drives the transmitter multiplexer. As there are no post-PLL buffers, there are no newly added sources of power supply induced jitter as well as static phase offset due to process/parasitic mismatches. As well, since the PLL directly drives the capacitance of the multiplexer and the metal wiring capacitance, the power consumed is substantially less as the load is resonated into the VCO. Therefore, there are no immediate clock buffers to dissipate wasted energy. While there are some advantages with this direct modulation method, there are some key problems that need to be addressed as well. This chapter will describe some of the key advantages/disadvantages of this new method.

3.2 Oscillator Design

The 2:1 multiplexing and 1:2 demultiplexing architectures require the use of a 10GHz complementary clock. The critical aspects in the design of this clock are low power supply sensitivity, low intrinsic random noise (i.e. device noise, 1/f noise), low static

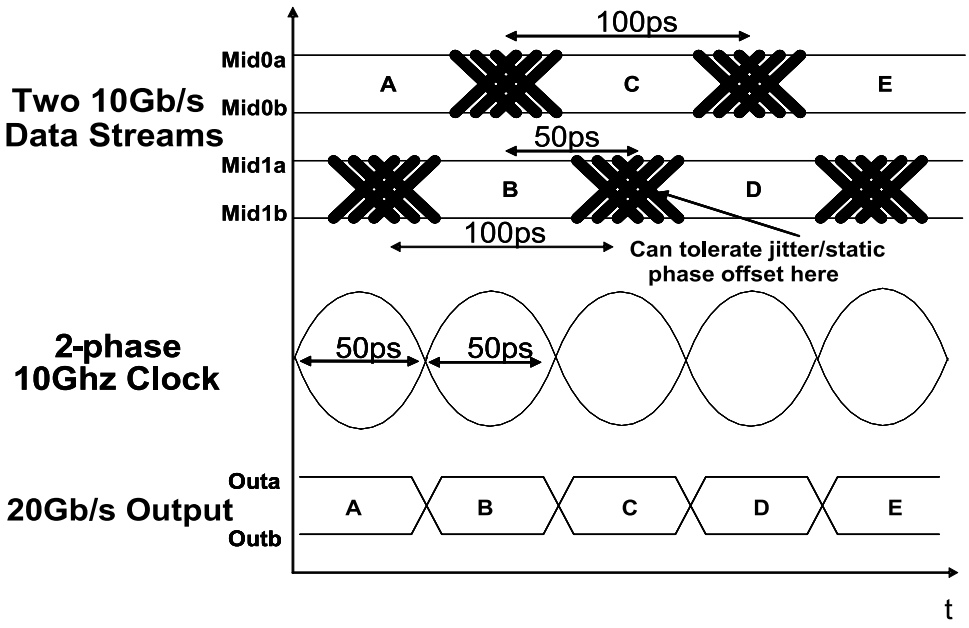

Figure 3.4: Timing Waveforms of the New Architecture Using the LC Oscillator for High Speed Transmitter Multiplexing

phase offset, and low power dissipation.

Two common oscillators are the inverter-based ring oscillator [45] as well as a source-coupled, load-based delay stage [18]. The fastest voltage controlled oscillator using inverter elements is a three stage ring oscillator, where the oscillation frequency is 1/(6*Td), and Td is the delay through one stage. The FO2 delay in a 0.13um CMOS technology is above 20ps, meaning that the fastest oscillation frequency is 8.8GHz. Delay control mechanisms, such as power supply regulation control or variable capacitance loading, will additionally reduce the maximum oscillation frequency. If the static phase offset and high intrinsic phase noise due to device noise are also considered, an inverter-based oscillator for this architecture will be difficult to implement. Oscillators using source-coupled, load-based delay can achieve much higher oscillation frequency than inverter-based oscillators, due to their lower voltage swing and hence, higher bandwidth. However, one major drawback is the high power consumption due to the large current required for adequate output voltage swing in each delay element. Finally, ring oscillators, either inverter-based or resistively-loaded CML based, have

a very large phase noise [10] and may not be suitable for use in a 20 Gb/s serial link where timing precision/jitter is critical.

3.3 LC Oscillator

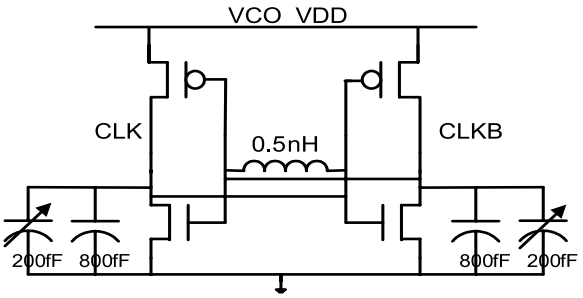

Figure 3.5: LC Oscillator Description

In this proposed thesis architecture, a LC oscillator is implemented. One advantage of LC oscillator structures is that they have extremely high oscillation frequencies in deep submicron CMOS [11, 25, 26, 27, 28]. Maximum LC oscillation frequency is not an issue since the inductance can resonate out the load capacitance. As a result, the oscillation frequency can approach the Fmax of a MOS transistor, which continues to improve with CMOS scaling. A second advantage is that due to the relatively high Q values for monolithic inductors(around 10 at 10GHz), power is conserved as the energy is recycled in the passive inductor – for example, the power dissipation is about 5mW in our implemented VCO. Also, since this power is consumed in a complementary fashion between the two nodes, the current dissipation from the supply is relatively constant as opposed to a inverter-based ring oscillator, reducing contributions and susceptibility to power supply noise. Finally, the intrinsic random noise from 1/f noise or device noise is considerably less than ring oscillator VCOs. From the literature references listed above, phase noise characteristics and long term jitter are considerably smaller.

There are two major penalties for using an LC oscillator. The first is the penalty in area. A 0.5nH inductor occupies an area of 120um x 120um, which might use as much as 1/3 of the area allowed for a serial link, assuming the link size is limited by pad

pitch. However, if we consider higher serial link speeds in future CMOS processes, the size of the inductor will reduce in area in order to achieve a higher oscillation frequency. As well, future processes will have many layers of metal (i.e. 0.13um has 8 metal layers) and thick top layer metal (i.e. greater than 1um thick top layer metal), implying that Q's should increase in future process due to less substrate losses and lower series resistance.

The other major problem is the limited tuning range of LC oscillators. As the oscillation frequency depends only on the sqrt(LC) and the inductor is fixed to a particular value, only the capacitance can be used to achieve tuning. For a 0.5nH inductor, the total capacitance loading on each phase is 1pF. There is considerable amount of fixed capacitance such as the capacitance of the negative Gm inverter stages, metal wiring parasitic capacitance, frequency divider input loading, and transmitter output multiplexer loading. Since an inversion-mode varactor has a maximum capacitance contrast ratio of 2:1, with a 200fF varactor capacitance, the tuning range is approximately 12 percent. For serial link applications where backward compatibility is necessary (for example, 1.25Gb/s, 2.5Gb/s, 10Gb/s, and 20Gb/s), an LC oscillator approach cannot achieve a large range of operating frequencies. One way to extend the tuning range of LC oscillators is to incorporate a dual tuning control loop, with one loop being a coarse, digital frequency controlled loop and the other loop being a fine varactor controlled tuning loop. This is typically done using a digitally switched capacitor bank for the coarse tuning. This has been shown to increase tuning range substantially in some PLL schemes [28].

3.4 Static Phase Offset of LC Oscillator

Removal of the post-PLL clock buffers results in the output of the phase locked loop directly driving the inputs to the transmitter multiplexer or receiver demultiplexer. This means that by eliminating the clock buffers, we have removed significant sources of static phase offset and susceptibility to high frequency jitter. Ultimately, two components of timing uncertainty still remain: the residual static phase offset of the PLL and the residual timing jitter of the PLL as a result of high frequency power

supply noise. The next section will discuss the two main sources of static phase offset – transistor mismatch and capacitance mismatch.

3.4.1 Transistor Mismatch

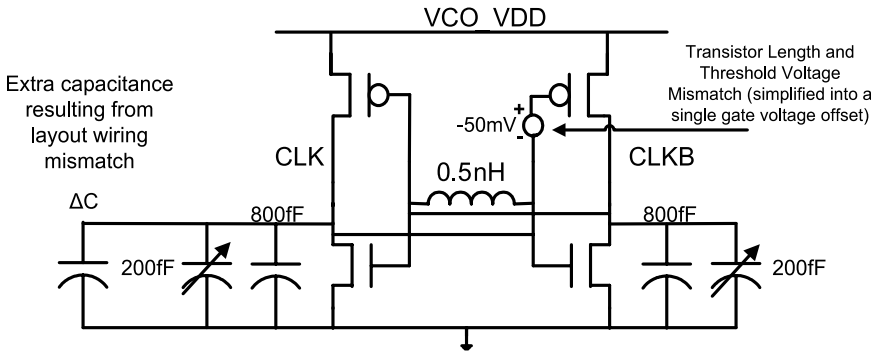

Figure 3.6: Transistor Mismatch, from Both Threshold Voltage Offset and Capacitive Mismatch of the Cross Coupled Pair

Figure 3.6 shows a diagram of how transistor mismatch may be modeled in the 10GHz LC oscillator. In the cross-coupled pair, even if device sizes are made large to reduce statistical dopant fluctuation in the transistor channel, threshold voltage (and in future CMOS technology, gate length variation) will cause mismatch in transistor properties, ultimately causing asymmetry in the rise and fall times of the complementary LC oscillator nodes. However, due to the effect of resonator Q (quality factor), unbalanced energy provided to the tank through the cross-coupled, mismatched transistors is small compared to the energy stored in the resonator. The result is that the static phase mismatch will be mitigated.

To simulate this effect, a deliberate -50mV threshold voltage was placed on one of the PMOS gates, creating an imbalance in the -gm on one side of the differential conductance.[2] This causes one side of the VCO to have a transistor pullup to VDD that is asymmetric with respect to the other. Since this pullup rise time is slower than its complementary node, it causes the other side of the differential pair to have

[2]-50mV is a fictitious upper bound on the total offset of the cross-coupled inverters, assuming the worst case offset condition for all four transistors. Another consideration in future CMOS processes is that transistor length variation may be more significant than threshold voltage mismatch.

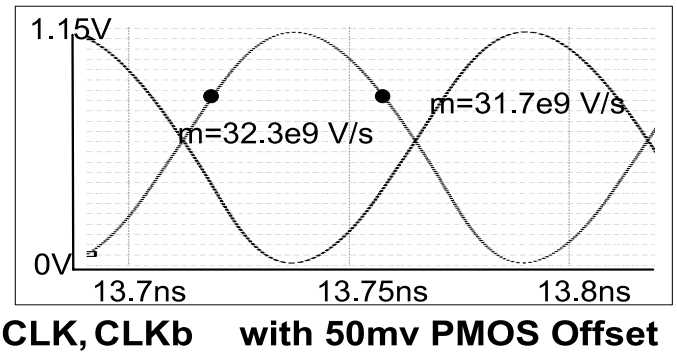

Figure 3.7: Simulated Transistor Mismatch (Threshold Voltage Offset) within LC Oscillator

a slower fall time, as the gate input to the other side of the -gm NMOS pulldown is a slower and smaller voltage. These two competing forces result in a rise/fall time asymmetry between the two complementary nodes, as shown in Figure 3.7. However, since the Q of the oscillator still ensures that the energy must oscillate between both nodes evenly, the static phase offset remains very small – in simulation, less than 1.5 ps.

3.4.2 Capacitive Mismatch

While threshold and transistor mismatch can cause asymmetry in rise/fall times of the complementary nodes, the relatively high Q of the resonator kept the static phase offset small, less than 1.5ps. However, the problem of layout/capacitive mismatch may still exist between the two differential nodes, which may result in an even more significant problem for static phase offset.

Mismatch in capacitance arises when there exists a difference in the capacitive loading between the two differential nodes. For example, the LC oscillator needs to fanout its wiring to relatively long wiring lengths, to both the divider as well as the transmitter multiplexer. Figure 3.8 shows the layout of the complementary nodes from the oscillator to these various places. Notice that even with extreme care of layout, design, ground shielding, etc., it still is extremely difficult to maintain exact capacitive symmetry between the two, whether it is parasitic sidewall wiring

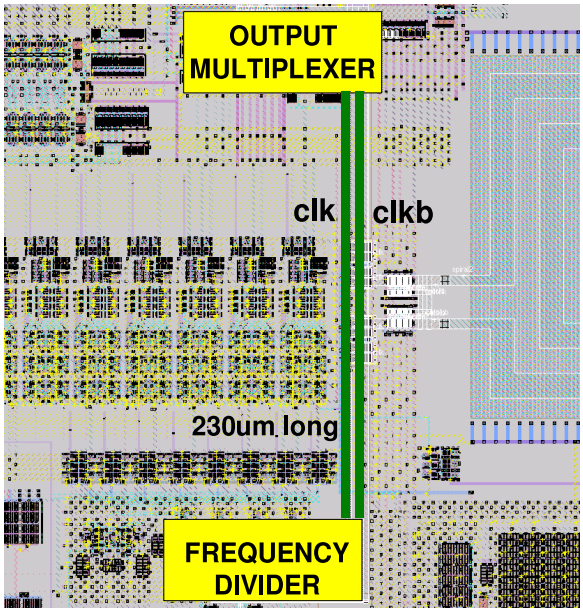

Figure 3.8: Layout of LC Oscillator to Transistor Gate Capacitances of the Frequency Divider and Output Multiplexer

capacitance to DC wires or possibly to signal wires that may have a dynamic coupling effect.

Fundamentally, this capacitive difference can be illustrated by the diagram in Figure 3.6. Here, the inductance can be seen connected between the two capacitive nodes, with the complementary nodes having different capacitive values. Here the current oscillates and recirculates between the two differential nodes. Since the voltage seen on each node is proportional to the charging time of the current into each capacitance, the differential voltage between the two nodes will be different.

Figure 3.9 shows a simple diagram of how the LC oscillator diagram can be transformed into a LC resonator where the two different nodes form a capacitive divider. To first order then, the voltage difference between each node will be exactly linearly proportional to the capacitive difference. In practice, the voltage difference is less than the capacitance mismatch. The reason for this is that devices are biased beyond just the minimum -gm necessary in order to insure oscillation across worst case process corners. Increasing this positive feedback inverter gain mitigates this capacitive-dependent voltage mismatch since the swing of the complementary nodes

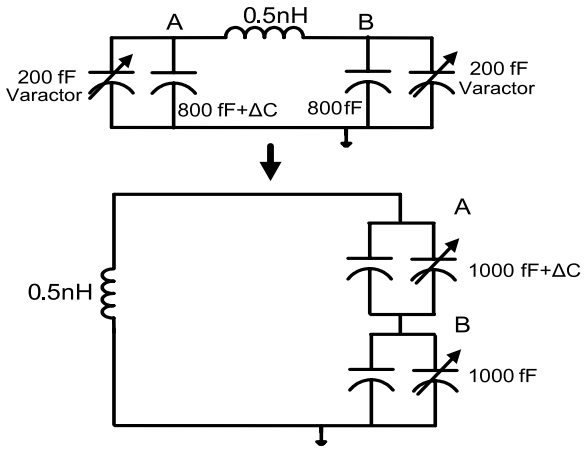

Figure 3.9: Transformation of Resonator into a Capacitive Divider

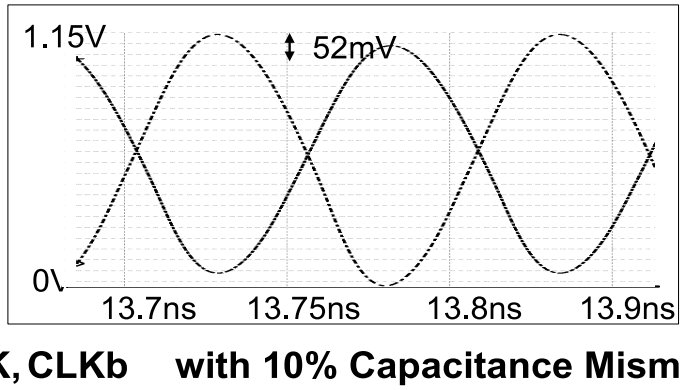

Figure 3.10: Simulated Capacitive Mismatch within LC Oscillator (Including -50mV Threshold Offset)

actually becomes voltage limited by the supply.[3]

Figure 3.10 shows the complementary voltage nodes where there is a 50mV PMOS offset voltage and a 10 percent capacitance mismatch that results in less than 4 percent voltage amplitude mismatch and static phase offset less than 2ps.

In conclusion, this section describes how static phase offset may exist in the LC oscillator when it directly drives the transmitter multiplexer. However, the resulting

[3]As explained by Hajimiri [10], when the oscillator is in the voltage limited regime as opposed the current limited regime, the VCO outputs are saturated close to the supply voltage, regardless of the current source bias. Hence, any capacitance mismatch resulting in voltage mismatch is mitigated by operating the oscillator in the voltage limited regime.

phase offset is simulated to be less than 2ps even with nonideal process and layout mismatches. Hence, the static phase offset problem due to this architecture has largely been mitigated when compared to a conventional time-interleaved clock design.

3.5 Power Supply Sensitivity

While the residual static phase offset of the LC oscillator direct drive method is small, the residual power supply sensitivity of the LC oscillator still exists, and can affect the subsequent timing jitter at the transmitter output. In this architecture, since there are no clock buffers between the point of generation (phase locked loop) and the point of use (transmitter multiplexer, receiver demultiplexer), we need to consider what the residual power supply sensitivity is of a VCO in this direct drive architecture. The power supply sensitivity was compared for a number of VCO topologies: simple ring oscillator [50], source coupled delay line [18], regulated supply delay line [50], inverters using power supply compensation technique [21], and this work that uses a LC oscillator. The figure of merit for comparing these devices is to observe the delay change / supply change (fraction of period changed vs. fraction of supply voltage change). The analysis was done by applying a voltage step to the supply (10 percent), and then measuring the change in the delay relative to its oscillation period. Since this voltage step has a very fast rise time (50ps), the worst case delay change occurs right after the voltage step occurs.

It can be seen from Figure 3.11 that the non-supply regulated LC oscillator topology performs as well as the other designs (Ring Oscillator as designed in [50]; Source Coupled Delay Line in [18]; Regulated Supply DLL in [21]. The minimal power supply sensitivity occurs because the oscillation frequency depends only on the sqrt(LC). When the power supply is perturbed, to first order the tank resonator inductance and capacitance is not altered. Further improvements to supply sensitivity is possible by actively regulating the supply voltage of the LC-VCO, effectively low-pass filtering any noisy digital voltage supply.

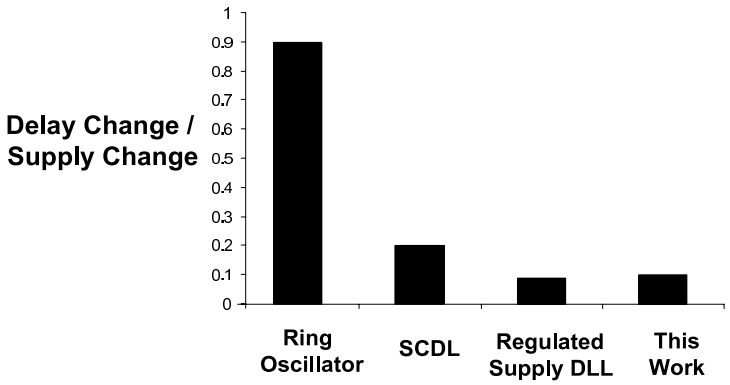

Figure 3.11: Simulated Comparison of Power Supply Sensitivities of Various VCO Topologies

3.6 Possible Negative Parasitic Effects

3.6.1 LC Resonator Q Degradation

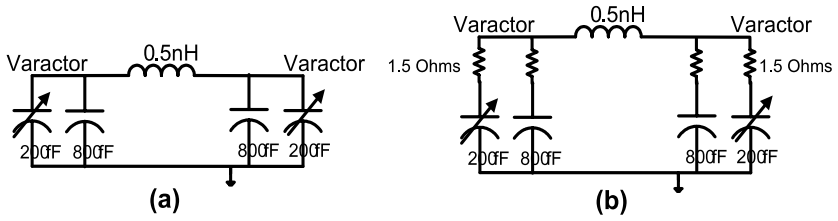

Figure 3.12: Ideal Resonator, and Resonator with Q Reduction Due to Series Varactor Resistance

Since the input capacitance of various blocks (i.e. frequency divider, transmitter multiplexing) is subsumed into the LC tank, any series resistance to these corresponding capacitances can cause significant degradation of the resonator Q. While the inductor Q is easily determined by EM simulation, the capacitance Quality factor is more difficult to estimate due to its distributed nature. Figure 3.12 (a) shows an ideal case for the resonator, where the 0.5nH inductor resonates 800fF of loading capacitance and 200fF of variable, varactor capacitance. In actuality, the Q of each capacitor is reduced due to the series resistance of the gate poly as in Figure 3.12 (b). For example, Figure 3.13 illustrates the layout of each varactor. Though the

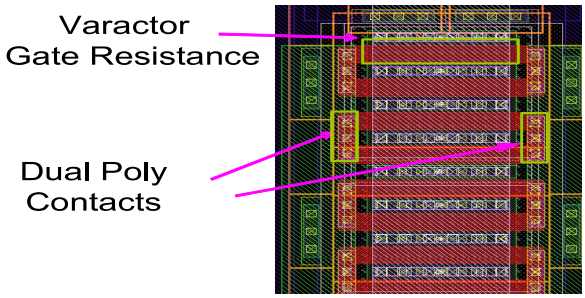

Figure 3.13: Layout of Varactor Illustrating Series Resistance

transistor channel length of each inversion PMOS device is 0.52um long, the width is 2.56um long. Because the poly resistance/square is about 0.3 Ohm/square, the resulting varactor resistance (in series with the varactor capacitance) is 1.5 Ohms.

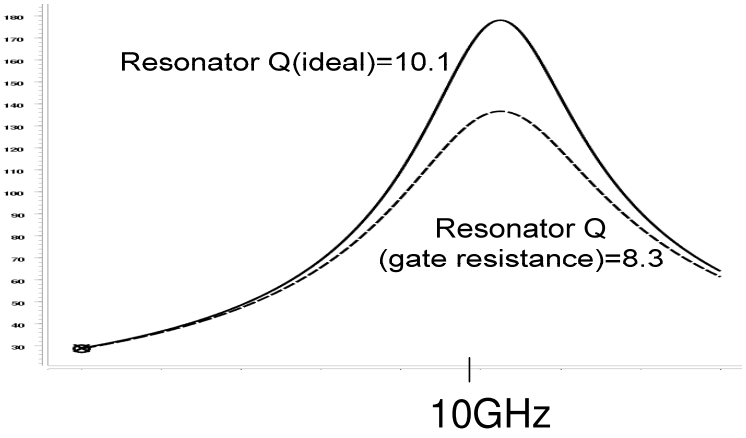

Figure 3.14: Simulation of Quality Factor for an Ideal Resonator and for Parasitic Resistance Resonator

Figure 3.14 shows the simulated resonator Q with and without including the series varactor capacitance – the Q falls from 10.1 to 8.3 due to the series varactor resistance. The -gm gain of the VCO feedback inverters is designed to tolerate this Q degradation. From a layout standpoint, series resistance of the varactor (and any other input capacitance) can be ameliorated by: a) keeping drawn transistor lengths long resulting in a lower number of resistance squares and b) tapping poly contacts on both sides of the transistor gates, effectively reducing the series resistance of the capacitances by a factor of 2.

