SINGLE-ENDED 16x8 GBPS DATA BUS IN 90NM CMOS

By

SAURABH MANDHANYA

A thesis submitted in partial fulfillment of

the requirements for the degree of

MASTER OF SCIENCE IN ELECTRICAL ENGINEERING

WASHINGTON STATE UNIVERSITY

School of Electrical Engineering and Computer Science

DECEMBER 2009

To the Faculty of Washington State University:

The members of the Committee appointed to examine the thesis of SAURABH MANDHANYA find it satisfactory and recommend that it be accepted.

_____

George La Rue, Ph. D., Chair

_____

Deuk Heo, Ph. D.

_____

Patrick Pedrow, Ph. D.

ACKNOWLEDGEMENT

First I would like to thank the *Center for Design of Analog-Digital Integrated Circuits* (CDADIC) for funding the 'Single-Ended High Speed Data Bus' project.

I would like to acknowledge the contribution by the faculty at Washington State University who provided me with the knowledge and skills required to complete this thesis. I would like to thank the members of my committee, Prof. George La Rue, Prof. Deuk Heo, and Prof. Partha Pande, who have taken time out of their busy schedule to review this thesis. A lot of credit goes to Prof. George La Rue who has been my advisor for the past two years and whose help and guidance at every step have made this thesis possible.

I would like to thank Prasanna Upadhyaya and Wei Zheng for teaching me many of the skills necessary to be a part of the analog circuits design group. I would like to thank Ding Ma for all his help in the layout of the final design. I would like to thank Peter Oshroff for his work on the coding schemes and the data generation and collection parts which make the testing of this design possible.

I want to thank my parents, Shree Gopal Mandhanya and Jyoti Mandhanya for supporting me in all my endeavors and making me the person that I am today. I dedicate the completion of this degree to them.

SINGLE-ENDED 16x8 GBPS DATA BUS

IN 90 NM CMOS

ABSTRACT

By Saurabh Mandhanya, M.S.
Washington State University
DECEMBER 2009

Chair: George S. La Rue

This thesis presents the design of a high-speed single-ended 16 channel data bus for chip-to-chip communication with each channel capable of providing up to 8Gbps of throughput and adding up for a total bus throughput of 128 Gbps. By using singe-ended signaling instead of the commonly used differential signaling, each channel requires only one I/O pin and one physical channel for transmitting data compared to two I/O pins and two channels required for differential signaling. Although some overhead is added, the single-ended bus still requires only 21 physical channels for transmission instead of the 32 that would be required for a differential bus.

A printed circuit board was fabricated to model the physical channels and to decide the best possible arrangement for the channels. Some design features commonly used for high-speed data communication like pre-emphasis and equalization are employed. The equalization schemes are tunable for each channel and therefore allows for different channel lengths. Also, some issues associated with single-ended signaling, like higher crosstalk, absence of common-mode rejection and power supply bounce, were investigated and methods to resolve them were implemented to ensure reliable communication.

TABLE OF CONTENTS

Page

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1

1. Introduction

With the scaling of CMOS technologies, semiconductor integrated circuits(IC's) have become progressively faster and denser. At the same time the improvement in Input/Output (I/O) capacities for IC's has been slow. This has created a bottleneck over time where it is difficult for I/O's to handle the vast amount of data that needs to be sent and received by an IC. The situation has been further worsened by the advent of multi-core and multi-Gigahertz (GHz) processors which along with other IC's on the system, demand very high I/O capacities for communicating with each other.

Another problem that arises due to technology scaling and the integration of several functions on the same IC is that the number of I/O pins available remains limited. These IC's, however, need more pins to implement the increased functionality. This provides a great incentive to look for ways to reduce the number of pins used for the purpose of data communication.

The data buses used for chip-to-chip communication today have speeds of up to a few Gigabits per second (Gbps) which needs to increase to close to 10 Gbps to keep up with the current technology trends. Also, it is common practice to use differential signaling for data buses operating at such high speeds where an information bit is transferred over two wires in the form of two complementary signals. Differential buses are used because of their inherent noise immunity and common-mode rejection. They are also somewhat immune to crosstalk problems that become very prominent at high speeds. However, the disadvantage of using differential

buses is that two I/O pins and two I/O lines are needed for every signal that needs to be transferred.