3.6.2 Phase Lag Due to Distributed Resonator Impedance

One other potential problem of direct resonator modulation of the loading capacitance is the distributed nature of the wiring loading capacitance. This may result in a phase lag/difference between the oscillator core and the actual point of use at the end of the wiring load.

While the inductance of the resonator consists of a localized spiral inductor, the resonance capacitance can be distributed across a long distance, effectively reducing the Q. For example, in Figure 3.8, the wiring route from the LC-VCO to the frequency divider is unusually long (300um), with 20fF of frequency divider gate capacitance lumped at the wire end. If the distributed capacitance has a large lossy capacitance/impedance, the clock oscillations at the input to the frequency divider will be phase lagged compared to the oscillations at the VCO core. ASITIC parasitic extraction showed that these 300um M8 thick metal traces (300um x 4um) at 10GHz can be estimated to be a 0.1nH, 0.2 Ohm, 20-fF transmission line. Transient simulations of such a line show less than 1ps of phase lag, and a 1% degradation in resonator Q due to this lossy trace. Therefore, the distributed nature of the wiring capacitance can still be treated to first order as a single lumped capacitance at one voltage node.

3.7 Jitter Accumulation

Likely the worst possible disadvantage with this direct drive VCO architecture is the possibility for data-dependent kickback. Since no isolation exists between the clean VCO and the data multiplexer/demultiplexer, there may be some coupling mechanisms that can introduce data-dependent fluctuations from the multiplexer directly back into the VCO, ultimately resulting in a decrease in VCO purity and thus increased jitter. To simulate the effect of this jitter degradation, the entire 20Gb/s transmitter output was simulated as in Figure 3.15. 10Gbs data streams are sent to the differential input pairs, to create 20Gbs random data at the output nodes (Outa-20g, Outb-20g). Data dependent jitter can arise from clock coupling between the

differential nodes to the differential pair tail nodes (tail0 and tail1). These nodes, to first order, act as virtual grounds–however, when the nodes Ina0/Inb0 and Ina1/Inb1 switch, the capacitance from gate to source causes a capacitive charge fluctuation across the nodes tail0 and tail1, resulting in a small but noticeable data dependence voltage fluctuation on these tail nodes.

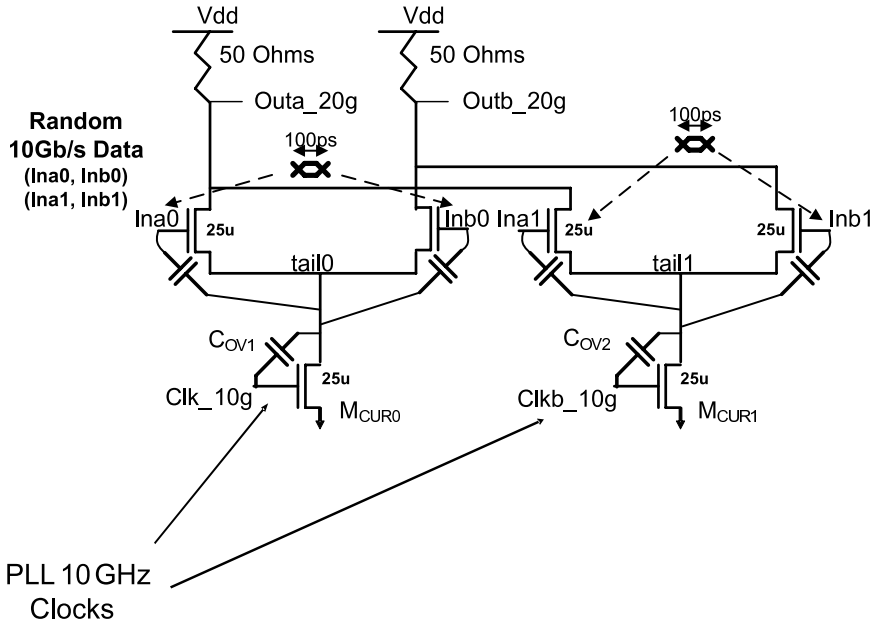

Figure 3.15: Simulation Diagram of Output Driver with Random Data Stream and its Effect on Accumulated Jitter

A long HSPICE transient simulation was done (20,000 clock cycles), where random PRBS data is fed into the differential pair gate nodes, and a 10GHz locked PLL is modulating the tail nodes. Seen in Figure 3.16, the virtual grounds of the current sources have a 50mV voltage fluctuation between no data sent and a PRBS data pattern transmitted. This tail node modulation results in a charge fluctuation that then is kicked back into the LC-VCO through the gate overlap capacitances (Cov1 and Cov2).

Figure 3.17 shows the resulting long term jitter accumulation on the 10GHz PLL with and without data being transmitted. (Effectively, comparing a quiet link and a link with data dependent kickback). Over the 20,000 clock cycles, the pk-pk jitter was measured. While HSPICE cannot accurately calculate long term jitter nor phase

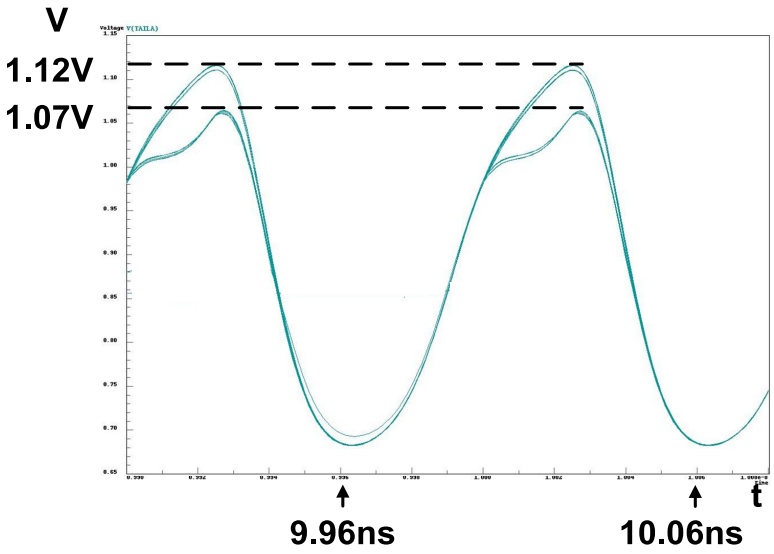

Figure 3.16: Transient Simulation of Tail Nodes of Output Driver Differential Pairs– One Pair with No Input Data Change and the Other with Random Data Input

noise, it can simulate the PLL jitter accumulation due to this data dependent, capacitance kickback fluctuation. The jitter measured with a quiet link is approximately 1.78ps; with transmitted random data the jitter is 3.98ps. Essentially, the jitter has increased, but is on the order of 2ps, which is sufficiently small as compared to a bit time of 50ps.[4]

To first order, we can approximate the maximum kickback jitter into the VCO by calculating the amount of drain-gate overlap capacitance modulating the charge of the LC-VCO. A 40-50mV tail node modulation on a 20fF gate-drain overlap capacitance causes the charge in the entire VCO resonator to be disturbed, modulating the effective resonance capacitance by 1fF (assuming 1V sinusoidal swing). This results in an instantaneous frequency shift, amounting to a maximum, instantaneous period shift in the VCO of 1.2ps. However, this charge injection only occurs during a single data transition (for example, from TX output from '0' to '1'). However, during the opposite TX data transition (from '1' to '0'), the opposite charge injection occurs, causing an instantaneous frequency shift in the opposite direction and bounding the

[4]Using the output stage of the 2nd generation transmitter in Chapter 5 almost completely removes any kickback of data-dependent jitter into the VCO, due to the use of clock isolation from using a separate voltage bias reference. However, this output stage also burns 2x the power

jitter increase. In addition, the PLL loop bandwidth also attempts to reduce any residual, long-term jitter increases.

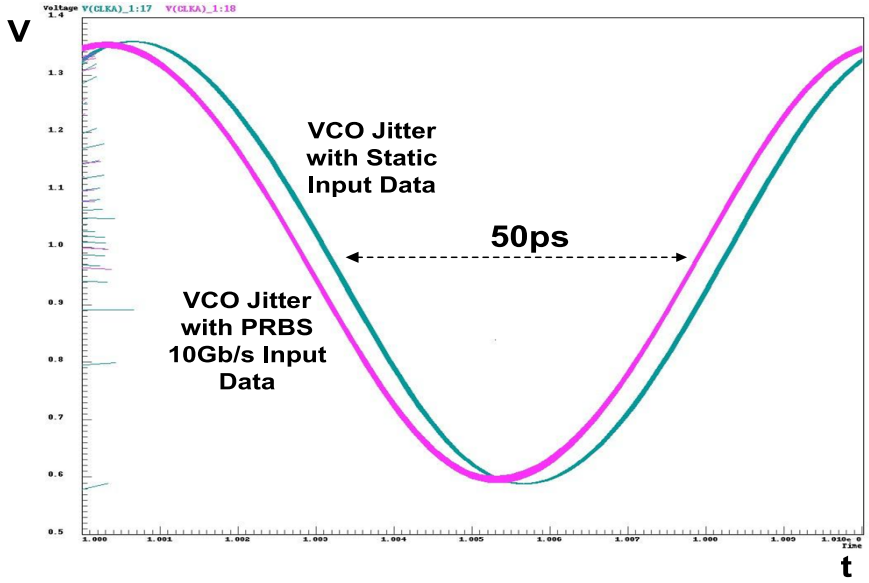

Figure 3.17: Transient Simulation of LC-VCO Long Term Jitter–With and Without Data Transitions

Chapter 4

First Generation Transmitter Architecture

We saw in Chapter 3 how the proposed architecture eliminates the clock buffers between the PLL and the transmitter multiplexer, thereby reducing one of the most significant causes of timing uncertainty in a transmitter output data stream. The LC oscillator, which resonates the actual wiring and loading capacitance of the output multiplexer, may have some advantages in regards to power, static phase offset, and power supply sensitivity. This chapter describes the block and transistor level implementation of this direct drive transmitter architecture.

The transmitter architecture uses a hybrid multiplexing structure to achieve a 20Gb/s output data rate–the transmitter first uses a 4:1, input multiplexed structure to achieve two 10Gb/s data streams, followed by a 2:1, output multiplexed structure to finally multiplex the 10Gb/s data streams to 20Gb/s. Compared to a tree-multiplexing architecture, this hybrid multiplexing relaxes the timing constraints on the feedback clock division and its synchronization with the lower speed data. However, this 4:1 multiplexing stage introduces significant intersymbol interference at the 10Gb/s data rates due to increased capacitive loading. This bandwidth reduction is compensated for by introducing a two-tap FIR filter at 10Gb/s, improving the 10Gb/s and final 20Gb/s eye openings.

4.1 Top Level Architecture Description

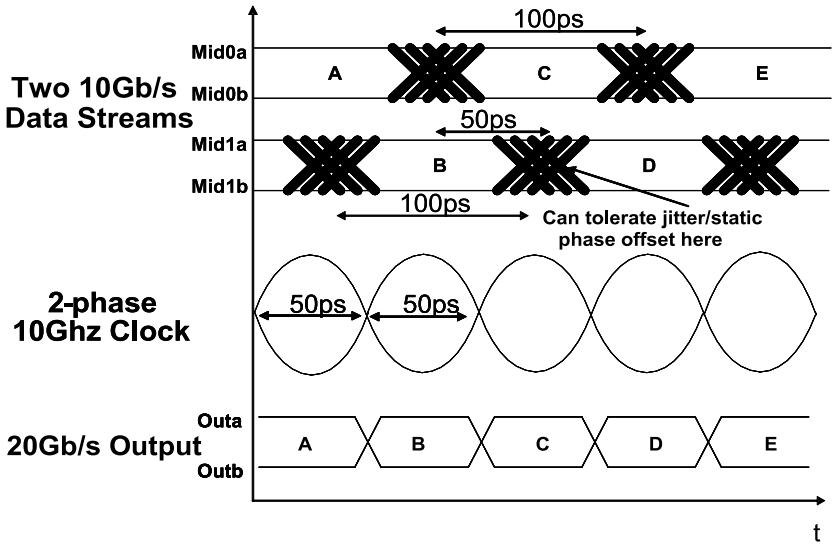

Figure 4.1: First Generation Transmitter Waveform Diagram

Figure 4.1 shows a top level waveform diagram of the first generation transmitter, while Figure 4.2 shows the implementation of the complete transmitter architecture. Two 10Gb/s data streams are created, offset in phase by 50ps – these data streams are created by noisy clock buffers and thus have significant amounts of timing uncertainty along its data transitions. However, a clean, two-phased, complementary LC oscillator clock is used to directly multiplex those two data streams to a 20Gb/s output, avoiding the use of noisy clock buffers. Hence, these LC oscillator phases have much less timing uncertainty, resulting in less jitter at the final 20Gb/s output.

Figure 4.2 shows how these timing waveforms are implemented. A 20-bit, pseudo-random bit sequence (PRBS) generator creates an eight-bit wide data stream at 2.5GHz. These eight data streams are retimed and sent to two sets of 4:1 multiplexers, creating two 10Gb/s data streams offset in phase by 50ps. The two 10Gb/s data streams are retimed to the 10GHz clock domain by two analog sample-and-hold circuits. The final 2:1 output buffer multiplexes the two 10GHz data streams to 20Gb/s. This architecture allows for the final transmitter jitter generation to depend solely on the jitter of the complementary 10GHz clock (CLK and CLKb).

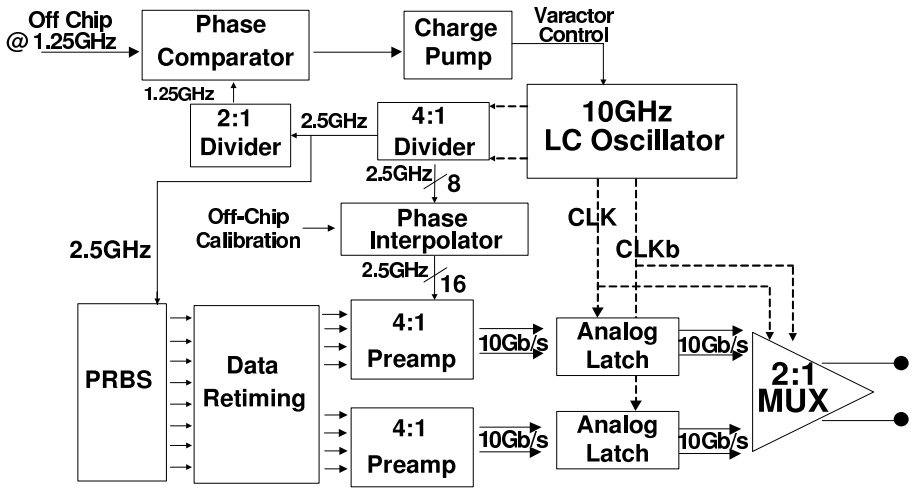

Figure 4.2: First Generation Transmitter Top Level Architecture

Notice that a significant amount of delay exists between the sixteen 2.5GHz clock phases and the complementary 10GHz phases, making it difficult to insure valid setup/hold time of the two 10Gb/s data streams in relation to the 10GHz complementary clocks. Each of the clock phases suffers from large delay variance as the phases pass through low-to-high level conversion, interpolator stages, and fanout buffering before arriving at the 4:1 10Gb/s multiplexer. Hence, a 5-bit DAC phase interpolator is needed to adjust each of the sixteen phases for accurate clock to data alignment. This will be described in more detail later in this chapter.

4.2 Creation of Two 10Gb/s Data Streams

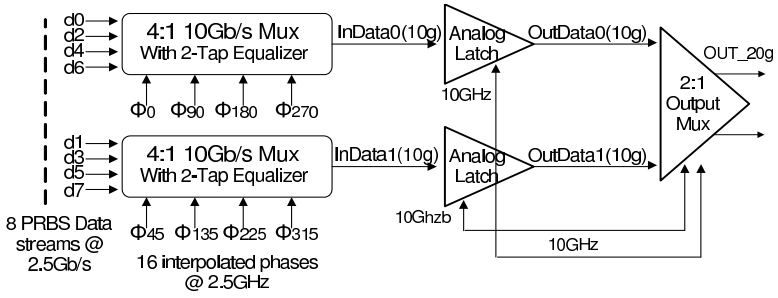

Figure 4.3: Transmitter 10Gb/s Multiplexing Architecture

Seen in Figure 4.3, the eight phases of 2.5GHz PRBS data d[7:0] are 4:1 multiplexed into two two-tap equalized 10Gb/s data streams, using the 16 interpolated clocks from the phase locked loop. These two 10Gb/s data streams are retimed to the clock domain of the 10GHz LC-VCO clock by the use of 10GHz analog latches. Then the two 10Gb/s retimed data streams are then sent to the final 2:1 output multiplexer, achieving 20Gb/s data rate at the differential output. Notice that the 10Gb/s muxes contain two-tap equalization that are used to compensate for significant bandwidth roll-off at node InData[1:0].

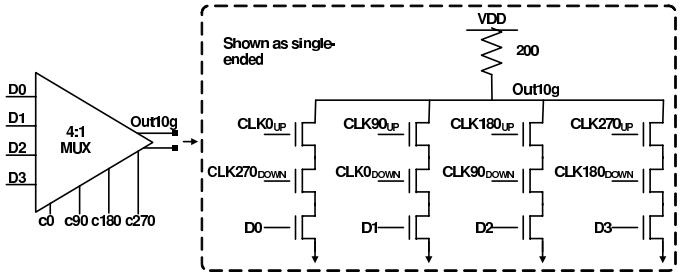

Figure 4.4: 4:1 Transmitter Multiplexing Circuit Design Structure

Each 10Gb/s data stream is created using a 4:1 output multiplexer as seen in Figure 4.4. One multiplexer is designed with four 2.5GHz clocks CLK0, CLK90, CLK180, CLK270, that are offset in phase by 100ps. The other multiplexer is designed with phases 50ps offset from the first multiplexer – CLK45, CLK135, CLK225, CLK315. Notice in this design that the clocks used are both rising and falling edge; for example, CLK0up and CLK0down. Therefore, each multiplexer is designed with eight total clocks. Since each of these phases (rising and falling edge for each phase) is independently controlled by phase interpolators, the design is simplified because each bit duration is defined by its own independent clock. For example, if only four phases are used for each multiplexer, adjusting the rising clock CLK0 will also affect the phase position of falling edge of CLK0. The penalty for making each rising/falling edge independent is burning twice the power and area.

As seen in Figure 4.4 the pulldown path occurs through three NMOS devices, which results in larger device sizes and significant bandwidth degradation. Alternatively, the use of a two-stack NMOS pulldown multiplexer can be achieved using a

dynamic NAND structure [37], which merges the clock phase with the data, as opposed to this implementation of a three-stack NMOS multiplexer[1]. However, such a multiplexing scheme adds two stages of transistors between the clock signal and the multiplexer, allowing for possible static timing mismatch at the multiplexer that cannot be accounted for by the preceding phase interpolator. Using a three-stack NMOS pulldown multiplexer ensures that if the eight interpolated multi-phases are aligned with the two 10GHz complementary clock phases, timing closure can be achieved.

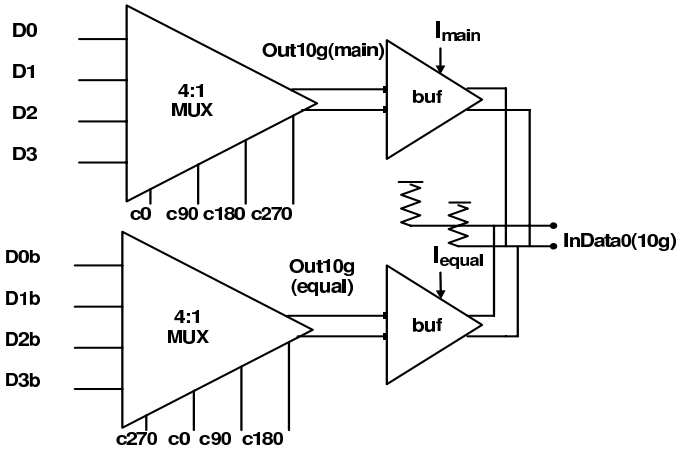

Figure 4.5: 10Gb/s Multiplexer with Current Summing Two-Tap Equalizer

4.2.1 10Gb/s Multiplexer with Current Summing 2-Tap Equalizer

While this implemented three-stage NMOS pulldown multiplexer minimizes the possible static phase offset susceptibility, gain/bandwidth of the multiplexer is negatively affected. Compounding this problem is the fact that aperture time/bandwidth limitations of the following 10GHz analog latch increases intersymbol interference (ISI), translating into ISI at the final 20Gb/s output eye. Two-tap equalization is employed to relax the bandwidth constraints, by equalizing the low pass channel characteristics. Figure 4.5 shows the current summing at the output of the 10Gb/s data streams, using a delayed version of the data to achieve 10Gb/s two-tap equalization. The benefits

[1]As seen in Figure 2.13 of a conventional transmitter multiplexer

of this equalization will be described in the next section.

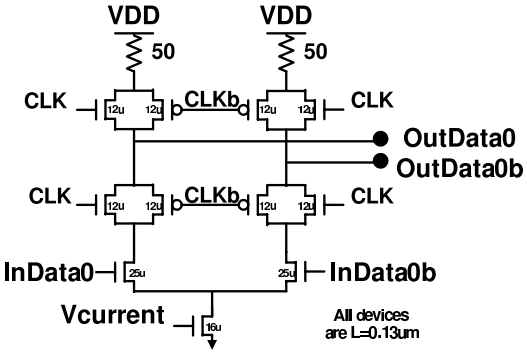

Figure 4.6: 10GHz Analog Latch

Figure 4.6 shows the schematic of the 10GHz analog latch used. As the complementary 10GHz clocks CLKb and CLK fall, the differential signal Indata0 is sampled onto the intermediate nodes Outdata0, Outdata0b. Full pass gates are used to mitigate charge injection loss onto the sampled nodes. Other types of high speed analog latches were deemed less effective. For example, a pass-gate analog sampler is ineffective, as a large(400-600mV) voltage swing is sampled onto node Voutdata0, which will affect the on-resistance of the switches. Another alternative latch is to use a current-mode logic(CML) latch, but its performance is not as robust since charge loss occurs in its differential pair when the output discharges through the resistive load. Full pass gates in this implementation isolate the resistive path from the output, such that when the clock falls, the data is fully sampled onto the differential node Outdata0. Also, full pass gates with the same size NMOS and PMOS gives symmetric clock kickback injection.

While the aperture time for this implemented latch is very fast, one problem that results with this full-pass gate switch is that hysteresis exists at the differential output node. Since only 2-phase 10GHz clocking is used, one phase of clock is used for sampling and the other for holding; no clock phase exists to reset the sampled nodes and remove the hysteresis, as typically done with conventional multi-phase sample-hold switches. This hysteresis can be observed as bandwidth reduction or intersymbol interference (ISI), as the previous sampled bit can not be discharged from the internal nodes and hence, residual memory exists for previous sampled bits.

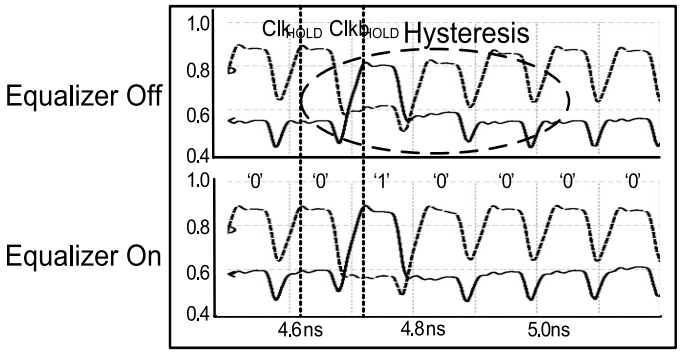

Figure 4.7: Pre-Emphasis Waveforms for 10GHz Analog Latch Hysteresis

As seen from previous research work [50], ISI can be compensated for using analog equalization – in this implementation using a preceding 10Gb/s two-tap equalization. Figure 4.7 illustrates the sampled differential analog latch voltage with and without pre-emphasis. For the '1' symbol being latched, clearly significant differential signal memory is seen for three bit times. With pre-emphasis of the 10Gb/s data streams on, the differential voltage observed is constant for the '1' symbol.

Figure 4.8 shows the utility of pre-emphasis of the analog latch. Note that the dotted line is for voltage waveforms with equalization, solid line for waveforms without equalization. Figure 4.8 (Panel 4) shows the analog latch input with and without pre-emphasis. The output of the latch in Figure 4.8 (Panel 7) shows the latched analog 10Gb/s data streams, with significant ISI during the fast edge transitions. 4.8 (Panel 12) shows the 20Gb/s transmitter output with and without equalization of the analog latches. A significant amount of eye closure (approximately %15) is observed in the 20Gb/s output eye for the non-equalized latch(solid line) as compared to the equalized case(dotted line).

4.3 Final 2:1 20Gb/s Output Multiplexer

The retimed 10Gb/s data streams are multiplexed by the final 2:1 output driver, as shown in Figure 4.9 The 2:1 output driver is implemented using source coupled

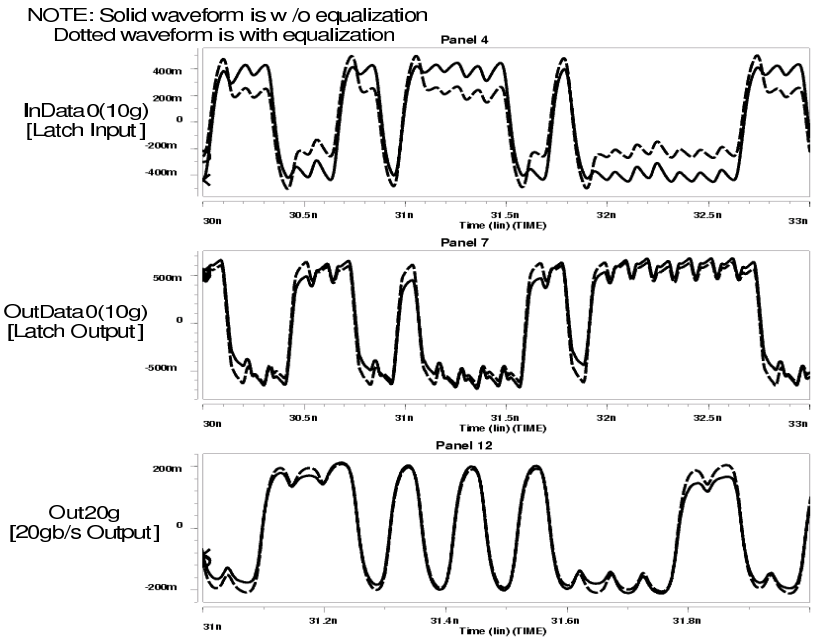

Figure 4.8: Multiplexer Waveforms Illustrating Utility of Internal Bandwidth Equalization

pairs, current summing on both differential pairs through 50 Ohm on-die poly termination resistors. Notice that the possibility of data-dependent kickback can occur, when the output voltage modulation causes the differential pair tail node to also modulate–essentially, the tail node no longer acts as a perfect virtual ground. This data dependent modulation of the tail node will then couple back into the resonator nodes through parasitic gate-drain overlap capacitance, essentially injecting data-dependent noise into the clean LC resonator. As explained in Chapter 3, the effect of data dependent, accumulated VCO jitter is small, due to the small sizes of the parasitic drain-to-gate capacitance of the current source nodes (20fF), versus the large capacitance of the complementary resonator nodes (1pF). Also, the data-dependent modulation of the differential tail nodes is about 50mV, making the amount of data-dependent charge kickback to the resonator small, resulting in a minimal amount of increased jitter (less than 2ps pk-pk).

Figure 4.10 shows a simulation of the single-ended, transmitter output post-layout, with an ideal channel response of over 1,000 bit times. Notice the voltage ripples

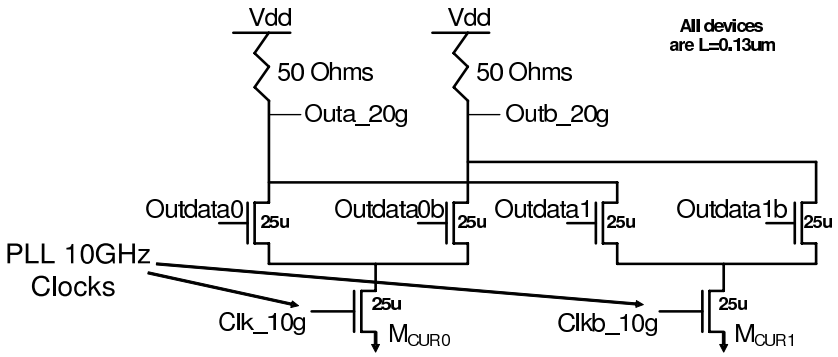

Figure 4.9: Final 20Gbs 2:1 Output Multiplexer

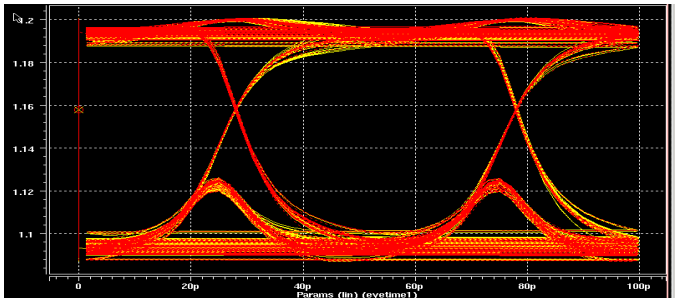

Figure 4.10: 20Gb/s Simulated Eye Diagram with Parasitics

occurring during the zero crossings of the bit symbols (between 1.12V and 1.1V, occurring @ 20GHz repetition frequency). This problem occurs due to non-idealities in the current source nodes, where the inputs are the complementary 10GHz clock nodes. When the complementary phases are at opposite voltage magnitudes, where one clock is at its apex and the other at its nadir, the current steering differential pair works well, as only one of the NMOS current source pulldowns is on, forcing the current to move to the GND return through only one NMOS. However, a problem is observed during the zero crossing of the complementary clocks, as seen in Figure 4.11. If the common mode voltage of both complementary clocks is high during the zero crossing, both NMOS pulldowns will act as current sources, causing the output nodes to modulate into larger differential voltages during this moment. On the other hand, if the common mode level of the complementary clocks is low, during the clock zero crossing moment neither current source is on, causing eye opening reduction at

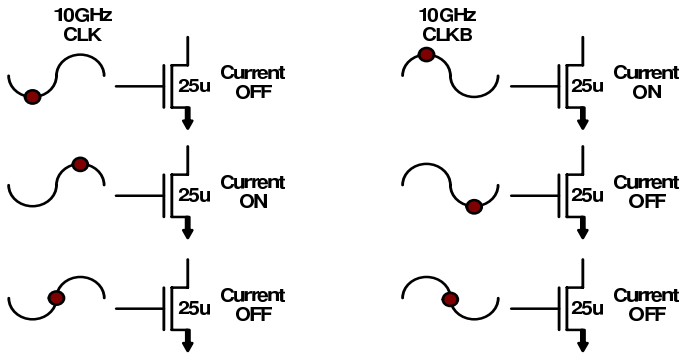

Figure 4.11: Current Source Modulation During Clock Zero Crossing

this transition – the tail node of the differential pairs pulls up, causing the differential outputs to pullup to VDD as there is no other path for the current to flow. In the eye diagram, this voltage modulation magnitude is about 40mV (differential) or 22 percent of the total eye size. However, since this modulation occurs during the zero crossings, this effect on bit error rate is small since these perturbations occur out of phase from the receiver clock that retimes the bits at the center of the eye. While this noise may ring around the transmission line corrupting later bits, this noise is attenuated significantly by the channel loss at this frequency (20GHz), as the noise has to traverse a round-trip delay of the transmission line.

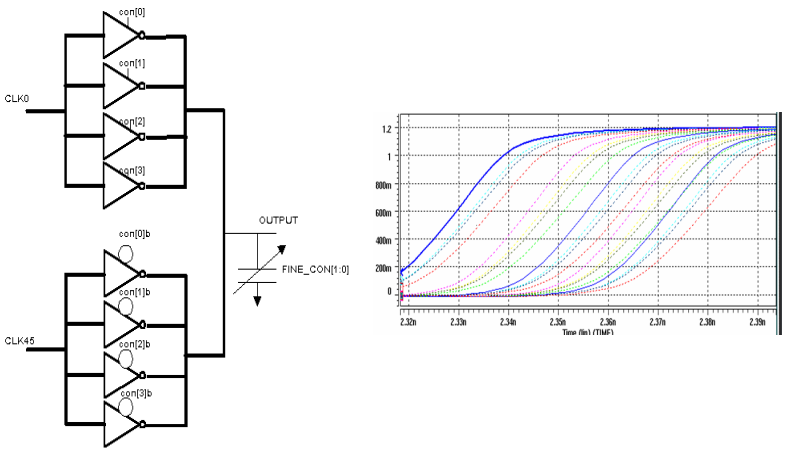

Figure 4.12: Interpolator Schematics and SPICE Simulations

4.4 Clock Alignment

As mentioned previously, a significant amount of delay exists between the sixteen 2.5GHz clock phases and the complementary 10GHz phases, making it difficult to ensure valid setup/hold time of the 10GHz analog latches, with respect to the two 10Gb/s data streams which are sampled by the 10GHz complementary clocks. Each of the 16 clock phases suffers from large delay variance, as the phases pass through low-to-high level conversion, interpolator stages, and fanout buffering before arriving at the 4:1 10Gb/s multiplexer. Hence, a 5-bit DAC phase interpolator is needed to adjust each of the sixteen phases, creating the necessary phase shift for valid setup time of the following 10GHz analog latches (Figure 4.12). Coarse interpolation is achieved using tri-state buffer current summing, while fine interpolation is done through capacitive trimming, resulting in a minimum interpolation step of 7.3ps, covering a total range of 50ps.

While the interpolator allows individual adjustment of each of the 16 phases that create the two 10Gb/s data streams, there is no closed loop feedback to insure adequate setup/hold time of the 10Gb/s data streams with the 10GHz complementary clocks of the retiming latches. Also, static phase mismatch between the 16 different clock phases will cause uneven 10Gb/s data streams into the two complementary 10GHz latches, giving uncertain setup/hold times for the 10GHz analog latches.

Calibration logic to align the 16 interpolated clock phases with the two-phase 10GHz clock is relatively difficult to design, as it adds additional complexity due to the measurement circuits and increased power dissipation resulting from the phase interpolation overhead needed for all 16 phases. Therefore, calibration logic to adjust the 16 clock phases in relation to the 10GHz clocks was not implemented. In the actual implemented prototype, simulation of the clock delay from the PLL to the data multiplexers was used to determine the optimum value of the phase interpolator setting. Then, by observing the data eye opening recovered at the oscilloscope, the phase setting could be shifted incrementally from this default value.

A more exact, on-die approach is to use a statistical eye sweeping approach, similar to a sub-sampling oscilloscope as described by Weinlader in [35]. The eye height and

width of the 10Gb/s data, as well as each of the two 10GHz complementary phases, can be swept by a subsampling clock, and therefore the zero crossing of each can be measured and corrected.

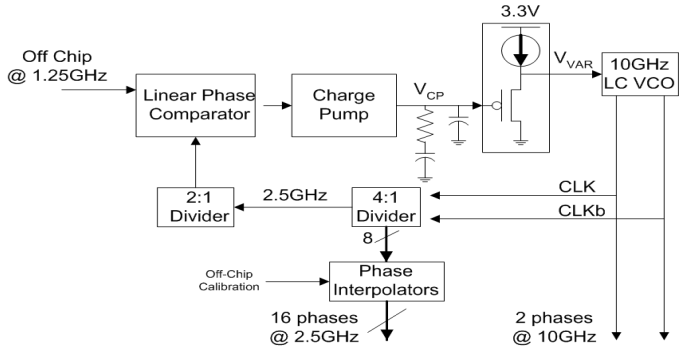

Figure 4.13: First Generation Transmitter Phase-Locked Loop Architecture

4.5 First Generation Transmitter Phase-Locked Loop

Figure 4.13 shows the block diagram of the 10GHz phase-locked loop that multiplies the off-chip reference of 1.25GHz by a factor of eight using the on-chip LC oscillator. The design of the linear phase comparator and charge pump are conventional and are not described here in any great detail.

The first division stage of the 10GHz PLL is a 4:1 division directly from 10GHz to 2.5GHz. This is accomplished by placing four CML-type latches in a ring orientation similar to Figure 4.14. To improve bandwidth, the NMOS current source transistors are removed reducing the number number of stacked transistors between the supply rails.

One non-conventional modification of this PLL design is the addition of a source follower level converter following the charge pump. The purpose is to reduces the VCO gain relative to signal swing as the charge pump output is converted from 1.2V to 3.3V voltage domain. This technique also allows for a larger tuning range – without conversion to a higher voltage regime, the charge pump output can only swing from approximately 0.2V - 1.0V, resulting in a smaller tuning range of the inversion-mode

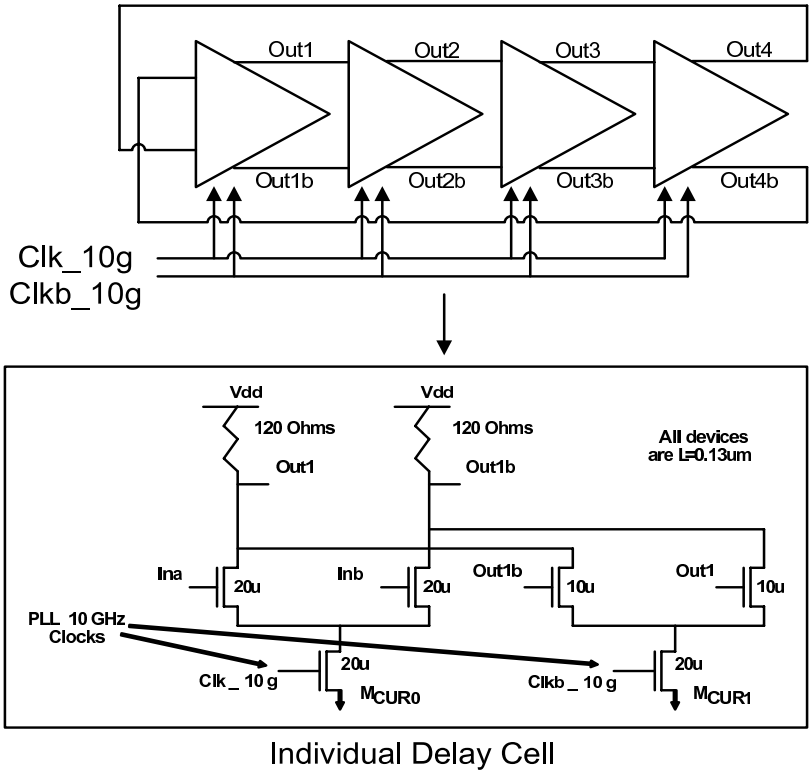

Figure 4.14: 10GHz to 2.5GHz Frequency Divider

varactors. However, after conversion to the 3.3V supply, the DC swing of the charge pump can now approach 0.6V-2.5V (depending on the current bias of the source follower). As the swing of the LC-VCO sinusoidal voltages is between 0 and 1.2V, one side of the varactor voltage midpoint (the node connected to the VCO clock) is approximately DC biased at 0.6V. If the charge pump can only swing between 0.2V and 1.0V, the largest DC magitude on the varactor is -0.4V to 0.4V. By converting to the 3.3V supply, the DC range on the source follower output is now 0.6V to 2.5V, which switches the varactor through its large capacitance switching region (0V to 1.9V).

Secondly, the level converter can also act as a voltage buffer that isolates the critical charge pump voltage from the varactors and the VCO. Since the VCO complementary nodes have a voltage swing close to VDD (around 1V), these large signal voltage swings can capacitively couple directly through the two varactors, causing a

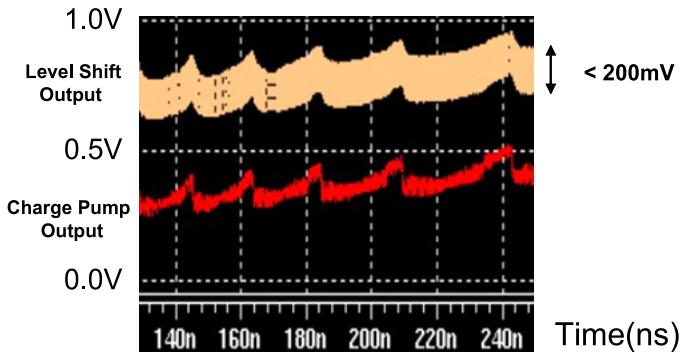

Figure 4.15: SPICE Simulation of Charge Pump Output and Level Shift Output

ripple to occur on the charge pump at a frequency of 20GHz. Essentially, the magnitude of the ripple is proportional to the size of the varactors relative to the size of the charge pump low pass filter – in simulation, this ripple is on the order of 400mV. This type of ripple may have an effect on phase noise since the varactors now see a DC voltage that oscillates with a 400mV noise floor, modulating the varactor DC bias, thereby changing the varactor capacitance value, resulting in a change in the LC resonance frequency. Addition of the level converter reduces the varactor DC modulation due to this large signal kickback by 2x, less than 200mV, as can be seen in Figure 4.15.

Chapter 5

Second Generation Transmitter Architecture

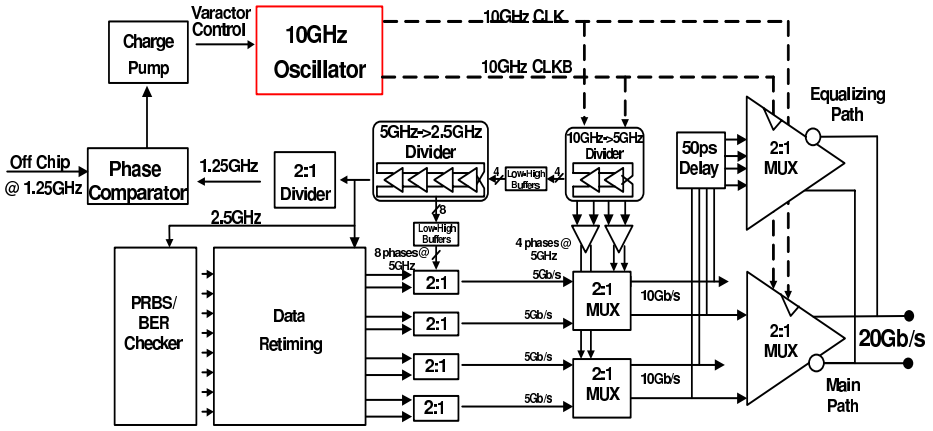

Figure 5.1: Second Generation Transmitter Top Level Architecture

One problem that was seen in our first generation transmitter was that there existed some difficulty in aligning the two 10GHz complementary phases with the two 10Gb/s data streams. Since there exists a significant delay between the location of the sixteen 2.5GHz clock phases (used in the 4:1 multiplexers to create two 10Gb/s data streams) and the 10GHz clock phase edges, a non-trivial feedback loop must be designed that can accurately align these two different clock domains.

To address this synchronization problem, a second generation transmitter was

designed that also exhibits a few other advantages. First, a tree multiplexing architecture is used to multiplex data to the higher transmitter output rate, as seen in Figure 5.1 and 5.2.[1] Essentially, data is up-multiplexed from 2.5Gb/s - 5Gb/s - 10Gbs - 20Gb/s, while the clocks are derived from the PLL divider chain – from the 10GHz LC-VCO clock (two-phase complementary clocks at 10GHz), to the 10GHz clock divider(four phases @ 5GHz), and finally to the 5GHz clock divider(eight phases at 2.5GHz). Notice from the figure that this transmitter does not use explicit latching and retiming of the final 10Gb/s data streams. Instead, the delay through the 10GHz-to-5GHz clock divider/clock buffer is simulated extensively to insure that this 10GHz clock falls within the period for which the data is valid. This simulation of the clock delay paths relative to the data multiplexing is done for every stage of data upmultiplexing.

One other added feature of this new second generation transmitter is the implementation of a two-tap 20Gb/s equalization filter to equalize frequency dependent attenuation of the channel, which is very likely to occur at high data rates. This equalizer is implemented by delaying the two 10Gb/s data streams by 50ps, and then current summing this delayed path with the main 2:1 final output driver. This two-tap equalizer at 20Gb/s improves signal-to-noise ratio (SNR) at the receiver when used in a moderately lossy channel (-6.5dB at 10GHz).

The final major modification of this transmitter is the implementation of a revised output driver. In the previous transmitter, there was no explicit way to control the output voltage swing. Hence, this new transmitter has an explicit current source control to adjust the main current tap and therefore the output voltage swing. This revised output driver also completely isolates the output data from the clock using a parallel differential pair with one node biased to a particular voltage bias. The current commutation occurs when the differential pair inputs compare between the clock voltage and the bias reference voltage, similar to a fT doubler. Since the clock node is isolated from the lower frequency data switching noise, data dependent kickback is reduced significantly. Finally, this new output driver to first order rejects the

[1]Note that we are still directly driving the final output multiplexers directly from the LC-VCO, thus removing any timing uncertainty traditionally caused by clock buffers being used before the transmitter output.

dependence of the output swing on the common mode of the differential clocks– the 20GHz voltage ripple seen in the first generation transmitter can be reduced by optimizing the bias voltage in relation to the clock swing.

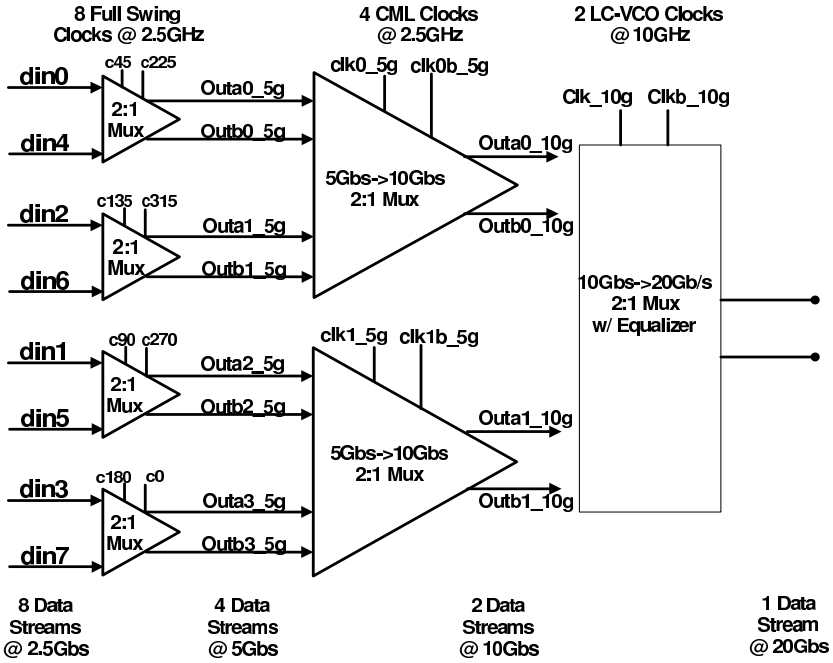

Figure 5.2: Tree Multiplexing from 2.5Gb/s to 20Gb/s

5.1 Transmitter Multiplexing from 2.5Gb/s to 10Gb/s

The 2.5Gb/s to 5Gb/s multiplexer is shown in Figure 5.3(a). DIN0/DIN4 are full-swing digital data coming from the 2.5GHz PRBS generator. clk/clkb are full-swing 2.5GHz clocks created from the PLL, where the 10GHz VCO phases are passed through two divider stages, low-to-high level converters, and finally inverter buffering before being used to up-multiplex data from 2.5Gb/s to 5Gb/s. The output swing of this pseudo-CML multiplexer (outa0_5g) results in about 0.5*VDD swing, essentially a gain of 0.5.