Single-ended signaling is the alternative to differential signaling for data buses which uses just one I/O pin and one line to transfer a signal. Although this saves a lot of channels and I/O pins, it lacks the noise immunity and robustness of differential signaling. Crosstalk also needs to be taken care of since it becomes a major problem. Hence, designing a single-ended data bus at similar speeds is much more complex, as these issues have to be resolved.

But some physical limitations exist for the channel itself, irrespective of the type of signaling. The physical channels, which in the case of chip-to-chip communication are usually copper traces on FR-4 printed circuit boards (PCB's), have frequency dependent losses which become significantly high beyond 2 Gbps. Also, the skin effect comes into play, which distorts the signal and spreads it out in time leading to Inter-Symbol Interference (ISI). As a result, the signal starts degrading as it travels down the channel, and for practical channel lengths signal integrity gets severely impaired [3]. Several types of equalization techniques have been used historically to overcome these losses which basically try to compensate for the lossy characteristics of the channel.

Another issue that needs to be resolved is crosstalk. When I/O pins and the channels are in close proximity, some energy gets coupled from each channel to every other channel that is close to it. The amount of coupling increases dramatically with frequency and it becomes the main limiting factor for achieving very high data rates. Crosstalk can be kept low if the separation between channels is increased but this would increase the area of the bus.

1.1 Project Description

In this dissertation, a 16 bit wide, single-ended data bus design is proposed which can achieve data transfer speeds of up to 8 Gbps per channel, which is nearly equal to the existing differential buses in use. The final design uses 21 channels over a FR-4 PCB to transfer 16 bits of data hence saving 9 PCB traces and 9 I/O pins over a differential design. The design was focused on resolving the issues that are inherent in single-ended signaling:

1. Power Supply Bounce: Single-ended I/O's are not necessarily balanced at all times, meaning that they can have more lines that are high than low, or vice versa, at any given time. This causes more current to be drawn from the supplies in some cases and less in other cases. As a result the power supplies may not remain at constant voltage potentials. This in turn can cause variations over the single-ended bus and even worse, it could affect other signals on the chip that are using the same power supplies. A simple coding scheme is used to overcome this where the 16 data bits are encoded on to 21 bits which are balanced to within one. Perfect balancing is then achieved by using a dummy driver.
2. Crosstalk: Since single-ended buses are not as immune to noise as differential buses, and crosstalk is the main limiting factor in achieving high data rates beyond 5 Gbps, active crosstalk cancellation techniques are essential. Although crosstalk occurs from every channel in the bus to every other channel, it was found that for some transmission line layouts only the crosstalk from the nearest neighbors (NN crosstalk) was significant and the crosstalk from the other channels can be ignored.
3. Common-mode Rejection: Since the common-mode level of a signal is essential in determining whether it is a low or high signal, a way to accurately determine this is required. This is not needed in differential signaling since there are two complementary

signals for every transmitted bit and the common-mode level lies in between those signals at all times. Based on this observation, the simplest way to acquire common-mode information is by including a differential bus and extracting the common-mode voltage. Also, instead of one, two differential buses are included in the single-ended bus and the common-mode signals from these channels are distributed to nearby single-ended signals on the bus.

Apart from the features stated above for single-ended signaling, the bus also features a tunable equalization scheme for each channel at both the transmitting and receiving ends for compensating for line losses. Also, a variable delay line at the receiver end allows the receiver clock to be adjusted individually for each channel. This combined with the tunable equalization scheme allows each channel in the bus to have a slightly different length.

1.2 Organization of Thesis Dissertation

This dissertation is organized to maintain a smooth flow of ideas and help the reader understand it clearly. Chapter 2 comes next and it provides some background information related the design of the bus. It provides a brief overview of lossy channels, equalization schemes and crosstalk cancellation methods. Chapter 3 discusses the system level architecture of the bus and the testing methodology. Chapter 4 goes into the details of the individual modules at the transistor level. Chapter 5 illustrates the simulation results and the layout of the circuit. Chapter 6 is a brief summary of the work, along with a discussion of the major contributions of this work and recommendation for future work.