Figure 5.3(b) shows the diagram of the 5Gb/s to 10Gb/s multiplexer. The inputs (Outa0_5g, Outb0_5g and Outa1_5g, Outb1_5g) come from the previous 2.5Gb/s to

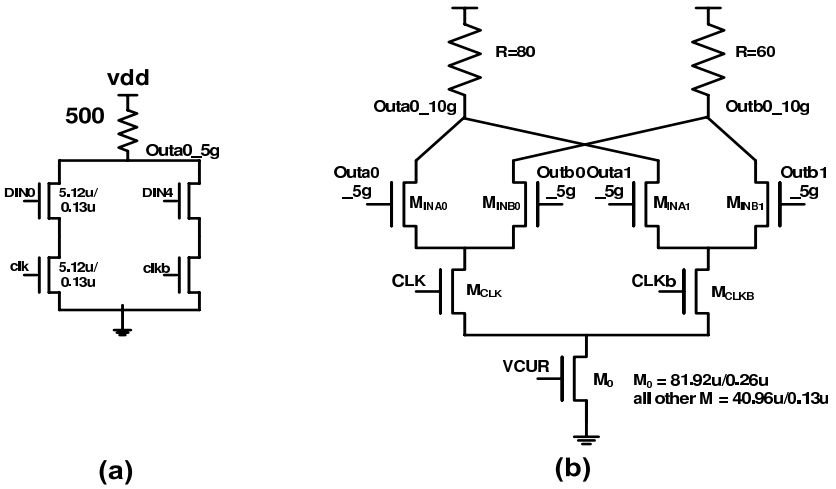

Figure 5.3: (a) 2.5Gb/s to 5Gb/s Data Multiplexer (b) 5Gb/s to 10Gb/s Data Multiplexer

5Gb/s multiplexers. CLK/CLKb signals are CML low swing, 5GHz differential clocks, created from the PLL after the first divider stage, and through one stage of CML buffering to improve the clocks' rise time and swing. Vcur controls the swing on this node, compensating for process variation. The gain is converted from about 0.5Vdd at the input to 0.3Vdd at the output, with an output data rate of 10Gb/s.

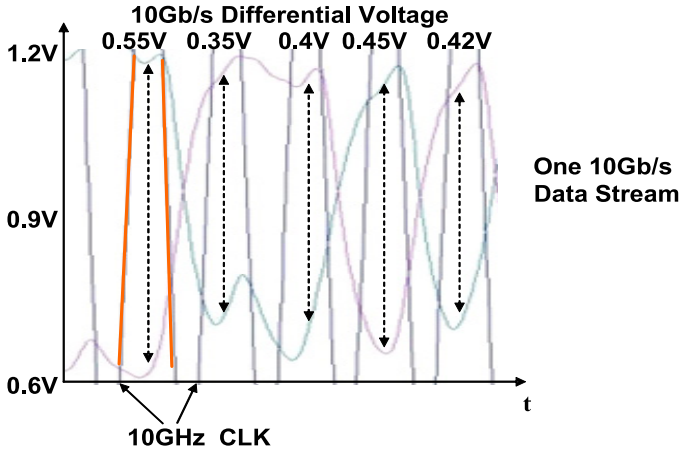

Figure 5.4: Output of 10Gb/s Multiplexer, to be Further Latched by 10GHz Clock

Figure 5.4 shows a SPICE simulation of one differential data stream that has been upconverted to the 10Gb/s data rate (Outa0_10g, Outb0_10g). This 10Gb/s data

stream will then by multiplexed by one phase of the 10GHz clock, resulting in the final 2:1 20Gb/s transmitted output, while the other 10GHz clock phase multiplexes the other 10Gb/s data stream. Notice that the 10Gb/s data eye has quite a large amount of voltage ripple. For example, the 10Gb/s differential output voltage might vary from 0.55V to 0.35V. This is a result of limited bandwidth and ISI occurring at the output resistor node at the 10Gb/s data rate. If the resistor is too small there is not enough gain from this stage to sufficiently drive the next stage 2:1 20Gb/s multiplexer. At the same time, a large resistor results in limited bandwidth which causes causing larger residual ISI of this output, causing further problems during the multiplexing at the next stage. Already, the small voltage difference between 0.35V and 0.55V causes the next stage multiplexer to incompletely switch the output current, causing residual ISI and rise-time degradation in the 20Gb/s output eye.

Hence, while this multiplexing seems relatively straightforward, difficulty occurs in ensuring adequate timing setup/hold margins for the clock distribution and bandwidth of the CML multiplexers. Possible solutions are to improve the bandwidth of the CML gates, either by inductive peaking techniques or voltage limiting techniques, such as shunt-shunt feedback topologies similar to a Cherry-Hooper amplifier [7]. But these techniques require increased design time, larger area, and increased power dissipation.

5.2 Two-Tap Transmitter Equalization Implementation

After the two 10Gb/s data streams are created, they are again multiplexed to a higher data rate by a 20Gb/s 2:1 multiplexer, resulting in the main output driver tap. These two 10Gb/s data streams are also sent through two stages of CML delay, as in Figure 5.5, effectively creating a 50ps analog delay. These two delayed 10Gb/s data streams are finally sent through an identical 2:1 20Gb/s multiplexer with opposite polarity, current summing with the main tap and creating a two-tap analog filter at 20Gb/s.

Notice that using a 10GHz non-interleaved, complementary, clocking structure

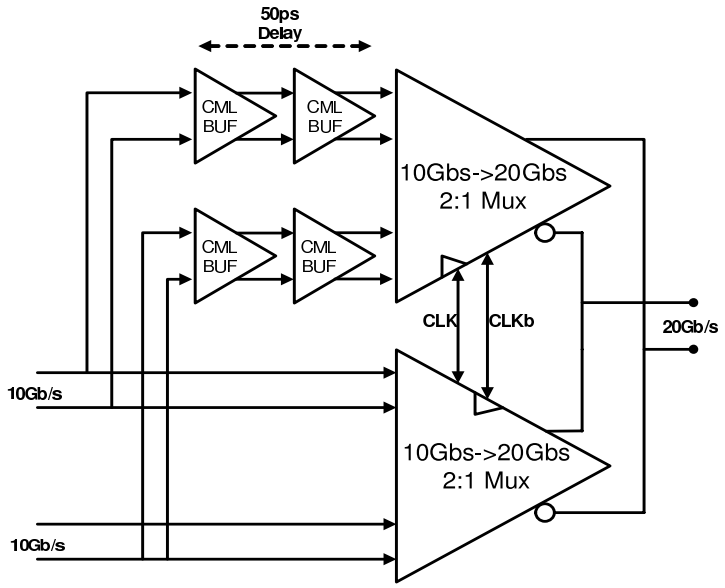

Figure 5.5: 10Gb/s to 20Gb/s Multiplexer with Two-Tap Equalization

does not allow for a simple design of the transmitter equalization as compared to using a conventional multi-phase transmitter architecture. Since the final transmitted output is at the 20Gb/s line rate, there are no delayed copies of the analog inputs, making equalization very difficult, especially at these high data rates. The implementation of the two-tap equalization above uses very minimal hardware for equalization, as it only requires a second output driver and CML gates for the 50ps delay. The strength of each tap is controlled by the current source in the output driver. The 50ps delay can be modified somewhat by the current consumption in the CML buffers, but the tuning range of this 50ps delay is not very large as the delay is mostly determined by the RC time constant of the CML output load.

Figure 5.6 shows the SPICE simulation of the 10Gb/s output waveform along with the delayed 10Gb/s data stream that is seen after the two stages of CML delay. Notice that the simulated delay through the CML gates is about 45ps, within 10 percent of the nominal 50ps delay needed. One disadvantage with this analog delay method is that the propagation delay through resistively loaded CML gates can vary significantly in the presence of process variation. One future possible solution for the process spread of the 50ps delay cell is calibration of this delay cell by using

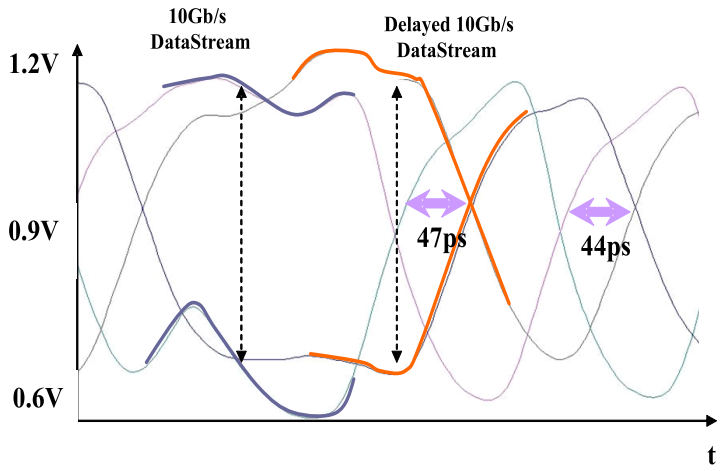

Figure 5.6: SPICE Waveforms of 10Gb/s Multiplexer Output and Delayed Output (for Equalizer Input)

digitally programmable switch capacitors and variable resistance PMOS output loads for the CML loading. This would effectively change the RC time constant of the delay cell, making it adjustable and programmable and thereby allowing for large process tolerance.

5.3 20Gb/s 2:1 Output Driver

The final 2:1 multiplexer is seen in Figure 5.7. This architecture is slightly modified from the first generation output stage, in that there are three, stacked NMOS transistors, which are used for differential current steering. There are two disadvantages and a few advantages of this new output driver. The first disadvantage is that the gain/bandwidth of this stage is less, due to the lower gate overdrive for the three NMOS devices, resulting in the device sizes being larger. Second, the power dissipation is doubled for this output stage – approximately 14.4mW total. (7mA per differential pair current). However, the benefits outweigh these costs. First, by including the steering of the clock current to the opposite side from the data, data dependent jitter kickback into the LC-VCO oscillator is almost completely removed. That is, there is no path for capacitive kickback or data-dependent modulated charge

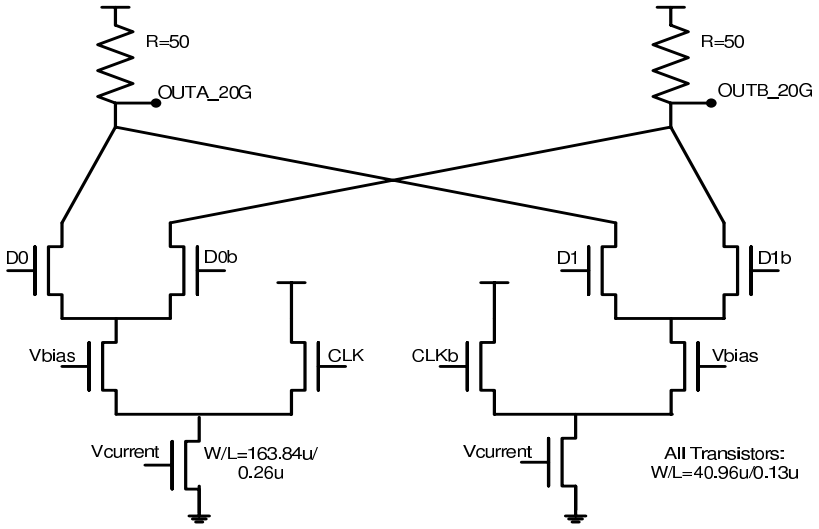

Figure 5.7: Second Generation Output Driver with Current Source Nodes

to disturb the stability of the LC resonator. Second, this architecture has the ability vary the output drive strength, due to the bottom NMOS current source switch. This is an important feature as now there is a mechanism for controlling the main and equalizer filter tap weights. Finally, the clock steering current pairs allow modu-

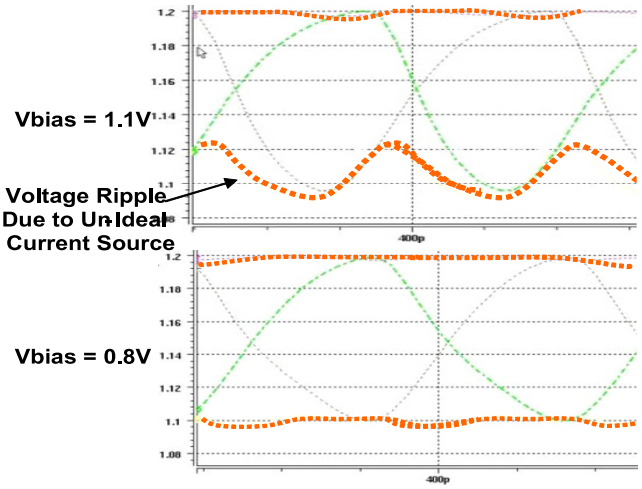

Figure 5.8: 20Gb/s Output Waveforms with Different Values of Vbias

lation of the output waveform based upon the common-mode level of the clock input when compared to the voltage Vbias. In the previous architecture, too low of a clock

common mode voltage caused the output to have a noise modulation when the zero crossing of the complementary clocks occurred. Both NMOS current sources would be conducting current during the clock zero crossing, but neither current source would be fully "on". Hence, the tail node voltages of the current sources would modulate, ultimately resulting in a 20 percent voltage ripple in the output eye diagram. By enabling a knob to adjust the common levels of the complementary clocks as well as the Vbias voltage, the output noise waveform is reduced. The noise reduction is observed in Figure 5.8, where the 20Gb/s output is seen to have different voltage ripples based on the Vbias voltage. When Vbias=0.8V, the differential pair commutates the current faster in regards to the 10GHz clock voltage. However, when Vbias=1.1V, the current switch turns on slower with regards to clock voltage, resulting in the differential pair no longer acting as a current source and subsequently the large voltage ripple seen during the zero crossings.

5.4 20Gb/s Transmitter Eye Diagrams

The following simulation diagrams show the post-layout simulated results of the second generation transmitter in various configurations: a) ideal transmitter output, b) transmitter output in a low pass channel (-6.5dB @ 10GHz), c) transmitter output with two-tap equalization in an ideal channel, and d) transmitter output with two-tap equalization through a low pass channel. Figure 5.9(a) shows the simulated 20Gb/s output eye diagram over about 10,000 bit times, using SPICE and overlaying the eye over a 50ps bit time. The differential eye height in an ideal transmitter is 80mV high and 43ps wide. One important feature to notice is the appearance of the voltage ripple occurring during the clock zero crossing, that is smaller than the same phenomena that occurs in the first generation transmitter. Also notice that since this voltage ripple occurs at 20GHz frequency, it is attenuated significantly when sent through a low-pass channel as seen in Figure 5.9(b).

A second feature to notice is the appearance of a bimodal response in the eye diagram during the zero crossings. This bimodal response arises due to the limited

bandwidth of the 10Gb/s data streams that precede the final 20Gb/s output multiplexer. As mentioned in Section 5.1, the variation in the 10Gb/s output swing causes the final 20Gb/s output multiplexer to have varying output current commutation, resulting in varying output slew rates. Seen in the SPICE waveforms of Figure 5.4, due to the low pass response, the 10Gb/s data voltage magnitude will be smaller after a data transition than for two consecutive data values being sent. For three or more consecutive data values, the voltage magnitude limits to that of two consecutive data values. Since the ability to switch the output current depends on the magnitude of the differential input voltage to the final output multiplexer, a smaller input differential swing results in a slower rise/fall time of the commutating current. This is a difficult effect to prevent at this high data rate. Equalization and bandwidth enhancement techniques such as shunt peaking at the 10Gb/s output can help ameliorate this rise/fall time problem. It should also be noted that the bimodal response causes larger ISI in the form of a tri-modal response when sent through a low-pass channel, as described in [6] and seen in Figure 5.9 (b).

Figure 5.9(b) also shows the ISI increase and loss of SNR due to having a low pass, single-pole channel response where the loss is -6.5dB @ 10Ghz.

Figure 5.10(c) shows the 20Gb/s eye diagram through an ideal channel but with the two-tap equalizer on, with a second tap value of 0.37 (main tap=1). Due to speed difficulties in building a two-tap equalizer at 20Gb/s data rate, the equalization shows closure of the eye in the horizontal direction. Essentially, the current summation for the main tap and the equalizer tap do not occur at exactly the same time, causing the tri-modal response in the equalized eye. [2] However, even with this difficulty in building an accurate two-tap filter, the equalization is still shown to have a positive effect in a low-pass channel. In Figure 5.10(d), the SNR is seen to be 2x larger than without the equalizer(33mV vs. 62mV).

[2]The output driver of the equalizing path has different 10Gb/s voltage magnitudes as compared to the main path, due to the two stages of CML delay buffering. Hence, the commutation of current for the equalizing path is different from the main path, resulting in time variations of the output current summation.

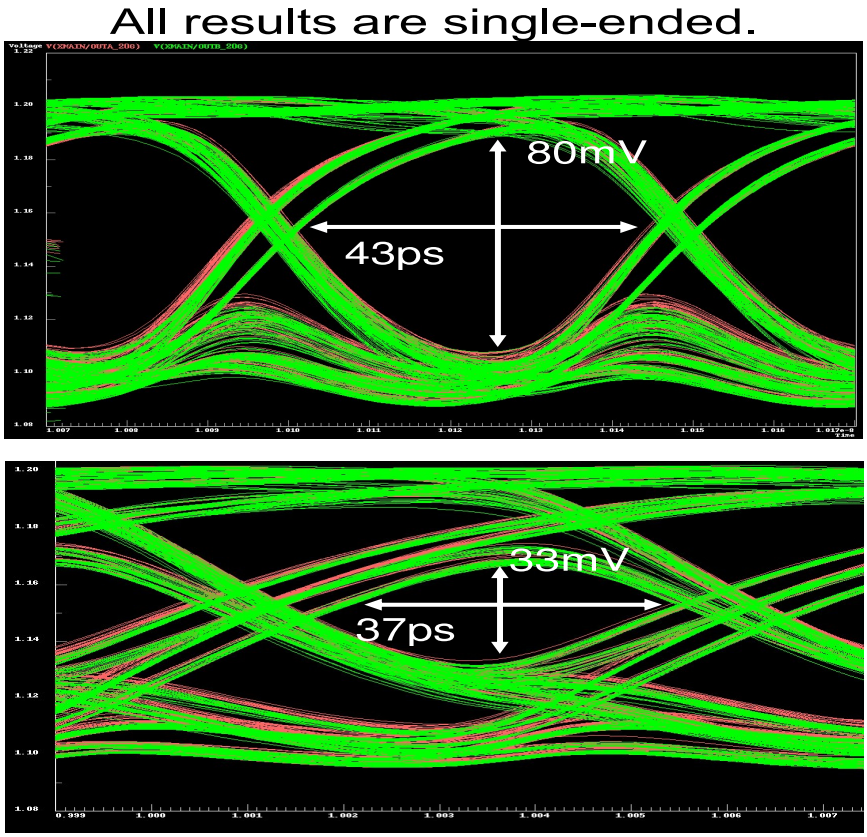

Figure 5.9: (a) Simulated 20Gb/s Eye Diagram with No Channel Loss (b) Eye Diagram with -6.5dB Loss at 10GHz and Two-Tap Equalizer Off

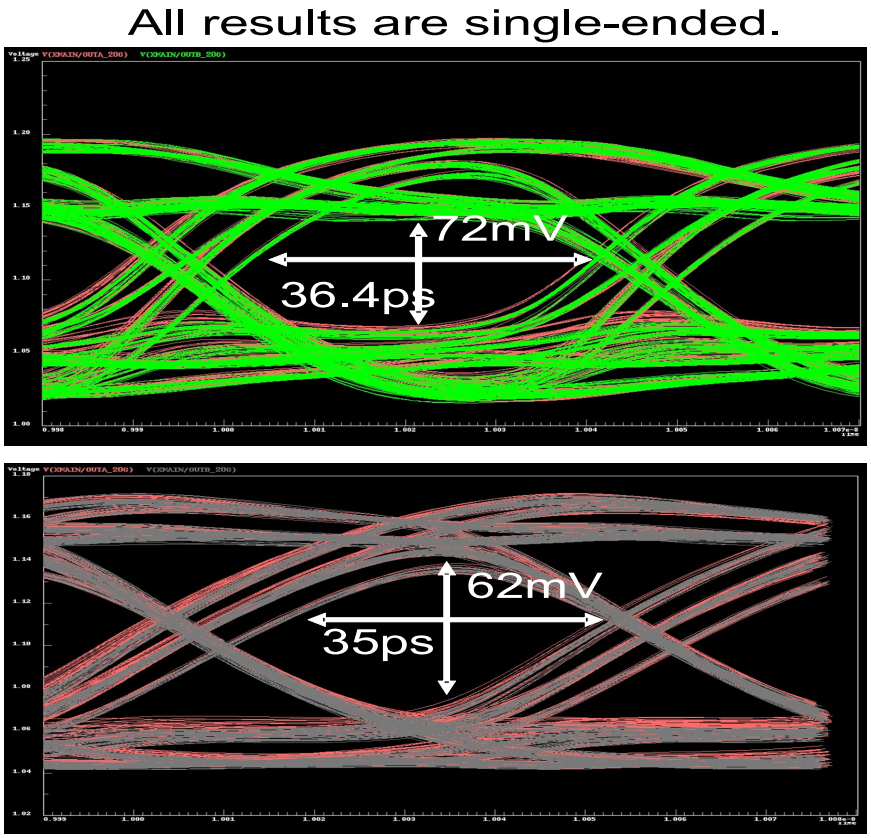

Figure 5.10: (c) Simulated 20Gb/s Eye Diagram with No Channel Loss and Two-Tap Equalizer On (A = 0.37) (d) Eye Diagram with -6.5dB Loss at 10GHz and Two-Tap Equalizer On (A = 0.37)

Chapter 6

Receiver Architecture

It was shown in the two previous chapters how a LC-VCO can directly drive the final transmitter output multiplexer, resulting in an increase in timing margin as the transmitter output jitter is due solely to the quiet LC oscillator. Due to its low static phase offset, low power supply sensitivity, and reduced power dissipation, this direct drive LC-VCO architecture addresses some of the timing uncertainty issues present in conventional high speed transmitter architectures.

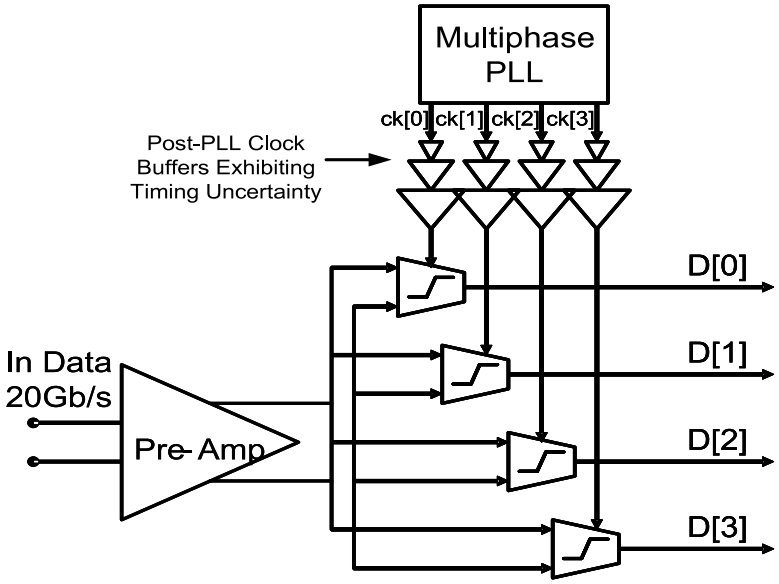

Figure 6.1: Conventional Receiver with Timing Uncertainty in Post-PLL Clock Buffers

Just as a conventional transmitter using multi-phase clocking is vulnerable to increased timing uncertainty, the same timing issues also exist in a conventional receiver. Figure 6.1 shows a block diagram of a conventional receiver where a multi-phase, time-interleaved PLL is created to achieve time-interleaved, slicer quantization. Note that the post-PLL clock buffers that follow the PLL output clock may exhibit both increased static phase offset and power supply sensitivity, resulting in uncertainty of the clock edge when it finally quantizes the input data. Ultimately, this timing uncertainty is the limiting bottleneck for low bit error rate, data recovery at high data rates.

The receiver developed in this thesis addresses this timing uncertainty issue as well as a few other concepts. First, the implemented receiver reduces timing uncertainty by using a direct LC-VCO to drive the front-end receiver demultiplexer. Similar to the transmitter, the LC oscillator's complementary clock phases directly sample the 20Gb/s differential input, with no post-PLL buffers between the VCO and the input samplers. This effectively removes the problem of static phase offset and power supply induced jitter introduced by clock buffers, leaving only the intrinsic jitter and phase offset of the phase locked loop.

A second new concept of the implemented receiver is the use of a downsampling demultiplexing technique to convert the fast 20Gb/s analog input into eight digital data streams at 2.5Gb/s. Due to gain-bandwidth limitations, the 20Gb/s analog input is downsampled to two 10Gb/s analog data streams, and then further downsampled to four 5Gb/s analog data streams before digital quantization to eight 2.5Gb/s digital data streams. This technique reduces the power and bandwidth limitations of the slicers to the lower 2.5GHz frequency, as opposed to digital quantization at 5GHz.

The third new receiver concept is the implementation of a source synchronous clock for data recovery. Here the transmitter sends a mesochronous clock to the receiver, where the sampling clock phase position is calibrated at chip startup using BER curve measurement, and then statically set during normal operation. This technique reduces complexity and power dissipation compared to a conventional receiver that generates a 2x oversampling clock for clock and data recovery.

6.1 Receiver Architecture

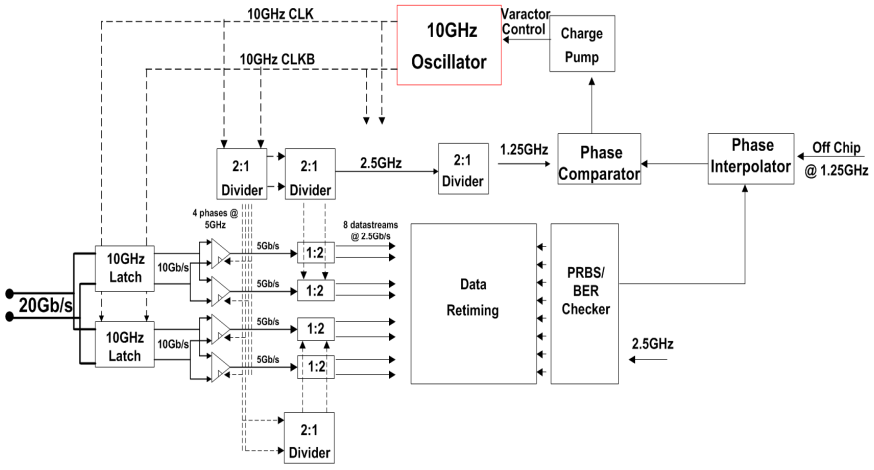

Figure 6.2: Receiver Block Diagram

The general block diagram of the receiver is seen in the Figure 6.2. A 20Gb/s input data stream is initially demultiplexed into two 10Gb/s data streams by two 10GHz latches. Notice that the two front-end 10GHz, sample-and-hold circuits are directly driven by a LC oscillator. This removes the possibility of increased static phase mismatch as well as power supply induced jitter as a result of the addition of clock buffers, similar to the technique seen in the transmitter. After the two 10GHz sample-and-hold gates (Av(gain) = 0.8), the two 10GHz values are further demultiplexed to four 5Gb/s data streams (Av = 2). Demultiplexing the data streams down to lower speed, analog streams relaxes the speed and power constraints of the quantizing digital slicers. For example, without this further demultiplexing from 10Gb/s to 5Gb/s, time-multiplexed regenerative latches need to run at 10Gb/s, which is extremely difficult to achieve without burning excessive power or area (such as using inductive shunt peaking). These 5Gb/s data streams are finally quantized by eight offset compensated, regenerative Strongarm latches. Data retiming and PRBS checking is achieved at 2.5GHz. Notice that there is no explicit 2x oversampling clock at the front-end. Therefore, data recovery is done during an initial calibration stage where the optimum clock sampling point is swept across an entire bit duration, and the BER is measured at each point across this sweep, effectively building a BER

bathtub curve. The interpolator used for this sampling time sweep is built with high redundancy since the minimum phase step size is important. Therefore, a combination of capacitance as well as current summing is used to achieve fine phase resolution. After measuring the optimum sampling point, the most optimum phase for low BER is set and the receiver can now run in active, demultiplexing mode. Notice that such a receiver only works for mesochronous systems, where the receiver clock has no frequency slip from the transmitter data.[45] The following section describes the challenges in building the front-end 20Gb/s 1:2 demultiplexing.

6.2 20Gb/s 1:2 Demultiplexing

6.2.1 Receiver Input Bandwidth

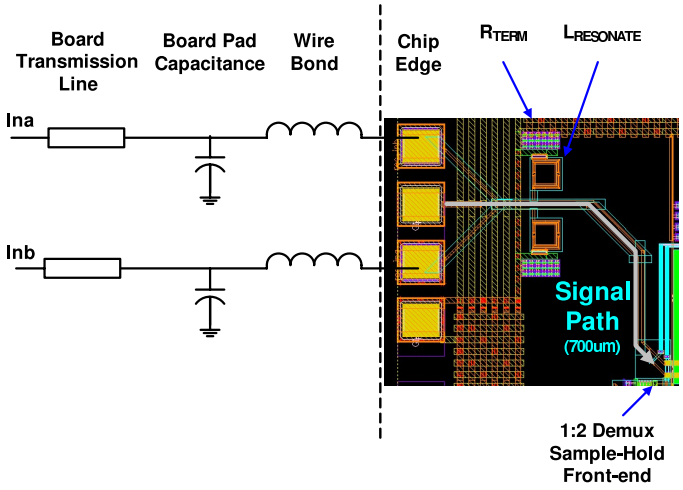

Figure 6.3: Viewpoint of Parasitics when Signal Travels from Board to On Chip

The input bandwidth of the receiver differential input, not the front-end 10GHz NMOS sample-and-hold circuits, limits the maximum performance of the link. Figure 6.3 (a) illustrates the impedance coming in from off-chip. The differential signal travels through the 5mil, 50 Ohm transmission lines, connecting to small board pads, passing through short (L less than 0.3nH) wire bonds, and onto the chip input. At the chip input, the signal then passes through 70um x 70um RF pads, through 900um

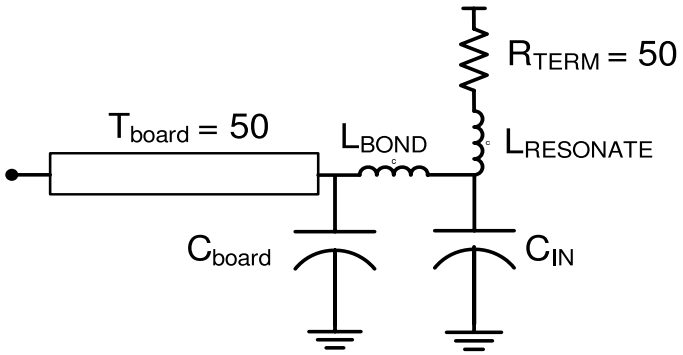

Figure 6.4: Viewpoint as Lumped Parasitics as Signals Travel from Board to On Chip

of top layer M6 metal traces (5um wide / 5um spacing), before actually reaching the first 1:2 NMOS demultiplexing sample-and-hold circuits. 6.4 (b) illustrates the schematic lumped model of the front-end. Tboard is the 50 Ohm board transmission line; Cboard is the small, board via capacitance; Lbond is the short bond wire inductance; Rterm is the 50 Ohm on die poly resistor termination resistor; Lresonate is the shunt peaking inductor; and Cin is the parasitic extracted input capacitance of the differential lines before they arrive at the receive demultiplexing front-end. Cin is relatively large, as the differential nodes need to travel a few hundred microns to get from the bonding pads to the receiver input. The die is wire-bonded directly to the board by adding an additional bondable gold layer on the board surface. As well, the layout of the chip-on-board allows for direct Cascade Microtech contact to the differential front-end, which improves bandwidth by removing some of the frequency attenuation of the bondwire.

Figure 6.5 illustrates the simulated bandwidth after parasitic extraction of the receiver front-end with various values of shunt peaking Lresonate: 0nH, 0.4nH, and 0.8nH. Figure 6.5 illustrates this bandwidth with Lbond=0nH, such as when the Cascade Microtech probe touches the input pads, and there is no bandwidth reduction due to wire bonding from the chip to the board. From simulation, the -3dB bandwidth increases from 14GHz to 20GHz using a 0.4nH shunt inductor, representing bandwidth enhancement of more than 43 percent.

Figure 6.6 illustrates this bandwidth with Lbond=0.4nH, which is the value of the bondwire inductance when chip-on-board is used. The -3dB bandwidth increases

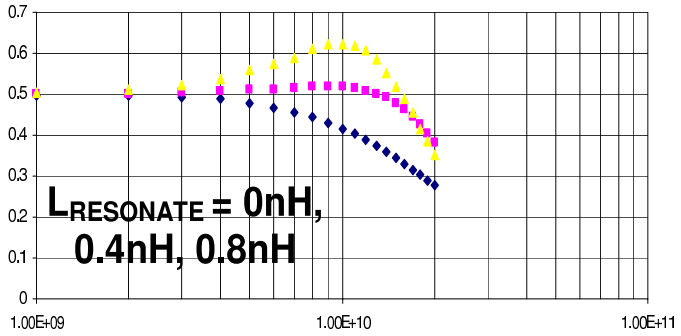

Figure 6.5: Input Bandwidth when Bond Wire L=0nH, as in Direct Cascade Microtech Probe Connection

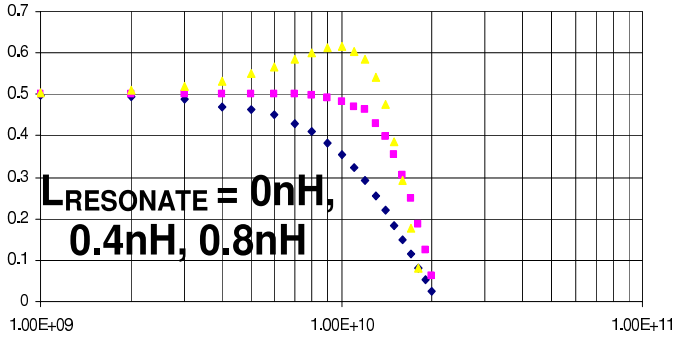

Figure 6.6: Input Bandwidth when Bond Wire L=0.4nH

from 10GHz to 15GHz with the 0.4nH shunt inductor, increasing input bandwidth by 50 percent. Due to these significant improvements in bandwidth of the receiver input, a small spiral inductor of value 0.3nH is used for shunt peaking.

Since passive inductors exhibit a significant area disadvantage, the use of a spiral inductor for shunt peaking was compared with using a small area, active transistor NMOS gyrator. Figure 6.7 shows that the gyrator structure has inferior frequency response characteristics, due to the additional capacitive loading of the NMOS parasitic capacitance and the limited gain bandwidth of the gyrator transistor. Additionally, the gyrator structure has a nonlinear behavior that depends on the output swing level, since the current through the gyrator is modulated based upon output swing. This current modulation changes the gyrator's gm and therefore reduces the output bandwidth enhancement. Finally, due to the threshold voltage drop of the NMOS, the

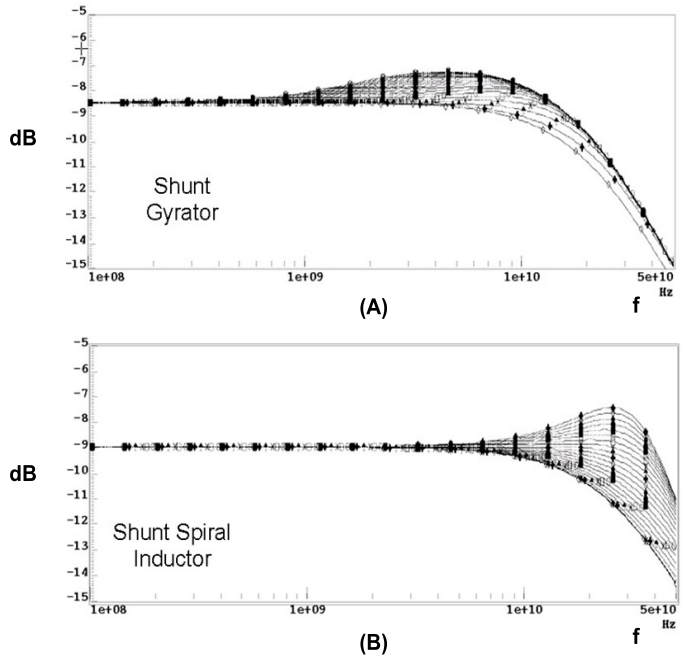

Figure 6.7: (a)Bandwidth Extension of Input Front-End Using Varying NMOS Gyrator Structures (b) Bandwidth Extension of Input Front-End Using Varying Shunt Spiral Inductors

gate must be biased higher than the supply for the circuit to work close to the power supply swing, making the bias conditions for the gyrator difficult to achieve. For all these reasons, a small spiral inductor is preferable for shunt peaking as compared to a NMOS gyrator structure.

6.2.2 Intrinsic Noise Sources

kT/C noise [4] describes the fundamental noise fluctuation of the input receiver samplers, where the RMS noise voltage(s) is given by: where Cin is the input capacitance

$$\sigma = \sqrt{\frac{kT}{Cin}}$$

of the sample-and-hold latch. Assuming that the lower limit of parasitic sampling capacitance is 10fF, 10σ (thermal noise) is equivalent to 6mV, quite negligible if the expected differential input is around 50mV. We extracted from layout the parasitic

sampling capacitance of the NMOS switch, resulting in a sampling capacitor of 205fF – this includes 83fF of parasitic wiring capacitance, 22fF of the sampling switch Cgd, 8.2fF of sampling switch Cdb, and 90fF of proceeding stage latch input capacitance Cgs. From the equation above, the input-referred kT/C thermal noise (10σ) is 1.4mV.

$$\sigma = \sqrt{4kTRf}$$

Resistive thermal noise (seen above) in the 50 Ohm termination resistor is given by the equation above. For a frequency up to 20GHz, the 10σ noise is 1.3mV. Thus, both the capacitive and resistive thermal noise values give a total input referred noise of less than 2mV, such that intrinsic channel noise does not limit the input sensitivity.

6.2.3 Front-End 10GHz Sample-and-Hold Circuits

Simple NMOS sampling switches are used for demultiplexing the 20Gb/s input data into two 10Gb/s data streams. There are other possible front-end sampler architectures, including differential pair-type sample-and-hold amplifiers, such as CML or GAD[37]. There are a few advantages for using just simple pass gates compared to these alternatives.

The first is exceptionally high bandwidth. Since the fT of the device relates to the speed/time it takes for electrons to transit from drain to source, the NMOS pass gate allows for extremely high bandwidth, given the short channel length of 0.13um. With technology scaling transistor channel lengths, the sampling bandwidth of pass-switch, sample-and-hold structures will improve with scaling. We characterized the implemented NMOS switches for aperture time, using various rise/fall times of the sampling clock. This characterization is especially important because in the implemented receiver architecture the use of a 10GHz sinusoidal clock (as opposed to a square wave type clock coming from an inverter buffer) is used as the sampling switch clock signal. Figure 6.8 below illustrates the approach used to characterize the input bandwidths for these varying clock edges.

The 1:2 demultiplexer is first sent an input stimulus of a lone '1' pulse, a 100mV differential input with 10ps rise time and 40ps hold time. With the 1:2 demultiplexing

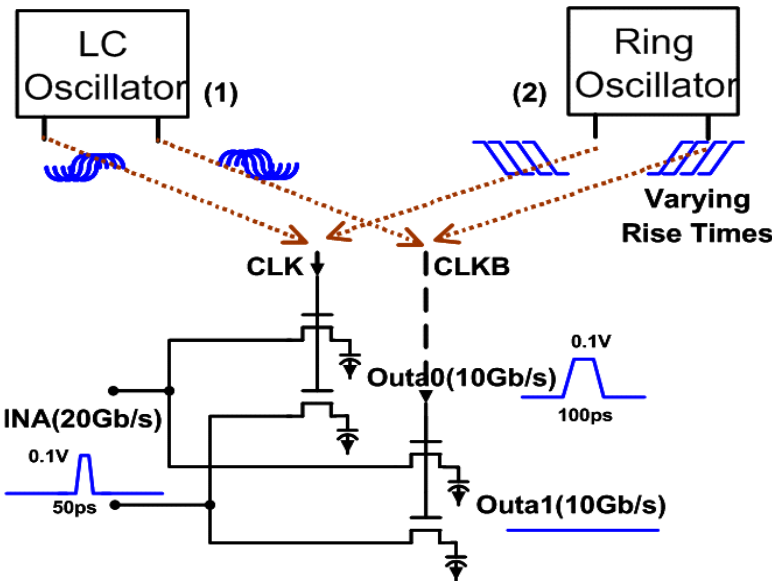

Figure 6.8: Switch Bandwidth Comparisons Between LC-Oscillator Sinusoid and Inverter Square-Wave Sampling

structure, the data is demultiplexed to 10Gb/s (Outa0) while the other data stream remains unchanged (Outa1). For the first experiment, the sample-hold switches are sampled by a 10GHz sinusoidal clock, whose sampling edge is swept across the entire bit time, from 0 to 50ps. The second experiment uses a complementary 10GHz square wave clock with varying rise/fall time, which is also swept across the 50ps bit time. These results are seen in Figure 6.8. The third experiment simulates a realistic CMOS implementation by taking a 10GHz, 20ps rise-time square wave and sending it to two stages of FO2 inverters (Tpd = 18ps, Trise = 35ps) before this buffered clock is used to sample the NMOS switch.

As can be seen from Figure 6.9, the degradation in bandwidth of the LC oscillator sampling relative to a pseudo-square wave clock sampling is minimal. In effect, the sampling aperture time of sinusoidal sampling is comparable to a 30ps rise-time square wave or a FO2 inverter-driven clock chain. While the 10GHz sinusoidal sampling might have bandwidth degradation due to its slow rise time (i.e. 50ps sine wave signal), at 10GHz most on-chip inverters in 0.13um CMOS already experience limited

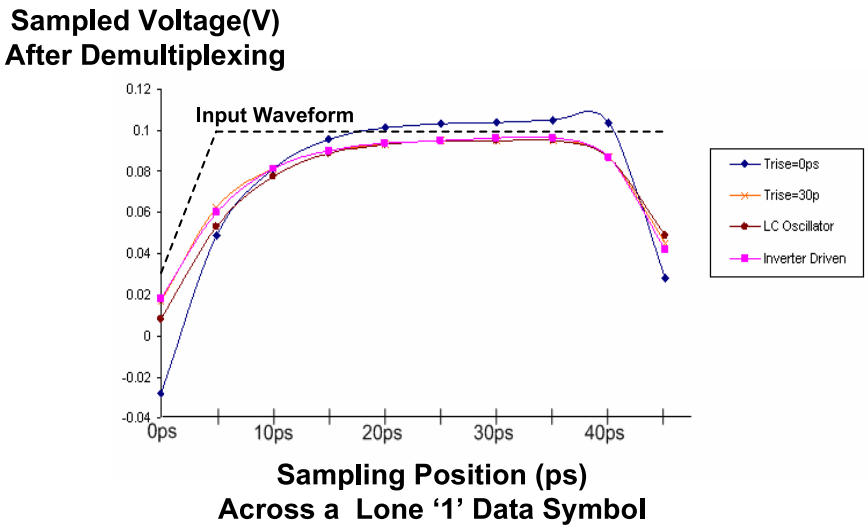

Figure 6.9: Sampled Voltage After Receiver Demultiplexing of a 20Gbs Transmitted, Lone '1' Pulse

rise-times. For example, with a two stage FO2 inverter driven clock tree, the performance (Tpd = 18ps, Trise = 35ps) is likely the fastest edge-rate possible in 0.13um technology. For even higher sinusoidal oscillation frequencies (i.e. 20GHz) used for sampling, the sampling bandwidth increases relative to a square-wave sampling, as the sinusoidal oscillation frequency is much higher compared to a technology limited RC-based inverter. For example, while an inverter rise time is limited to 35ps in 0.13um CMOS, LC oscillators have been shown to have resonant frequencies above 40GHz. [30, 31] Thus, MOS sampling using direct LC oscillator sampling may possibly achieve extremely high 1:2 demultiplexing frequencies.

The second benefit of a NMOS sampling switch multiplexer is the fact that threshold offset mismatch does not have a significant effect on sampling accuracy as opposed to a sample-and-hold circuit that uses a differential pair input. With a differential pair sampler, the threshold offset directly correlates to gain mismatch, resulting directly in an offset-amplified output. For example, 30mV of threshold mismatch between the differential pairs results in a minimum input sensitivity to the input data of 30mV. Therefore, this offset effectively determines the minimum quantization sensitivity of

the receiver front-end. Offset cancellation [50] helps mitigate this problem, but normally at the expense of increased complexity, bandwidth degradation, and residual offset due to finite quantization error.

With NMOS sampling gates, the threshold mismatch does not translate directly into output mismatch. The reason for this is that when the transistor operates in its linear regime, the gate overdrive voltage is very large, and subsequently any voltage mismatch has a minimal effect on the NMOS on-resistance. Figure 6.10 illustrates this phenomenon. In this simulation, the NMOS sampling switches(with a 0.6V common-mode voltage) are simulated with a differential threshold offset swept between 0 - 40mV, using an input stimulus similar to the bandwidth simulation in Figure 5.6. Notice that the differential output voltage is nearly the same for any value of differential threshold voltage offset. Hence, any mismatch between these gates (i.e. threshold variation, gate length variation) is to first order mitigated, as compared to a differential pair-type sample-hold circuit.