Chapter 2

2. Background

2.1 Signal Integrity

Signal Integrity (SI) is a measure of the quality of electrical signals that are transferred over a data communication channel. For communication over a conductor, the data stream is transmitted as a voltage or a current waveform. This works perfectly if the conductor is short in length and the speed of transmission is not very high. However, for higher speeds and longer distances, several physical effects cause the signals to degrade resulting in the possibility of data errors when the original data cannot be recovered at the receiver. Hence, analyzing and maintaining SI is a very important task. The main issues that are of concern related to SI are line losses, crosstalk and power supply bounce. Another issue associated with single-ended signaling is that it does not have any inherent common-mode rejection. The next few sections explore these issues in more detail and the possible approaches to resolve them.

2.2 Channel Characteristics

Standard backplanes in use today have a FR-4 substrate with copper traces as the channels for communication. These backplanes were designed to operate up to around 1 Gbps. In recent years, advanced materials, like those manufactured by Rogers, have been introduced but since they are more expensive, their usage is still rather limited. As a result, the need arises for using the existing FR-4 material for achieving speeds much higher than what they were originally designed for.

Any backplane, in general, has a frequency dependent lossy characteristic which is due to the lossy nature of both the conductor and the dielectric. The channel loss increases as the frequency increases meaning that when a signal is sent down a channel, the high frequency components in it are attenuated more severely than low frequency components. As a result, if an impulse is sent down this channel, it gets spread over time. For a communication channel, this means that the data stream being transmitted gets highly distorted by the time it is received by the receiver. This distortion needs to be minimized if the signals are to be received without errors. This loss is a result of the combined effect of conductor loss which is mainly due to skin effect and dielectric loss which can be explained by the loss tangent of the dielectric.

2.2.1 Skin Effect

This is a physical phenomenon that causes the current flowing through a conductor to migrate towards the periphery or the 'skin' of the conductor as the frequency increases. This causes the effective resistance of the conductor to increase at higher frequencies. This phenomenon dominates DC losses that are caused due to the constant resistance of the wire, and it is the main reason behind the losses at frequencies in the multi-gigahertz frequency range. The amount of penetration of current into the metal is inversely proportional to the frequency as shown below in eq. (2.1).:

$$\sigma = \sqrt{\frac{2\rho}{\omega\mu}} = \sqrt{\frac{\rho}{\pi f\mu}} \qquad \text{eq. (2.1).}$$

where ω and μ are the angular frequency and the permeability of free space respectively. ρ is the resistivity of the metal and f is the frequency.

As a result, the penetration decreases inversely as the square root of frequency and the current experiences more resistance at higher frequencies causing higher losses.

2.2.2 Dielectric Losses

The dielectric constant of a material is represented as a complex quantity as shown in eq. (2.2):

$$\varepsilon = \varepsilon' - j\varepsilon'' \qquad \text{eq. (2.2).}$$

where the real part is the standard value of the dielectric constant and the imaginary part signifies the loss. Furthermore, the loss tangent of a dielectric is defined as:

$$\tan |(\delta)| = \varepsilon' / \varepsilon \qquad \text{eq. (2.3).}$$

Also, the power in an electromagnetic wave decays as:

$$P = Poe^{-\delta kz} \qquad \text{eq. (2.4).}$$

where, z is the distance traveled, Po is the incident power and $k = 2 \Pi f$. Hence, the power decreases with increasing frequency and distance. As a result of the frequency dependence, the response of the channel looks very similar to a low pass filter.

Hence the first step towards designing the bus is to characterize the channels and their losses by measuring their frequency response using a vector network analyzer(VNA). The measurement needs to be accurate as most of the design and simulations are based on this. The next subsection describes this process in more detail.