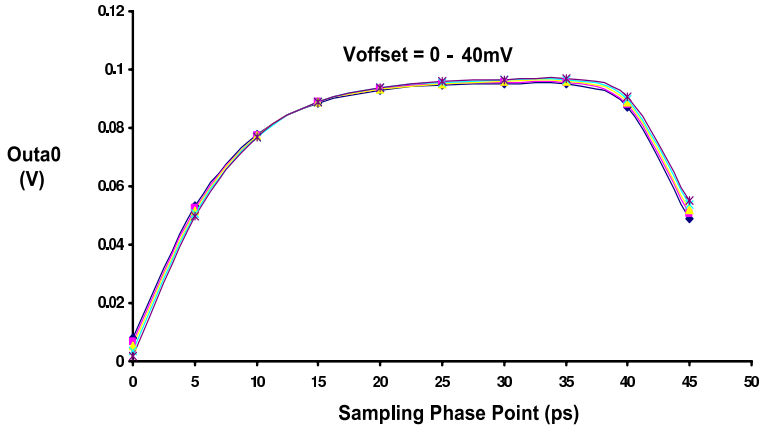

Figure 6.10: Sampled Differential Output for Varying Threshold Voltage Offset Between Sampling Switches

While the simple sampling switch has some advantages, there are a few disadvantages. One is that the NMOS sampling bandwidth depends strongly on the input common mode voltage, as the NMOS on-resistance is highly dependent upon DC bias and gate overdrive. Figure 6.11 shows the bandwidth of NMOS sampling for varying values of input common mode range. This problem can be alleviated by using a DC

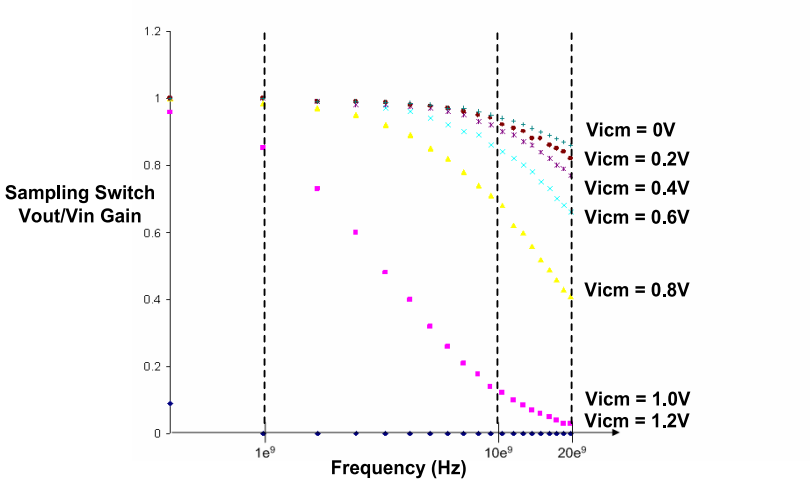

Figure 6.11: NMOS Sampling Bandwidth for Varying Input Common Mode Voltages

blocking capacitor for the transmission line and setting the local 50 Ohm termination voltage supply to 0.6V. The common mode of the input cannot dip too low, otherwise the analog amplifiers following the sampling switches will not have adequate input common mode range.

Notice that unlike most sample-and-hold receiver architectures, there are only two phases of complementary clocks, essentially sample mode and hold mode–there is no reset phase to remove hysteresis. This is unlike other conventional multi-phase clocking architectures, that allow for some interleaved clock phases to perform internal capacitor reset, level-shift boot-strapping, etc. Hence, two-phase complementary clocks require higher sampling bandwidth as sampler hystersis is a more significant problem. One possible revision of the design would be to add unity gain, current-steering amplifiers/buffers immediately after the NMOS sample-and-hold signals. This would reduce receiver NMOS loading capacitance and increase sampler bandwidth, as more than half of the NMOS parasitic capacitance results from the wiring capacitance connecting the demux to the next stage differential pair. However, such an addition would increase power consumption, area, and timing uncertainty due to amplifier delay.

6.3 10Gb/s to 5Gb/s Demultiplexing Sample-and-Hold Amplifiers

After demultiplexing the 20Gb/s input data into two analog data stream of 10Gb/s, the analog differential signals are further demultiplexed into four analog data streams of 5Gb/s that are then finally sent to the slicers. The schematic of this demultiplexer is seen in Figure 6.12. It should be noted that the initial design cycle was to demultiplex the data to two 10Gb/s datastreams, and then to quantize the two analog 10Gb/s data streams into four digital 5Gb/s signals. However, design iteration showed that timing margin was extremely difficult to meet, as the Strongarm latches have limited bandwidth when used with lower supply voltages. Also, the receiver PLL design created only four phases of clock at 5Ghz, making it difficult to have an explicit reset phase, which is integral towards regenerative latch design. Finally, it is also difficult to achieve fast transient edge rates at a 5GHz clock frequency, which is important for minimal Strongarm latch regeneration time.

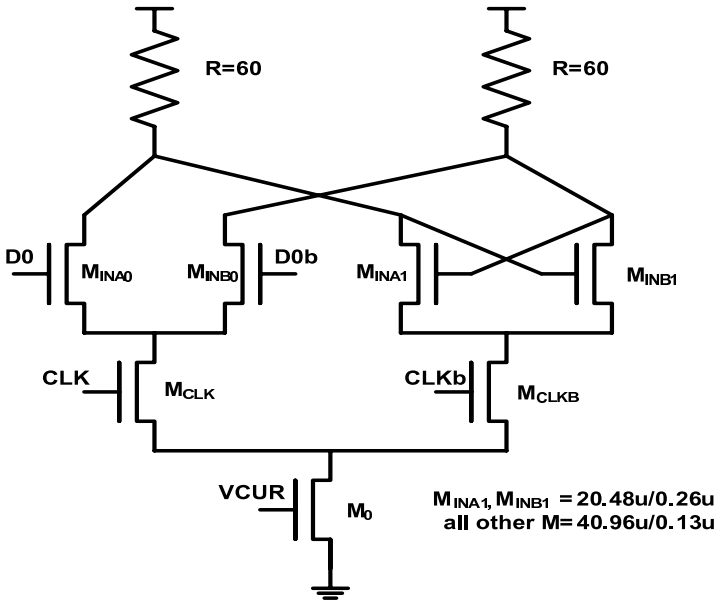

Figure 6.12: 10Gb/s to 5Gb/s Demultiplexer

Instead of going from two analog 10Gb/s data streams directly into four digital

5Gb/s data streams, this architecture further demultiplexes the four 10Gb/s analog data into 5Gb/s analog data, using four sets of current mode latches as seen in the figure. Default current consumption is 2.5mA, exhibiting a gain close to 1.3. The CML gate also provides level conversion, as the CML differential input is common mode referred to the differential common mode of the receiver termination, and the CML differential output is close to 1V, before the slicer input.

SPICE simulation of the demultiplexing from 20Gb/s to 5Gb/s can be observed in Figure 6.13. The 200mV 20Gb/s differential input is sampled and held in two 10Gb/s differential analog latches, exhibiting little gain loss. These two 10Gb/s differential inputs are further demultiplexed into four 5Gb/s analog data streams, with a gain of 1.5, resulting in a peak differential voltage of around 140mV. These analog voltages are then aligned with the eight full-rail clocks running at 2.5GHz, finally retiming the analog values to an eight bit digital word.

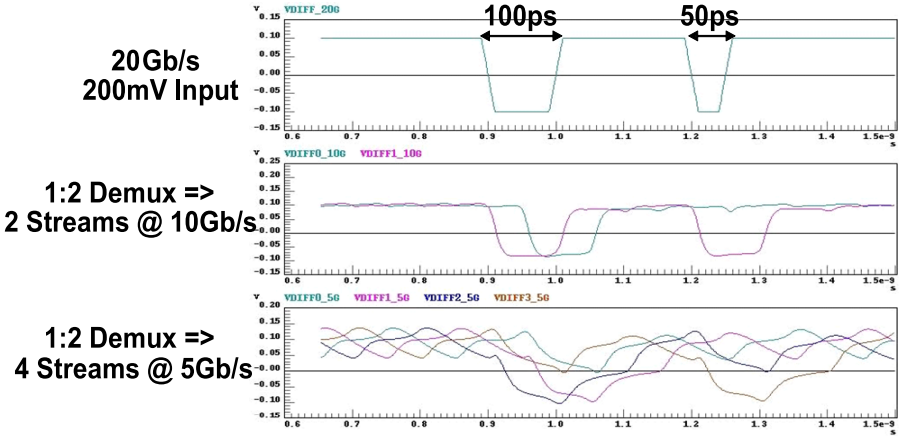

Figure 6.13: SPICE Simulations of Demultiplexing from 20Gb/s to 5Gb/s

6.4 Slicer/Latch Design

Figure 6.14 shows the circuit architecture for one bit slice of the slicer, demultiplexing a 5 Gb/s analog value into a 2.5Gb/s digital value. PMOS pass gates at the differential input allow for the inputs to be shorted together during calibration of offset

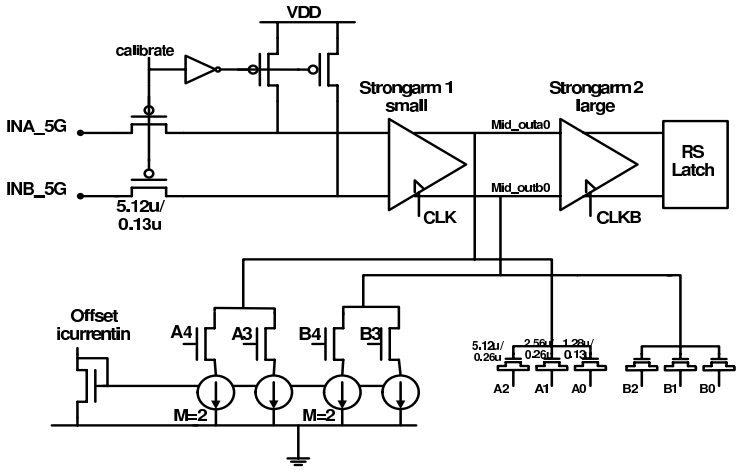

Figure 6.14: Circuit Schematic of Strongarm Slicers with Offset Compensation

compensation. The complete slicer uses two cascaded Strongarm latches operating on complementary clocks, before the regenerated signals are amplified to full value within the RS latch. The first Strongarm latch has devices slightly smaller than the next stage Strongarm to decrease the loading on the preceding stage. Offset compensation is achieved with two mechanisms, a coarse and fine-tune offset control. A 3-bit NMOS switch capacitor bank imbalances the parasitic capacitance of the first Strongarm latch, thereby achieving a minimum offset quantization of around 10mV. A coarse 2-bit current steering bank offsets the current pulldown path in the first Strongarm's output capacitance. As the current discharging path depends on the current source value, this current-steering bank allows for a wide dynamic range of offset compensation, programmable by the current source bias.

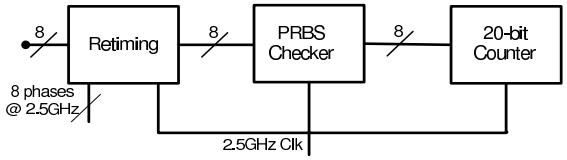

Figure 6.15: Block Diagram of Digital Back-End

6.5 On-Die Bit Error Measurement and Static Clock Recovery

In the previous sections, it was assumed that the clock sampling the analog differential input was optimally placed in the center of the data eye to achieve maximum differential swing, and hence maximum receiver sensitivity. However, determining the exact phase position of this sampling edge is very difficult as the eye width is only 50ps. In an ideal link, the sampling clock would be placed at exactly the center of this eye width, which is 25ps from either edge transition. However, due to unidealities in a real system, such as asymmetric rise/fall times, finite setup/hold/aperture times for sampling switches, and finite rise/fall times of the sampling clock, the ideal clock edge may not be exactly halfway between data edges. Many typical CDR systems use a 2x oversampled clock [25, 30], where the data is sampled both during its edge transition (to recover timing information) as well as after a 90-degree phase shift (for sampling the data voltage at its largest magnitude). This can be done either by generation of a 2x faster clock or by using a multi-phase clock. For example, a 20Gb/s receiver might use a quadrature locked 10GHz LC oscillator that generates four clock phases interleaved in time. The problem with these methods is that it is difficult to ensure precise clock phase symmetry due to static phase offset. Also, additional power is dissipated by creating this extra clock edge along with the associated latches this clock drives.

The receiver implemented in this thesis uses an open-loop, clock centering algorithm, where the clock edge is determined at startup by measuring the bit error rate at various sampling points of the differential input. The final sampling position of the sampling clock is statically set to where the measured bit error rate is lowest. This is analogous to an advanced source synchronous clocking approach, and is similar to that used by Casper for a 20Gb/s serial link transceiver in [3]. The one disadvantage to this static phase alignment clock recovery is the inability to track variations in sampling edge over time. For example, supply or temperature variations may cause the data edges to drift over time, such that the optimum sampling point must be adjusted more frequently. However, the channel is unlikely to vary excessively when the

serial link is finally installed in a large, stable computer system. Additionally, clock synchronization can be applied from a higher, system level. For example, periodic reset header packets (for example, after every 10us) can be used specifically for clock resychronization and alignment.

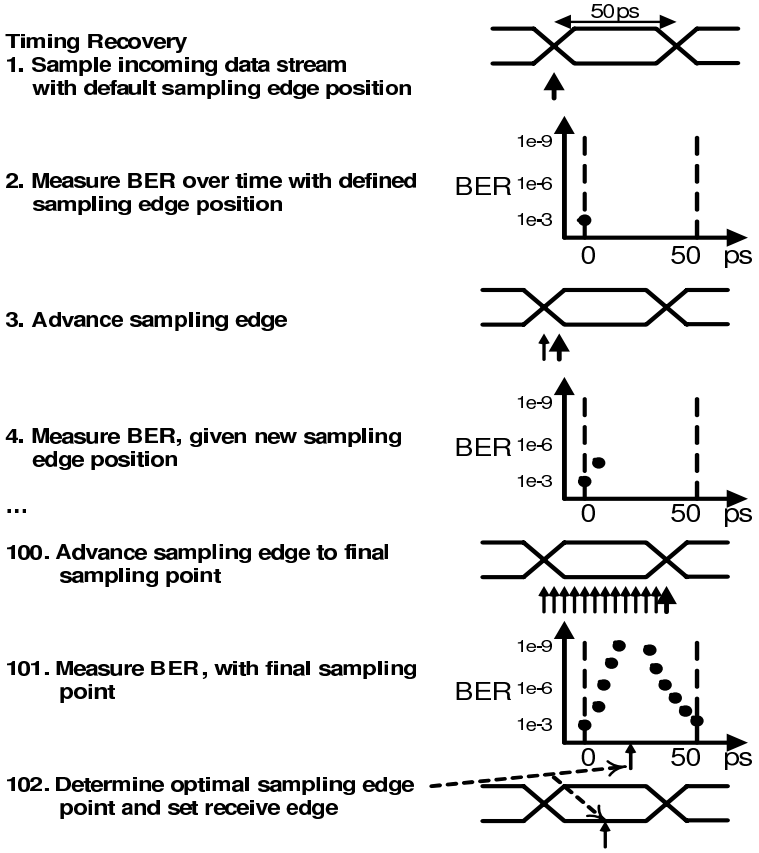

Figure 6.16: Algorithm for Determining Optimal Clock Sampling Point

The methodology for data recovery is seen in Figure 6.16. First, the receiver samples the incoming data at the default position of the sampling edge. After the analog demultiplexing downsamples the 20Gb/s data into eight 2.5Gb/s data streams, a single 8-bit word at 2.5Gb/s is created after clock retiming and reordering. This 8-bit word is sent to a PRBS checker that compares the bits received with the bits expected from the analogous 8-bit PRBS generator in the transmitter. If a bit error is determined, it is stored in a 20-bit ripple counter, allowing for accurate on-die measurement of bit error rate. A bit error 'valid' signal determines when the counter

is enabled to start counting bit errors, allowing for adjustability in the amount of time that is elapsed for the BER measurement. A finite state machine then moves the position of the sampling edge to the next quantization spacing and again repeats all the steps to measure the BER rate at the new clock position. This process is repeated for all positions across the whole range of one bit time, which is 50ps. For example, if the minimum quantization step size for each sampling edge is 5ps, the BER is measured at 10 points across the whole range. After every sampling position, the corresponding BER is recorded and a bathtub curve of BER is determined, such that placement of the optimum phase sampling position is obtained.

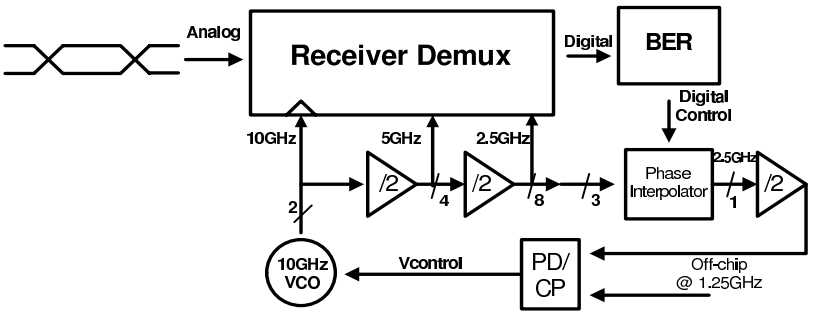

Figure 6.17: Block Diagram of Clock Recovery Using Optimal Clock Sampling

Figure 6.17 illustrates the receiver diagram of the CDR block diagrams. The BER checker and error counter digitally control the phase position of the phase interpolator, essentially stepping the 10GHz clock edge with relation the incoming 20Gb/s data stream. Hence, monotonicity and linearity of the phase interpolator are important such that the BER bathtub curve can be measured and the optimum phase selection point can be determined with accuracy.

6.5.1 Phase Interpolator for Clock Recovery

Phase position of the sampling edge is achieved by controlling a phase interpolator that advances the feedback clock into the phase comparator, thereby statically adjusting the phase of the 10GHz complementary clocks of the receiver front-end sampler. The phase interpolator consists of a coarse first stage, followed by a current summing fine second stage.

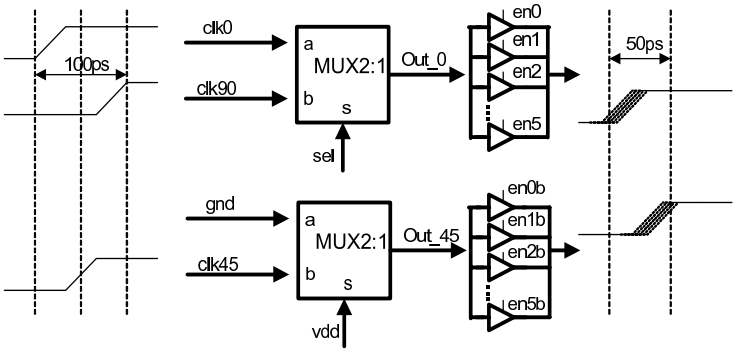

Figure 6.18: Coarse Phase Interpolator Design for CDR Loop

The first stage, coarse interpolation is seen in the Figure 6.18. The multiplexers on the front end determine which two phases, spaced 50ps apart, are used for the fine interpolator stage that follows. When sel=0, Clk0 and clk45 are chosen and sent to the current summing tri-state buffers. When sel=1, Clk45 and Clk90 are sent through to the tri-state buffers, giving a complete 100ps of phase range for the interpolator, which well exceeds the minimum phase rotation range of 50ps. The coarse current summing tri-state buffers can change the position of each phase by approximately 14ps.

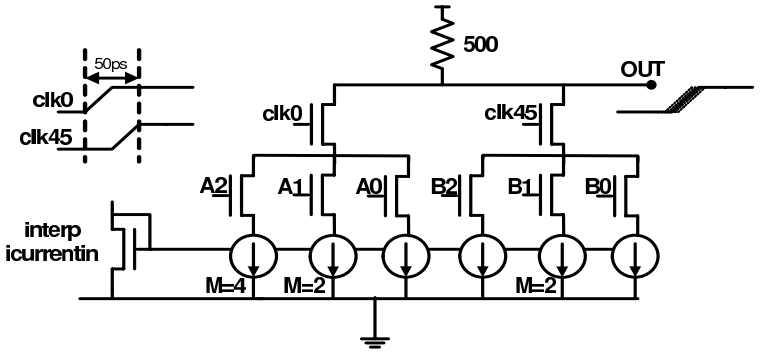

Figure 6.19: Fine Interpolator for CDR Loop

Figure 6.19 shows the fine stage interpolation that uses current summation to generate an interpolated phase between two clock edges spaced 50ps apart. The pulldown current sources can be adjusted independently to change the weighting of the fine current source bits. Simulation of the entire interpolator shows a minimum

phase step of 4ps and maximum phase range of 97ps.

6.5.2 Receiver Clocking Architecture

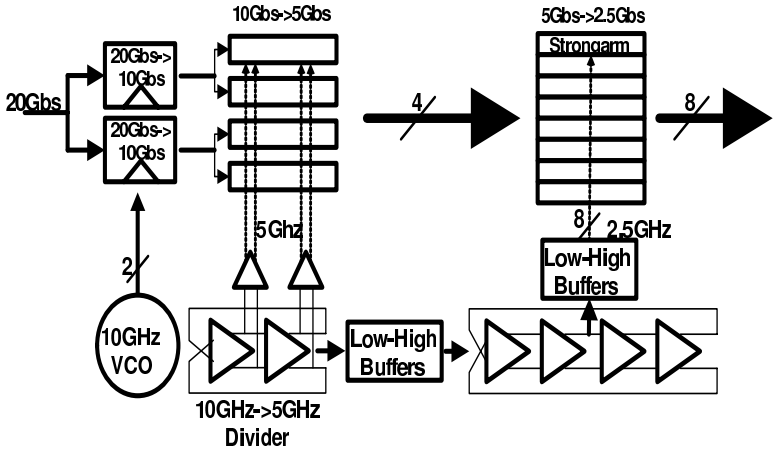

Figure 6.20: Clocking Architecture for Receiver

Critical to the success of the receive demultiplexing from 20Gbs to 2.5Gbs are the clock phases of the 10GHz clocks, 5GHz clocks, and the 2.5GHz clocks, and their alignment with the down-multiplexed data symbols. The delay from the clock dividers through the clock buffers to finally retime/demultiplex the data to a lower speed must be synchronized with the data passing through each demultiplexer, as seen in Figure 6.20. For example, consider when the data is retimed from 5Gb/s to 2.5Gb/s. While the 10Gb/s to 5Gb/s path has only a 5GHz buffer delay from the 10GHz divider, the 5Gb/s to 2.5Gb/s buffers have the delay of the 5GHz low-to-high buffers, the 5GHz-to-2.5GHz divider, and then the 2.5GHz low-to-high buffers before the clocks are used for demultiplexing. It is extremely difficult to match these clock delays for adequate setup/hold time for each stage, especially across process corners. One possible solution to this problem, not implemented in this thesis, is to use digitally adjustable delay verniers within each delay path. For example, adjustable resistor arrays in the CML buffers would change the RC delay in the 5GHz clock path compared to the lower frequency 2.5GHz clock path.

Chapter 7

Experimental Results

Two prototype chips were fabricated to validate the ideas in this thesis – a 20Gb/s transmitter prototype and a complete 20Gb/s transceiver prototype. Section 7.1 describes the experimental results of the 20Gb/s transmitter test chip and Section 7.2 describes the 20Gb/s transceiver prototype.