2.2.3 Channel Characterization

We first investigated the channel physical parameters including the transmission line widths and spacing between lines. A PCB was fabricated with three sets of 50 ohm transmission lines, with 5 channels in each set. Each set had a different channel width and channel spacing. The first set (set0) had 12.5 mil wide lines with 12.5 mil spacing. The second set (set1) had 13 mil wide lines with 17 mil spacing. The third (set2) set had ground lines in between the lines to shield the lines from each other and minimize crosstalk. The lines were each 12 mils wide and had a spacing of 13 mils to the ground shields. Simulations showed that the effectiveness of the shields was reduced if the distance between vias to ground on the shield lines were spaced more than 300 mils apart. Separate launch and receive boards with 5 SMA connectors were used to connect to the lines on both ends of the board. The fabricated boards are shown in figure 1.

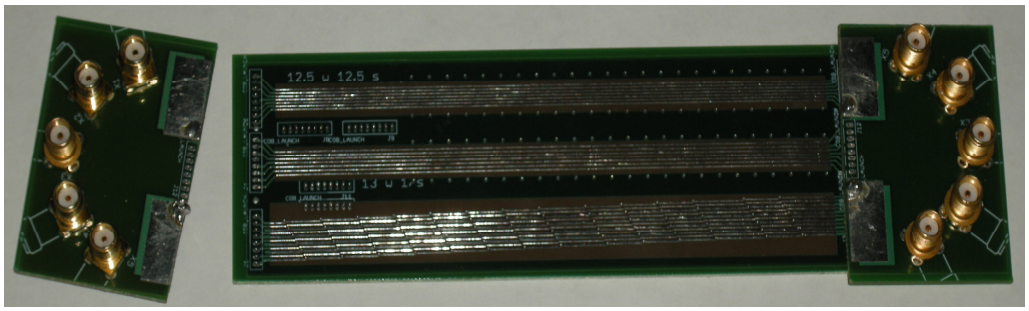

Figure 1 The fabricated PCB with three sets of lines: set 0 on the top, set 1 in the middle and set 2 at the bottom of the PCB. Launch boards shown on the sides.

There are mainly two methods to measure the characteristics of a channel: Time Domain Reflectometry (TDR) and S-Parameter Measurements. The primary difference between the two is that while TDR relies on measurements in the time domain, the S-Parameter (scattering parameters) measurements are performed in the frequency domain. TDR measurements are much easier to perform than S-Parameter measurements since they do not require extensive calibration

procedures. On the other hand, S-Parameter measurements provide more accuracy and modeling crosstalk is also easier with this method. Hence, S-Parameters were chosen to characterize the channels.

2.2.3.1 Defining S-Parameters

S-Parameters describe the response at each port of an N-port network for voltage inputs at each of the other ports. For example, if a 2-port network is considered, a two by two matrix containing 4 elements can be used to describe the system response characteristics: S11, S12, S21 and S22. The first subscript when naming an element is the responding port and the second one is the incident port. Hence, S12 describes the response at port 1 due to an input at port 2.

These S-parameters can be further categorized as transmission (S12 and S21) and reflection (S11 and S22) coefficients. It should also be noted that these parameters are actually vectors, meaning they are represented as a magnitude and a phase.

This concept can be applied to networks with more than 2-ports by simply adding more elements to the S-parameter matrix as shown below. A 10-port model was used in this work for simulation of 5 transmission channels.

$$(S_{11}) \quad (one-port)$$

$$\begin{pmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{pmatrix} \quad (two-port)$$

$$\begin{pmatrix} S_{11} & S_{12} & S_{13} \\ S_{21} & S_{22} & S_{23} \\ S_{31} & S_{32} & S_{33} \end{pmatrix} \quad (three-port)$$

$$Etc.$$

2.2.3.2 Vector Network Analyzers

Vector network analyzers (VNA) are the primary instruments used for S-Parameter measurements and a HP-8510 VNA was used in our measurements. It can measure full 2-port scattering matrices which contains the transmission and reflection scattering parameters for both the ports. This type of network analyzer consists of a sweep oscillator (almost always a synthesizer so that measurements will be repeatable), a test set which includes two ports, a control panel, an information display, and two RF cables to hook up the Design Under Test (DUT).

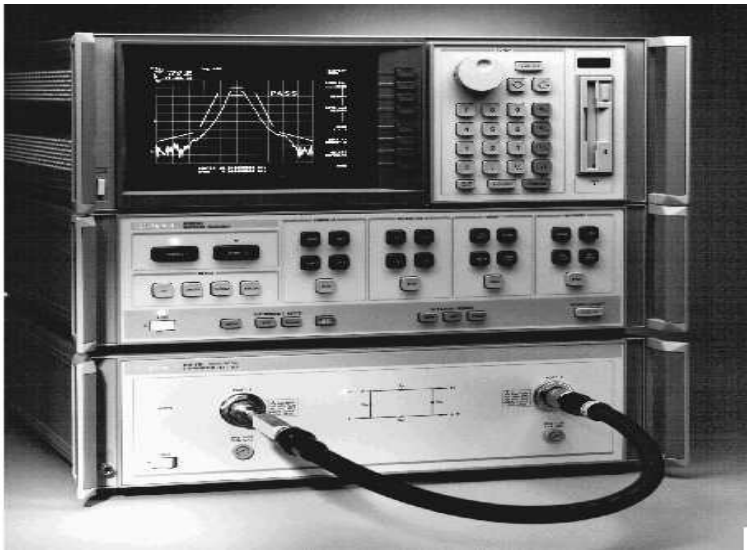

Figure 2 Vector Network Analyzer HP 8510

However, before making these measurements, the VNA needs to be calibrated for both reflection and transmission coefficients. Various types of methods are available for a full 2-port calibration and the most commonly used method is the Short-Open-Load-Thru Calibration. For this methodology, a standard calibration kit is available that has standard short, open, load and thru connectors. The VNA has an automated calibration procedure that prompts the user to

connect these standards to the test ports and once the calibration is completed, it applies the calibration results to all future measurements in real time.

The measured transmission characteristics S21 of the three sets of traces are shown in figure 3 below. Set 2 has the lowest loss out of the three which is due to the shielding traces, while set 1 and set 0 have similar characteristics.