7.1 Prototype #1: First Generation 20Gb/s Transmitter

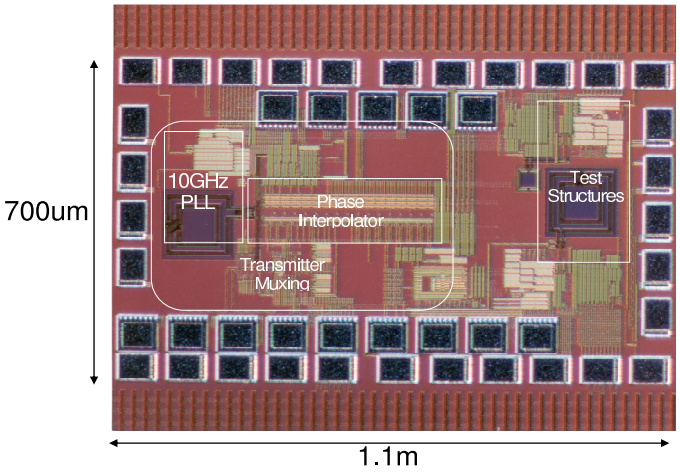

Figure 7.1: First Generation 20Gb/s Transmitter Die Photo (2003)

A 20Gb/s transmitter test chip was built in 0.13um standard CMOS technology, available through UMC (United Microelectronics Corporation). Figure 7.1 shows a die microphotograph of this chip. Including the bonding pads, the die measures 700um x 1100um. The two arrays of inner pads along the top and bottom of the die allow for on-die, high bandwidth 50-Ohm measurements using Cascade Microtech probes. The 20-Gb/s differential data stream travels through two sets of DC blocking connectors (one for the shield, the other for the transmission line), one meter of coax, and into the oscilloscope input. Total insertion loss at 10GHz is 3.4dB, as the line exhibits some amount of frequency-dependent attenuation.

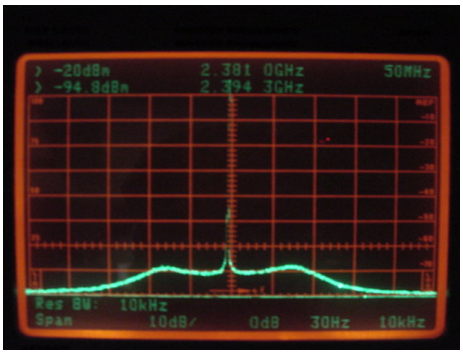

Figure 7.2: Power Spectral Density of 10GHz Synthesized Clock

Figure 7.2 shows the spectrum analyzer output of a 2.5GHz divide-by-4 clock from the standalone PLL, locked at 9.6GHz. The measured -3dB bandwidth of the locked PLL is 13MHz. There is bandwidth peaking of 3.2dB at 8.3GHz in the PLL. Due to an inappropriately sized loop filter resistor, the damping factor is observed to be less than 0.5 causing overshooting at 10MHz. The phase noise at 1MHz offset is -99.7dBc/Hz, observed in the presence of a large noise floor and -120dBc/Hz beyond 100MHz.

The standalone PLL jitter was measured using an HP54754A oscilloscope, and an Agilent 8133A waveform generator as the input reference clock. The jitter of a 1.25GHz clock source from the 8133A was 1.13ps (RMS), 8.9ps (pk-pk). The measured integrated jitter (50kHz-80MHz) of the 10GHz PLL was 4mUI (RMS) or 0.4ps.

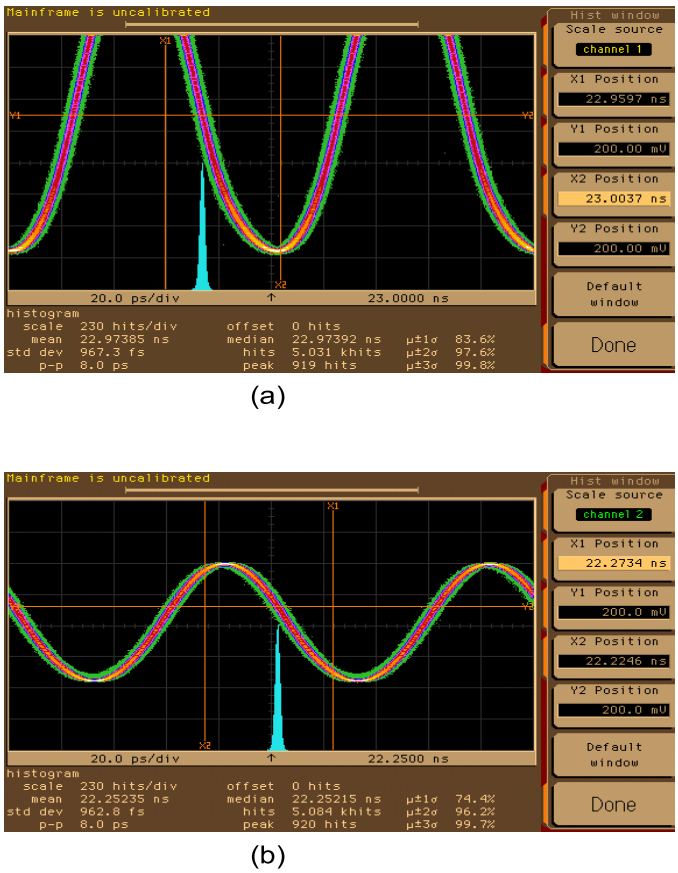

Figure 7.3: (a) Jitter Histogram of 10GHz Clock Waveform with Q=10 (b) Jitter Histogram of 10GHz Clock Waveform with Q=5

Figure 7.3 (a) shows that the jitter histogram of the 10GHz clock output(Q=10) is 0.97ps(RMS), 8ps(pk-pk). Since the PLL bandwidth of 13MHz is relatively high, much of the 8133A input reference jtter is passed directly to the PLL output. This suggests that a quieter input reference is needed to ascertain the true performance of the 10GHz PLL. One possibility for this is to use a quiet, low-phase noise RF sinusoidal generator, terminated to 50 Ohms at the chip input, after which the sinusoidal input is buffered up on-die. This was not done in measurement for the input clock since the board was not designed with its far-end terminated to Vcm=0.6V.

Figure 7.3 (b) shows the long-term jitter histogram of a low Q (Q=5), 10GHz LC clock output as 0.96ps (RMS), 8ps (pk-pk), which is comparable to the high Q (Q=10)

resonator – that is, the output jitter is limited by input clock reference jitter. While the power of the Q=5 resonator is 2.5x larger than the Q=10 inductor (15mW versus 6mW), since larger transistors are needed to compensate for the larger resonator loss), the area is 4x smaller (65um x 65um VS. 135um x 135um). This suggests that a lower Q oscillator can achieve comparable jitter performance (around 1ps RMS) with significantly less area, at the expense of slightly higher power dissipation. This observation is also confirmed by the LC oscillator designed by Intel [3], where a 10GHz LC VCO has a Q=5 and also very good jitter performance.

7.1.1 PSD Jitter Figures

Herzel and Razavi showed in [9] why this phenomenon might be true:

$$\Delta T_{PLL} = \frac{\sqrt{S_\Phi \Delta\omega}}{2\pi f_U} \frac{\Delta\omega}{\omega_0}$$

Where:

S_Φ = Power Spectral Density Of the Open Loop VCO

f_U = Phase Locked Loop Bandwidth

$\Delta\omega$ = Frequency offset from the Carrier

ω_0 = Frequency of the Main Carrier

	PSD(measured)	T_{PLL}
PLL0 (Q=10)	-99dBc/Hz	0.14psRMS
PLL1 (Q=5)	-92dBc/Hz	0.31psRMS

Figure 7.4: Conversion of Power Spectral Density to RMS Jitter

This equation shows the predicted phase locked loop jitter (RMS) given the power spectral density of the VCOs. Here, the predicted PLL output jitter for the Q=10 and Q=5 VCOs is 0.14ps and 0.31ps, respectively. As the experimental jitter is measured as 0.97ps(RMS), it is clear that the VCO jitter is dominated by the input reference clock and not the Q of the resonator. Notice that this analysis does not take into account the power supply rejection of the different resonators, as the higher Q VCO

will have better supply rejection. Along with the quiet supply measurements of the various PLLs with different Q oscillators, we also injected a 0.4V, 80MHz square-wave connected to a NMOS shorting switch (between VCO supply and ground), to determine the power supply rejection of each oscillator. Simulation shows that this NMOS shorting switch results in a 13ripple.

PLL Type	PLL Jitter (Quiet Supply)	PLL Jitter (Noisy Supply)
Q=10	1.295RMS, 11.1pk-pk	3.67RMS,22.2pk-pk
Q=5	1.262RMS, 11.1pk-pk	4.62RMS,22.2pk-pk

Figure 7.5: Measured Jitter Numbers with/without Power Supply Noise for LC Oscillator with Q=10 and Q=5

Figure 7.5 shows the measured jitter of the two LC oscillators (Q=10, Q=5), with both a quiet supply as well as the injected 80MHz square wave noise. Here it can be seen more clearly that with a quiet supply, the jitter is dominated by input clock reference jitter and not accumulated jitter, as suggested by the Herzel equation above. However, when noise is injected, the supply rejection of the Q=5 oscillator has worse RMS jitter, which can be explained by the fact that the size of the cross-coupled -gm inverters is almost 50 percent larger in size, and therefore exhibits higher coupling from the supply into the resonator. Since the implemented oscillator consists of back-to-back inverters without a tail current source node or supply voltage regulation, supply noise will also cause a large amount of jitter accumulation.

Open Loop VCO Phase Noise @ 1MHz	-97dBc/Hz
10GHz Jitter (RMS)	0.97ps
10GHz Jitter(pk-pk)	8.0ps
PLL Power	38.6mW
VCO Power	6mW
Tuning Range	1.14-1.31

Table 7.1: Measured PLL Performance

Table 7.1 summarizes the performance of the 10GHz PLL. The open loop VCO

phase noise at 1MHz frequency offset is -97dBc/Hz, resulting in a long term (5000 kilohits) jitter of 0.97ps(RMS) and 8ps(pk-pk). This noise is mostly input reference jitter limited, as the offchip Agilent 8133A clock generator has a measured jitter of 1.2ps(RMS). The VCO power is only 6mW due to the relatively large Q, while the total PLL power is 38.6mW, where more than half the power comes from the 1:4 frequency divider – the divider is a four-stage CML resistively loaded divider with significant current dissipation. The tuning range of this standalone PLL (no transmitter driver) is from 9.12-10.48 GHz, or 13 percent tuning range.

7.1.2 Measured Static Phase Offset of Multi-Phase Clock Generation

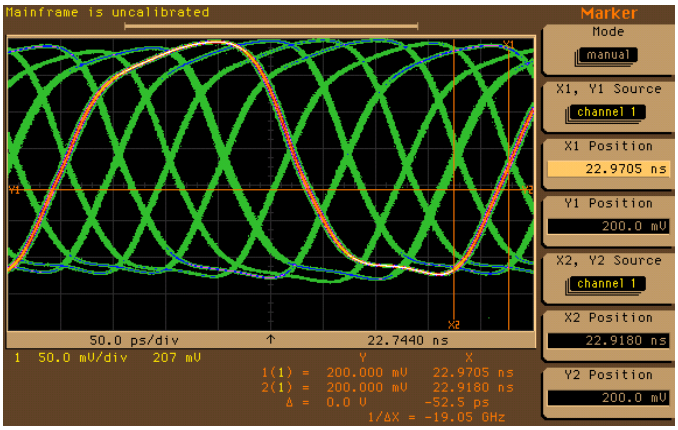

Figure 7.6: Measured Static Phase Offset of Eight Time-Interleaved Phases at 2.5GHz

In Figure 7.6, the output of eight 2.5GHz multi-phases are superimposed on to the same output buffer. These outputs would normally be used for a conventional output multiplexed transmitter architecture. The output phase step for these eight phases was 48.5ps, 51ps, 50ps, 60ps, 55.5ps, 61.5ps, 45.5ps, and 52.5ps, with a nominal phase step of 50ps. Even with care placed on matching capacitance and phase balance, the worst case eye will be shortened by 16ps, or 33 percent of a 20Gb/s eye width. Without interpolation and on-die feedback measurement/calibration to fix these phase imbalances, the transmitted eye diagram will not be symmetrically open

for all eight multiplexed data symbols.

7.1.3 Measured Eye Diagram

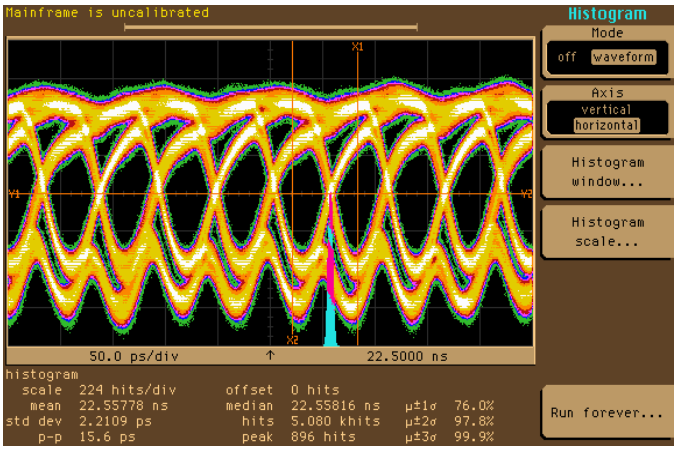

Figure 7.7: Measured 20Gb/s Eye Diagram

Figure 7.7 shows the eye diagram, at a data rate of 19.2Gb/s. Both outputs are superimposed on top of each other. The eye opening amplitude is approximately 105mV (single-ended) at the oscilloscope input. The measured jitter is 2.37ps (RMS) and 15.6ps (pk-pk). Measured static phase offset between eight consecutive bits is less than 2 ps. Notice that there is a 60mV voltage ripple seen at the bottom of the eye, predicted from simulation. This phenomenon occurs during the transition period of the sinusoidal complementary clocks, when neither tail NMOS is fully on, and the differential pair no longer acts with a constant bias current. Instead, the current drops during the transitions, causing the output nodes to pull up to the supply during the zero crossings of the clocks. The transmitter works from 16.32-19.2Gb/s, slower than our initial specification of 20Gb/s. This was actually observed in parasitic extraction simulation, showing the transmitter working from 16.8-19.6Gb/s. The cause of this is excessive capacitive loading of the complementary resonator clock nodes, with wide top thick metal(300um x 4um) connecting the frequency divider and the final 2:1 output multiplexer, as seen in Figure 7.8. More careful floorplanning of the layout would have placed the divider closer to the output multiplexer, reducing this clock

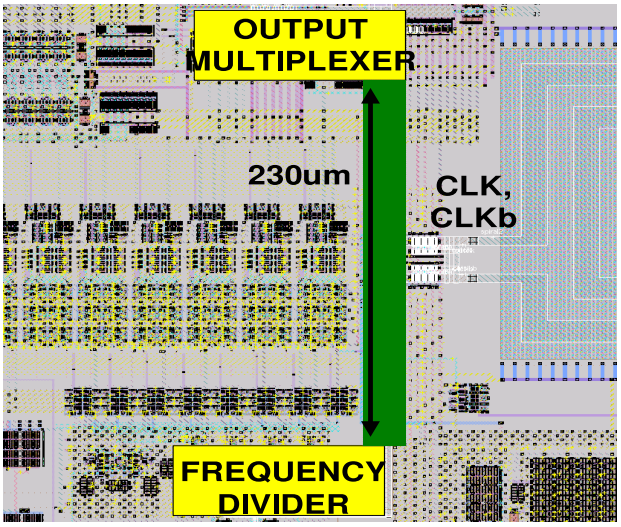

Figure 7.8: Long Clock Routing Resulting in Excessive Capacitive Loading of LC Oscillator

loading. Increasing PLL tuning range, such as using switch capacitor banks for the tank resonator, can also help alleviate this potential problem.

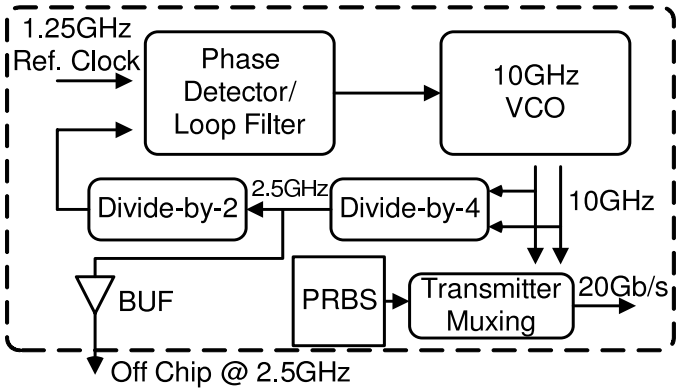

Figure 7.9: Measurement Setup of 2.5GHz Clock Jitter for Various Transmitter Configurations

We attempted to determine the effect of data-dependent kickback on PLL jitter accumulation by measuring the transmitter in various configurations. This was done by measuring the 2.5GHz divided clock coming from the PLL itself, as in Figure 7.9. If poor isolation exists between the transmitter VCO and the transmitter latches/muxes, we would expect the 2.5GHz PLL clock to have the largest jitter when PRBS data

Various Configurations	Jitter (rms)	Jitter (pk-pk)
Transmitter Off, PRBS Off	1.25ps	10ps
Transmitter On, PRBS Off (data:00000000)	1.25ps	8.9ps
Transmitter On, PRBS Off (data:00001000)	1.21ps	10ps
Transmitter Off, PRBS On	2.3ps	16.7ps
Transmitter On, PRBS On	2.25ps	15.6ps

Table 7.2: Measured PLL Clock Jitter for Various Transmitter Configurations

is transmitted. Table 7.2 illustrates the results of our measurements. When only the PLL is enabled, and the PRBS and transmitter multiplexing disabled, the measured jitter is 1.21ps, 10ps (RMS,pk-pk).

Active Area	Transmitter: 0.23mm^2 PLL: 0.073mm^2 Inductor: 0.017mm^2
Power (for 105mV differential swing)	Transmitter: 165mW PLL: 33mW LC-VCO: 6.4mW
Transmitter Tuning Range	16.32 – 19.2 GHz
Transmitter Output Jitter	2.37ps RMS, 15.2ps pk-pk
Static Phase Offset	< 2 ps

Table 7.3: First Generation Transmitter Performance Summary

With the PLL and transmitter muxing on (sending all zeroes) and the PRBS off, the jitter is 1.25ps, 8.9ps (RMS, pk-pk), illustrating no significant change in jitter[1]. With the same configuration, while sending a lone zero, the jitter again doesn't change, with a measured jitter of 1.25ps, 10ps (RMS, pk-pk). The jitter of the 2.5GHz clock increases substantially (2.3ps, 16.7ps) when the PLL and the digital PRBS are both on–even when the transmitter output is turned off. This illustrates that the dominant source for the increased jitter is the power supply noise, either/both in the divide-by-4 or the buffer chain to the output pin. Finally, when the PLL, PRBS, and the transmitter on, sending random 20Gb/s data at the output, the measured jitter is 2.25ps, 15.6ps. Since this measured jitter is roughly the same as the situation with the PRBS on and transmitter off, we can infer that the increase

[1] 1.21ps (RMS) and 1.25ps (RMS) are the quantization steps for the oscilloscope measurements.

in measured jitter is dominated by digital power supply jitter caused by the PRBS. Measured results using a spectrum analyzer do not uncover any other conclusions. These results suggest that such data-dependent accumulation is less significant and dominated by digital power supply noise. However, more work must be done to prove that data-dependent jitter accumulation is negligible for various oscillator topologies and quality factor values.

7.2 Prototype #2: 20Gb/s Transceiver

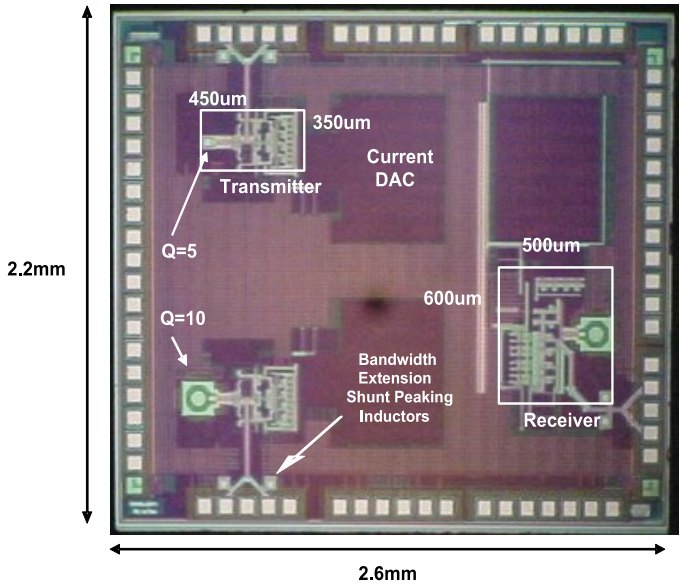

Figure 7.10: Second Generation Transmitter with Corresponding Receiver

Figure 7.10 shows a 20Gb/s transceiver built in ST Microelectronics 0.13um CMOS process. The image above shows the fabricated silicon–die size is 2.6mm x 2.2mm. Notice that the Q=5 inductor consumes less than one-third the area of the Q=10 inductor. Programmable current DACs of 8 binary bits are used to control the 12 individual current sources of each transmitter. Finally, 0.3nH single turn spiral inductors are used as shunt peaking inductors compensating for channel attenuation of the transmitters / receivers at the 20Gb/s differential outputs/inputs.

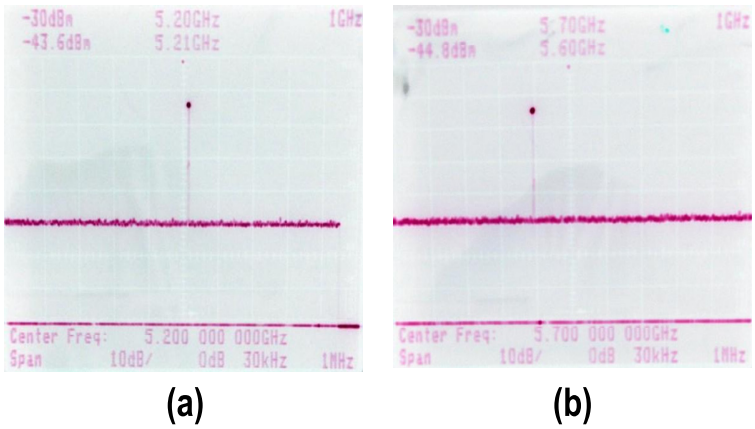

Figure 7.11: (a) 5.2GHz Divider Output (b) 5.6GHz Divider Output

The image of 7.11 shows the power spectrum from the output of the 10GHz divider. With the divider working between the frequencies 5.2-5.6GHz, this results in an operating frequency of the LC-VCO working from 10.4GHz - 11.2GHz and a tuning range of 7.5 percent. Unfortunately, the design was unable to be verified for performance beyond the operation of the LC-VCO and the 10GHz divider. The reason for this is two-fold. First, due to power/ground DC resistive losses, the maximum current capability of the current sources was 30 percent less than simulated from post-layout extraction. As the four-phase 10GHz divider and the eight-phase 5GHz dividers are resistively loaded CML flip-flops, they are very sensitive to voltage swing and therefore, current consumption. Since the eight-phase 5GHz dividers consist of four CML flip-flops in an inverting ring topology, the total current consumption is around 20mA for this divider. Due to DC supply losses, the voltage swing on these internal nodes is less than the optimum that is observed in simulations. Therefore, the next stage dividers/buffers cannot be driven effectively due to the poor power grid.