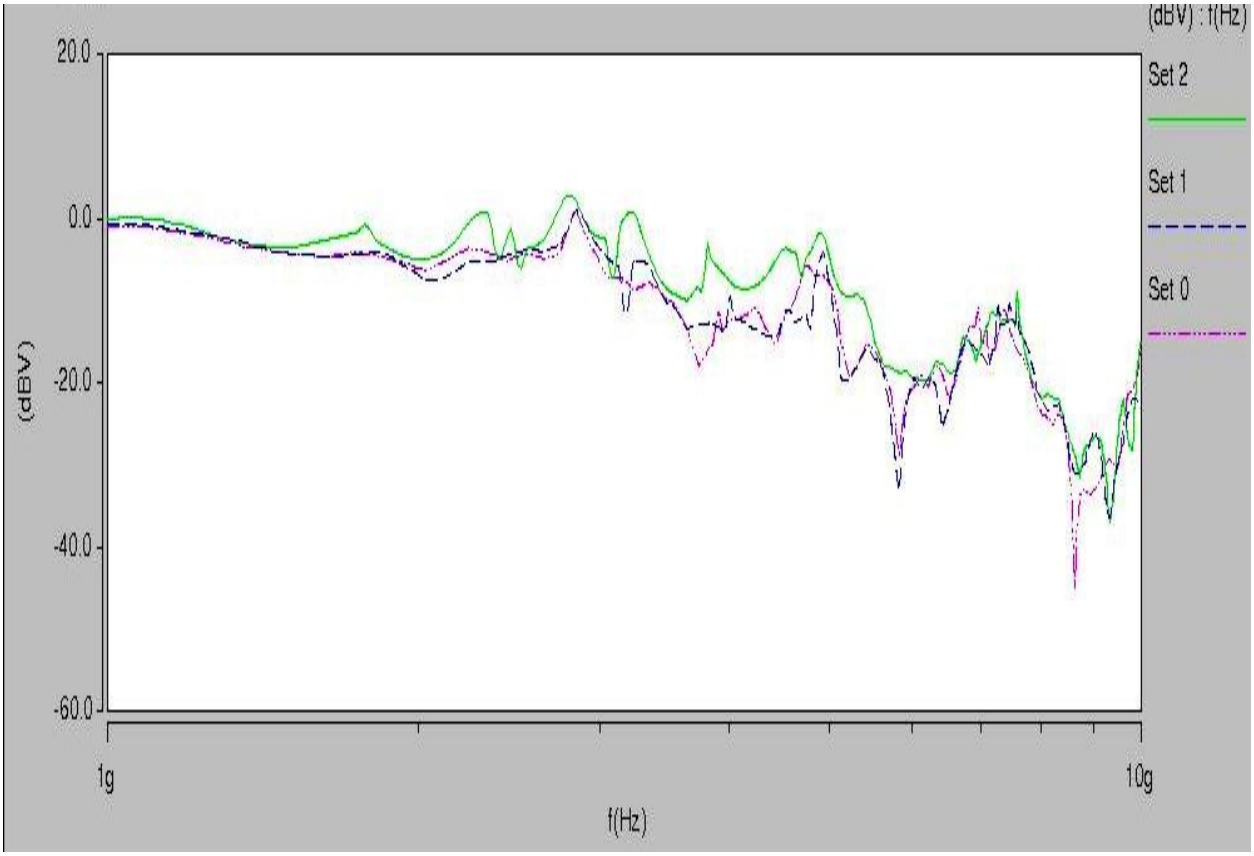

Figure 3 Channel transmission characteristics

2.3 Crosstalk

Crosstalk (XT) may be defined as a phenomenon by which the signal transmitted on one channel of a transmission system creates an undesired effect on another channel through the means of its electric or magnetic fields. The channel carrying data in this case is called the aggressor while the channel affected due to coupling is called the victim. Crosstalk in PCB's

occurs due to the coupling of conductors through the FR-4 substrate and this type of coupling is mainly capacitive by nature.

Crosstalk can be mainly categorized as Far-End Crosstalk (FEXT) and Near-End Crosstalk (NEXT). FEXT is the interference between two conductors measured on the same end of the conductors (transmission ends for both conductors) while NEXT is measured on opposite ends of the conductors (transmission end of one conductor while receiver end of the other). FEXT occurs mainly due to coupling of the traces on the transmission board. However, since the magnitude of FEXT is negligible, active cancellation techniques are used only for canceling NEXT in this work.

The coupling between conductors depends on the geometrical factors and frequency of operation. The most important factors that affect the amount of crosstalk are:

1. The spacing between conductors : crosstalk increases with less spacing
2. Permittivity of substrate material : crosstalk increases with increase in permittivity
3. Width of conductors : crosstalk decreases with wider conductors
4. Frequency : crosstalk increases with increasing frequency

The permittivity of the substrate is fixed in this case. The width of the conductor can not be varied significantly because the transmission line is supposed to have 50 ohms of impedance. The frequency of operation is also fixed to be at least 8 Gbps. Hence the only parameter that can be controlled is spacing. However, increasing the spacing increases the overall size of the bus on the PCB and is not practical beyond a certain point. As a result, we need to look at other methods to reduce crosstalk.

To determine the best possible architecture out of the three fabricated line sets, the crosstalk characteristics were measured for each of them. The same techniques and the VNA were used to measure the crosstalk characteristics as the ones used for characterizing lossy characteristics of the line sets. Instead of measuring the S-parameters by connecting the probes to the opposite ends of the same line, one probe is connected to the aggressor line and the other to the opposite end of the victim line. Hence, this basically measures the transfer characteristics of the aggressor line to the victim line. The plots for the NN crosstalk for the three line sets are shown in figure 4.