Exacerbating this particular problem is that the circuit design of the divider circuits is sub-optimum. Figure 7.12 illustrates the path of the clock from 10GHz to the lower frequency of 1.25GHz, which is the input into the phase comparator. The two complementary phases of the 10GHz clock first go to a CML master/slave divide-by-2 divider, creating four low swing clock signals. These low swing clock signals go

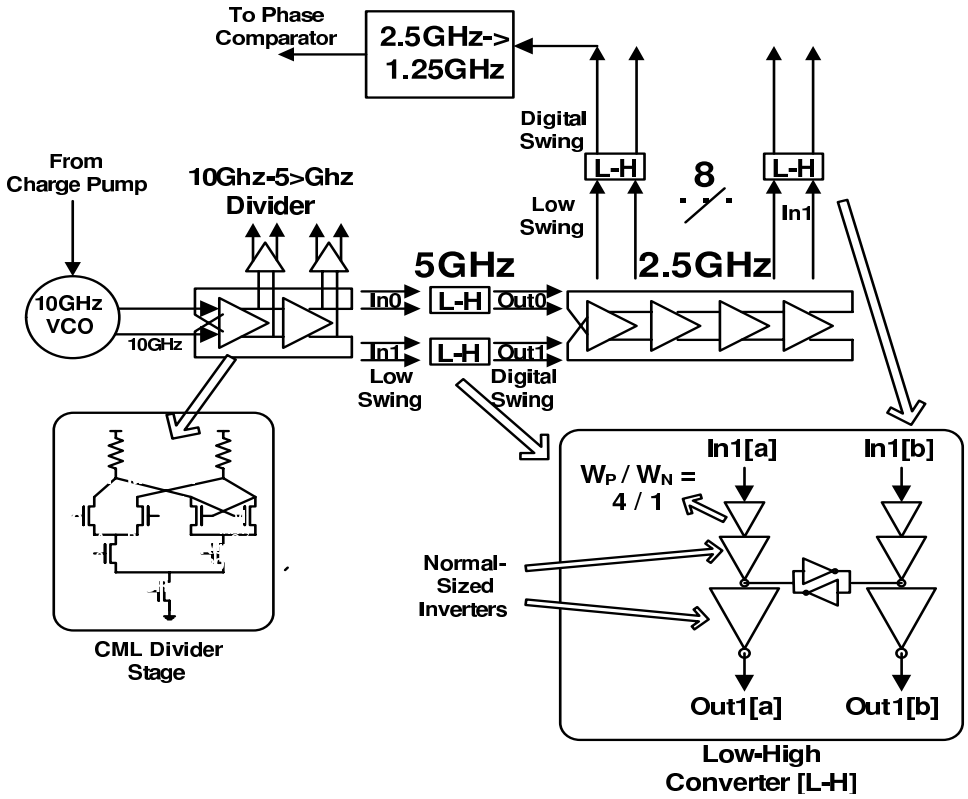

Figure 7.12: 10GHz Clock Divider Description and Circuit Design

through low-to-high converters (which as will be seen, are designed inappropriately) to four digital level clock phases. These four digital clock phases are then sent to another CML master/slave divide-by-2 divider, creating eight low swing phases at 2.5GHz. Again, these eight low swing phases are low-to-high converted to digital level, at which point one set of these phases is used in a simple, digital divide-by-2 divider, creating the 1.25GHz PLL feedback clock. The problem arises due to the design of the low-to-high converter. As seen in Figure 7.12, these L-H converters consist of an input buffer whose PMOS is sized twice as large as a normal sized inverter gate, essentially making the PMOS pullup stronger than the NMOS pulldown. The reason for this is that the preceding CML divider stage with resistively loaded outputs has a differential swing between Vdd(1.2V) and 0.5V. Since the CML output doesn't pull all the way to GND, sizing the PMOS gate larger can compensate for the reduced swing, at least in simulation. This digital L-V converter saves power substantially

compared to a CML output buffer.

This type of low-to-high converter is problematic, however, since it varies widely with process corner, and with the swing/common-mode voltage of the CML dividers. However, since there are resistive losses in the power supply, and the maximum current to the CML dividers is reduced by around 30 percent, the correct common mode and signal swing of the CML output cannot be achieved, especially to the inputs of the 2.5GHz dividers. Ironically, is it more difficult to build the eight-phase 5GHz divider than the four-phase 10GHz divider, since the spacing between phases is only 50ps, meaning that the bandwidth of the divider (and hence the current consumption) has to be optimum.

Another reason why this low-to-high converter operated poorly is that its switching midpoint depends heavily on the value of the supply voltage. During experimental testing, when it is observed that the current sources are 30 percent slower than expected, one possible solution is to raise the supply of the entire chip – for example, Vdd is raised from 1.2V to 1.8V. While this increases the swing and current through the CML dividers, effectively producing an effective 5GHz - 2.5GHz divider, raising the supply also affects the midpoint switching point of the following low-to-high converters, as this low-to-high converter is built simply sizing the inverter PMOS strength compared to the NMOS. Thus, raising the supply voltage does not improve the performance of passing the clock through these low-to-high converters. The solution to this problem is to use a more standard low-to-high converter, such as an NMOS pulldown with diode-connected PMOS loads.[2] Circuits in the divider should be built to be common-mode insensitive, so that exact voltage swings of the dividers are not necessary, and the design can be more robust to process variation. The implemented low-to-high converter with skewed PMOS/NMOS ratio is extremely sensitive to inverter tripping point, and therefore process mismatch.

[2]These circuits are not trivial to build, however, since the oscillation frequency is 5GHz, and it is difficult to achieve full-swing CMOS voltages at 5GHz with 0.13um CMOS. Moving to a more advanced technology, such as 90nm or 65nm, would help relax the design difficulties.

Chapter 8

Conclusion

As CMOS process scaling continues to increase the amount of computation possible on a single die, the ability to communicate information between different chips (using serial links) becomes critically important. For example, with future multi-core processors having tens to hundreds of independent threads, the memory and off-chip interconnect bandwidth need to also increase to supply data to these individual cores. As the clock frequency and data rate for these high speed links increases, clock timing precision becomes one of the key constraints to low bit error rate as the timing margin is reduced.

Two main causes for timing uncertainty are static phase offset and power supply induced jitter caused by post-PLL clock buffers. Static phase offset caused by device, process, and layout mismatch in the clock distribution results in DC timing inaccuracy. Power supply noise – arising from large digital circuit logic inducing supply current transients thereby causing power supply voltage variation – affects the timing delay of clock buffers. As a result, clock buffer timing jitter manifests as high frequency clock jitter and is therefore extremely difficult to alleviate using relatively low bandwidth techniques (i.e. low-bandwidth PLLs or linear power supply regulators).

Conventional serial links typically use a single, central, low jitter VCO clock source with optimized phase noise characteristics. However, this clean clock source needs to exit the synthesizer and eventually drive multiple serial link transmitters and receivers, requiring a buffer chain to increase drive strength. These post-PLL buffers

introduce possible sources of process mismatch and therefore static phase offset. Likewise, each buffer's delay is proportional to the power supply voltage, such that significant digital logic noise will create high frequency induced clock jitter at the clock buffer output. These sources of jitter, static phase offset and power supply induced jitter, are shown to increase with CMOS scaling, as process mismatch and power supply noise are expected to worsen in the future.

As it is observed that post-PLL clock buffers introduce timing uncertainty, this thesis describes a new clocking architecture for serial link transceivers using direct resonant VCO drive of multiplexers/demultiplexers. This scheme eliminates post-PLL buffers between the frequency synthesizer and the front-end multiplexer/demultiplexer, significantly reducing power consumption, static phase offset, and power supply induced jitter. However, as there is no isolation between the VCO and the transmitter multiplexer (or receiver demultiplexer), there exist some serious possible side-effects that must be considered. First, as the load capacitance is directly subsumed into the resonator load, the effect of Q degradation and reduced tuning range due to this additional capacitance is observed experimentally and also analyzed in simulation. Second, the residual static phase offset and power supply induced jitter of the LC oscillator are considered through simulation and experimentally verified. It is shown that static phase offset due to process, device, and capacitive mismatch are reduced substantially by the intrinsic characteristics of the LC resonator. This same narrowband property of the LC resonator also result in significant rejection of power supply induced jitter. Third, this "non-isolation" between the VCO and the front-end multiplexer/demultiplexer can result in increased accumulated phase-locked loop jitter, as the transmission of random data sequences can cause data-dependent charge fluctuation of the multiplexer/demultiplexer capacitance. Since this multiplexer/demultiplexer capacitance is subsumed directly into the LC resonator, there is the possibility of increased synthesized jitter. This effect of data-dependent accumulated jitter is mitigated by the use of fully differential circuit topologies, making any charge injection common-mode to both nodes of the LC-VCO. Simulation shows that the increased open loop VCO jitter over many clock cycles is less than 4ps.

To prove the validity of this direct modulation scheme for serial link transmitters

and receivers by using a LC-VCO, two 20Gb/s transmitters and one receiver prototype were designed in 0.13um CMOS. The first transmitter prototype achieves a 20Gb/s data rate by first implementing 4:1 10Gb/s multiplexers, two 10Gb/s analog latches, followed by a final 2:1 20Gb/s multiplexer. To overcome bandwidth limitations and latch hysteresis, a 10Gb/s transmitter pre-emphasis technique is implemented for the 10Gb/s multiplexers. Experimental results show a clean 20Gb/s transmitted eye diagram, with pk-pk jitter less than 16ps. Through experimental measurements, it is shown that data-dependent kickback VCO jitter is insignificant when compared with power supply induced jitter. The second implemented transmitter implements a 2:1 tree-multiplexing structure that has relaxed bandwidth constraints and does not require a final analog latch, therefore reducing design complexity. This second transmitter also implements a two-tap pre-emphasis filter at 20Gb/s, to compensate for channel loss at this high data rate. Post-layout simulation shows an almost 2x increase in signal-to-noise using this FIR filter in a low-pass channel. Finally, a 20Gb/s receiver prototype is implemented, where the front-end demultiplexer uses direct-drive sampling from a LC-VCO. After retiming to 10Gb/s analog data, the data is further demultiplexed into 5Gb/s analog voltages before quantization, relaxing the bandwidth constraints of the quantizers. Post-layout simulations show correct operation and demultiplexing of data using this new demultiplexing approach.

8.1 Future Work

Using this direct LC-VCO architecture in deep submicron CMOS has many advantages in regards to power supply rejection, low power consumption, low static phase offset, and low intrinsic jitter. While this architecture is seen for 20Gb/s transceivers and receivers, this architecture can be applied to other areas. For example, we are currently in the process of designing a 5-bit, 5GS/s, 50mW flash ADC using direct drive VCO clocking. Preliminary results suggest a 5x-10x reduction in power consumption (since half of the power consumption in conventional flash ADCs is consumed by clock power). This technique can also be extended to multi-gigahertz, low-resolution DACs and even further in the future, multi-gigahertz oversampling continuous-time

sigma-delta converters. Low resolution data converters have many applications, such as channel equalization for serial link transceivers, ultra-wideband RF systems, and software defined radio. Ultimately, solving these timing precision issues will enable designers to continue capitalizing on the increased transistor gain/bandwidth and reduced power consumption benefits of CMOS scaling.

Bibliography

[1] *International Technology Roadmap for Semiconductors*. 2005.

[2] J. Kim. Stanford Doctoral Thesis: Design of CMOS Adaptive-Supply Serial Links. 2002.

[3] B. Casper et. al. A 20Gb/s Embedded Clock Transceiver. *ISSCC Dig. Of Tech. Papers*, February 2006.

[4] Thomas L. Lee. The Design of CMOS Radio-Frequency Integrated Circuits. *Cambridge University Press, 2nd Edition*, December 2003.

[5] L. Lee, D. Weinlader, and C.K. K. Yang. A Sub-10-ps Multiphase Sampling System Using Redundancy. *IEEE Journal of Solid-State Circuits*, Vol. 41, No. 1, January 2006, pp. 265-273.

[6] J. Buckwalter and A. Hajimiri. Analysis and Equalization of Data-Dependent Jitter. *IEEE Journal of Solid-State Circuits*, Vol. 41, No. 3, pp. 607-619, March 2006.

[7] E.M. Cherry and D.E. Hooper. The Design of Wide-Band Transistor Feedback Amplifier. *Inst. Elec. Eng. Proc.*, Vol. 110, no. 2, pp. 375-389, February 1963.

[8] S. Palermo, A. Emamai, and M. Horowitz. A 90nm CMOS 16Gb/s Transceiver for Optical Interconnects. *ISSCC Dig. Of Tech. Papers*, February 2007.

[9] F. Herzel and B. Razavi. A Study of Oscillator Jitter Due to Supply and Substrate Noise. *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, Vol. 46, Issue 1, pp. 56-62. January 1999.

[10] T. Lee and A. Hajimiri. Oscillator Phase Noise: A Tutorial. *IEEE Journal of Solid-State Circuits* Vol. 35, No. 3, March 2000, pp. 326-336.

[11] U. Singh and M. Green. A 34Gb/s Distributed 2:1 MUX and CMU Using 0.18um CMOS. *IEEE Journal of Solid-State Circuits* Vol. 41, No. 9, September 2006, pp. 2067–2076.

[12] M.J. E. Lee et. al. An 84-mW 4Gb/s Clock and Data Recovery Circuit for Serial Link Applications. *VLSI Circuits Symposium* Kyoto, Japan, June 2001.

[13] S. Sidiripolous and M. Horowitz. A Semidigital Dual Delay-Locked Loop. *IEEE Journal of Solid-State Circuits* Vol. 32, No. 11, November 1997, pp. 1683-1692.

[14] L. Wu and W. C. Black Jr. A Low-Jitter Skew-Calbrated Multi-Phase Clock Generator for Time-Interleaved Applications. *2001 IEEE International Solid State Circuits Conference* February 2001, pp. 396-397.

[15] M.J.E. Lee. An Efficient I/O and Clock Recovery Design for Terabit Integrated Circuits. *Stanford University Doctoral Thesis* August 2001.

[16] R. Navid, et. al. Minimum Achievable Phase Noise of RC Oscillators. *IEEE Journal of Solid-State Circuits* Vol. 40, No. 3, March 2005, pp. 630-637.

[17] E. Egazi and A. Abidi. A Filtering Technique to Lower LC Oscillator Phase Noise. *IEEE Journal of Solid-State Circuits* Vol. 36, No. 12, December 2001.

[18] J. G. Maneatis. Low-Jitter Process-Independent DLL and PLL Based on Self-Biased Techniques. *IEEE Journal of Solid-State Circuits*, Vol. 31, No. 11, November 1996, pp. 1723–1732.

[19] L. Wu and W. C. Black Jr. A Low jitter Skew-Calibrated Multi-Phase Clock Generator for Time-Interleaved Applications. *ISSCC Dig. Of Tech. Papers*, February 2001, pp. 396–397.

[20] Chan-Hong Park, Ook Kim, Beomsup Kim. A 1.8-GHz Self-Calibrated Phase-Locked Loop with Precise I/Q Matching. *IEEE Journal of Solid-State Circuits*, Vol. 36, No. 5, May 2001, pp. 777–783.

[21] M. Mansuri and C. K. Ken Yang. A Low-Power Low-Jitter Adaptive-Bandwidth PLL and Clock Buffer. *ISSCC Dig. Of Tech. Papers*, February 2003, pp. 430–505.

[22] A. J. Drake KJN, T. Y. Nguyen, J. L. Burns, and R. B. Brown. Resonant Clocking Using Distributed Parasitic Capacitance. *IEEE Journal of Solid-State Circuits*, Vol. 39, No. 9, September 2004, pp. 1520–1529.

[23] T. H. Lee and A. Hajimiri. Oscillator Phase Noise: A Tutorial. *IEEE Journal of Solid-State Circuits*, Vol. 35, No. 3, March 2000, pp. 325–336.

[24] P. Heydari. Analysis of the PLL Jitter Due to Power/Ground and Substrate Noise. *IEEE Trans. on Circuits and Systems I*, Vol. 51, No. 12, December 2004, pp. 2404–2416.

[25] J. Cao, et. al. OC-192 Transmitter and Seceiver in standard 0.18-/spl mu/m CMOS. *IEEE Journal of Solid-State Circuits*, Vol. 37, No. 12, December 2002, pp. 1767–1780.

[26] L. Henrickson, et. al. Low-Power Fully Integrated 10-Gb/s SONET/SDH Transceiver in 0.13-um CMOS. *IEEE Journal of Solid-State Circuits*, Vol. 37, No. 10, October 2003, pp. 1595–1601.

[27] J. Savoj and B. Razavi. A 10-Gb/s CMOS Clock and Data Recovery Circuit with a Half-Rate Linear Phase Detector. *IEEE Journal of Solid-State Circuits*, Vol. 36, No. 5, May 2001, pp. 761-768.

[28] S.B. Anand and B. Razavi. A 2.75Gb/s CMOS Clock Recovery Circuit with Broad Capture Range. *International Solid-State Circuits Conference*, February 2001, pp. 214-215.

[29] B. Razavi. Prospects of CMOS Technology for High-Speed Optical Communications. *IEEE Journal of Solid-State Circuits*, vol. 37, No. 9, September 2002, pp. 1135–1145.

[30] J. Lee and B. Razavi. A 40Gb/s Clock and Data Recovery Circuit in 0.18um CMOS Technology. *IEEE Journal of Solid-State Circuits*, Vol. 38, No. 12, December 2003, pp. 2183–2190.

[31] Jaeha Kim, Bong-Joon Lee, Moon-Sang Hwang, Hyung-Rok Lee, Namhoon Kim, Deog-Kyoon Jeong, Wonchan Kim. Circuit Techniques for a 40Gb/s Transmitter in 0.13um CMOS. *2005 Digest of International Solid State Circuits Conference 2005*, February 2005, pp. 150–151.

[32] E. Alon, V. Stojanovic, and M. Horowitz. Circuits and Techniques for High-Resolution Measurement of On-Chip Power Supply Noise. *Journal of Solid State Circuits*, April 2005.

[33] M. Pelgrom, et. al. Matching Properties of MOS Transistors *Journal of Solid State Circuits*, Vol. 24, Issue 5, October 1989, pp. 1433-1439.

[34] A. M. Niknejad. Modeling of Passive Elements with ASITIC. *Radio Frequency Integrated Circuits Symposium*, June 2002, pp. 303–306.

[35] D. Weinlader, et. al. An Eight Channel 36G Samples/s CMOS Timing Analyzer. *2000 Digest of International Solid State Circuits Conference 2000*, February 2000, pp. 170-171.

[36] K. Yamaguchi, et. al. 12Gb/s Duobinary Signaling with x2 Oversampled Edge Equalization. *2005 Digest of International Solid State Circuits Conference 2005*, February 2005, pp. 70–71..

[37] W. Ellersick, et. al. A Serial-Link Transceiver Based on 8 Gsample/s A/D and D/A Converters in 0.25um CMOS. *2003 Digest of International Solid State Circuits Conference 2003*, February 2003, pp. 150–151.

[38] V. Stojanovic, A. Ho, B. Garlepp, F. Chen, J. Wei, G. Tsang, E. Alon, R. Kollipara, C. Werner, J. Zerbe, and M. Horowitz. Autonomous Dual-Mode PAM2/4 Serial Link Transceiver with Adaptive Equalization and Data Recovery. *IEEE Journal of Solid-State Circuits*, Vol. 38, No. 4, April 2005, pp. 2181–2190.

[39] R. Farjad-Rad, C.-K. K. Yang, M. Horowitz, T. Lee. A 0.3-um CMOS 8-Gb/s 4-PAM serial link transceiver. *IEEE Symposium on VLSI Circuits*, June 1999, pp. 150–151.

[40] Patrick Chiang, William J. Dally, Ming-Ju Edward Lee, Ramesh Senthinathan, Yangjin Oh, and Mark Horowitz. A 20Gb/s 0.13um CMOS Serial Link Transmitter Using an LC-PLL to Directly Drive the Output Multiplexer. *2004 IEEE Symposium on VLSI Circuits*, June 2004, pp. 272–275.

[41] Patrick Chiang, William J. Dally, Ming-Ju Edward Lee, Ramesh Senthinathan, Yangjin Oh, and Mark Horowitz. A 20Gb/s 0.13um CMOS Serial Link Transmitter Using an LC-PLL to Directly Drive the Output Multiplexer. *IEEE Journal of Solid-State Circuits*, Vol. 40, No. 4, April 2005, pp. 1004–1011.

[42] Patrick Chiang, William J. Dally, Ming-Ju Edward Lee. A 20Gb/s 0.13um CMOS Serial Link. *Hotchips 2002*, August 2002, pp. 272–275.

[43] Tse-Yu Yeh, Deborah T. Marr, and Yale N. Patt. Increasing the instruction fetch rate via multiple branch prediction and a branch address cache. *Proceedings of the 7th International Conference on Supercomputing*, pages 67–76, Tokyo, Japan, July 1993. ACM Press.

[44] K. Chang, S. Pamarti, K. Kaviani, E. Alon, X. Shi, T. Chin, J. Shen, G. Yip, C. Madden, R. Schmitt, C. Yuan, F. Assaderaghi, and M. Horowitz. Clocking and Circuit Design for a Parallel I/O on a First-Generation CELL Processor. *IEEE International Solid-state Circuits Conference*, February 2005, pages 150–151.

[45] W.J. Dally and J.W. Poulton. *Digital Systems Engineering*, Cambridge University Press, Cambridge, MA, 1998.

[46] C.-K. Yang and M. Horowitz. A 0.8um CMOS 2.5Gb/s Oversampling Receiver and Transmitter for Serial Links. *IEEE Journal of Solid-State Circuits*, Vol. 31, No. 12, December 1996, pp. 2015–2023.

[47] M. Horowitz, C.-K. Yang and S. Sidiropoulos. High-Speed Electrical Signaling: Overview and Limitations. *IEEE Micro*, January 1998, pp. 12–24.

[48] H. Hatamkhani and Chih-Kong Ken Yang. Power Analysis for High-Speed I/O Transmitters. *Digest of Technical Papers on the 2004 Symposium on VLSI Circuits*, June 2004, pp. 142–145.

[49] J. Kim and M. Horowitz. Adaptive Supply Serial Links With Sub-1-V Operation and Per-Pin Clock Recovery. *IEEE Journal of Solid-State Circuits*, Vol. 39, No. 11, November 2002, pp. 1403–1413.

[50] M.-J. E. Lee, W. J. Dally, P. Chiang. Low-Power Area-Efficient High-Speed I/O Circuit Techniques. *IEEE Journal of Solid-State Circuits*, Vol. 35, No. 11, November 2000, pp. 1591–1599.

[51] R. Farjad-Rad, et. al. 0.622-8.0 Gbps 150mW Serial IO Macrocell with Fully Flexible Preemphasis and Equalization. *Digest of Technical Papers on the 2003 Symposium on VLSI Circuits*, June 2003, pp. 63–66.

[52] K.L.J. Wong, H. Hatamkhani, M. Mansuri, and C.K.K. Yang. a 27-mW 3.6 Gb/s I/O Transceiver. *IEEE Journal of Solid-State Circuits*, Vol. 39, No. 4, April 2004, pp. 602–612.