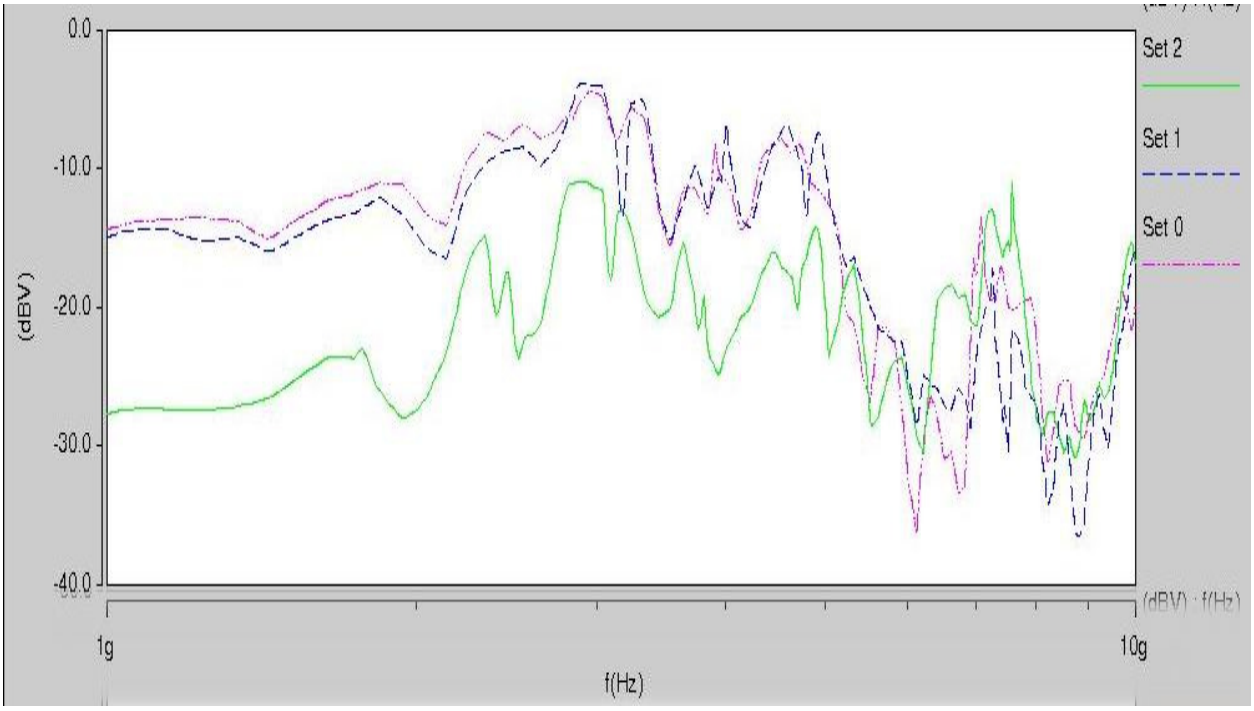

Figure 4 Crosstalk Characteristics for Nearest Neighbor (NN)

Based on the NN plots, it was found that the crosstalk in line sets 0 and 1 was as high as (-5dB). Line set 2 provided quite an improvement over the sets 0 and 1.

Also, the next nearest neighbor (NNN) crosstalk was measured and was found to be high on line sets 0 and 1. It was significantly lower on set 2 and was even low enough to be ignored. Hence, choosing line set 2 would provide a huge advantage as it would eliminate any extra circuitry needed to cancel NNN crosstalk. This saves a significant amount of power and also IC area. As a result of these observations, line set 2 was chosen to be used for this work. The NNN crosstalk characteristics for the three line sets are shown in figure 5 below:

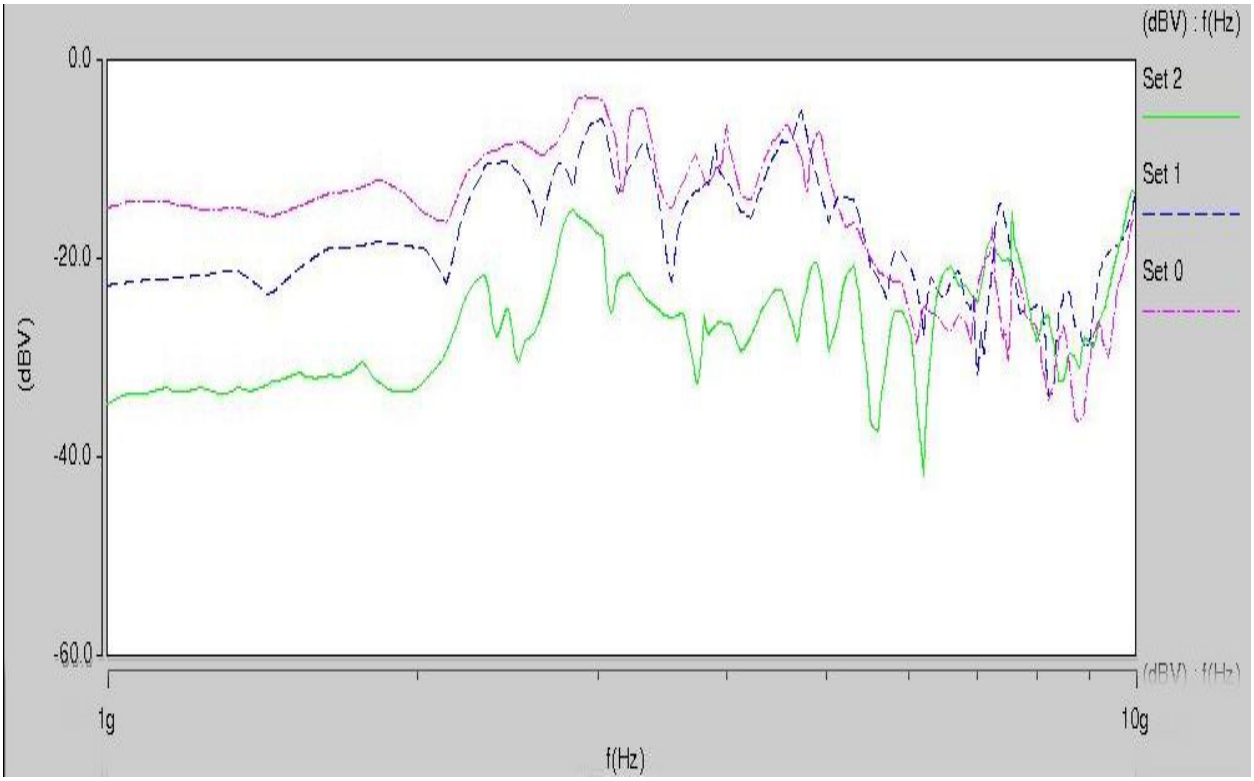

Figure 5 Crosstalk Characteristics for Next Nearest Neighbor (NNN)

Insertion of guard traces in between the channels is the first method that was employed in this work. This isolates the two channels, effectively decreasing the capacitive and inductive coupling coefficients. Also, it is important to ground these traces because if they are not grounded, they act as transmission lines themselves, carrying the coupled energy to the near and far ends. This energy could then couple to the channel that it is supposed to guard, thereby reducing its effectiveness. The line set 2 uses this approach to reduce crosstalk and is used for

this work. The main drawback for this being that the board area required almost doubles and reduces the amount of savings compared to a differential bus. A differential bus usually has more spacing between transmission line pairs so a single-ended bus with grounded shield traces may still reduce PCB area by about 30%.

However, even after using guard traces, the crosstalk is quite high which can be seen from the channel characteristics. Hence, a second method is employed which involves using an active noise cancellation techniques, like the ones described in [4] and [7]. The first technique, as described in [4], is based on the adjustment of zero-crossing points at the receiver to improve crosstalk induced jitter but does not attempt to reduce the crosstalk signal energy. The second technique described in [7] aims at reducing the crosstalk energy by using a discrete time FIR filter. The idea here is that if we can model FEXT as discrete-time sampled responses, we can then implement an FIR filter that can cancel it out. Also, this filter can be implemented as an extension to the pre-emphasis filter which is described later.

2.4 Power Supply Bounce

Another phenomenon that has an impact on signal integrity is power supply bounce. Variations in the power supply can cause the performance of all the circuits on the chip to degrade. It is also a well known fact that designs with good power supplies provide a better environment for signal transmission. Power supply variations can occur in all systems but the variations can be even larger for a single-ended design. This is because a differential bus is always transmitting the signals along with their complements and hence, is always balanced. A single-ended system on the other hand, could be transmitting more high signals than low or vice versa. In a worst case scenario, it could be transmitting all high signals and no low signals or vice

versa. This could result in large variations in the total current being drawn from the supplies and also variations in the voltage.

A way to get around this problem is to encode the 16 channels of the bus on to more lines. Encoding of a bus in such a way is quite common and there are several possible choices for codes. By encoding a fewer number of channels on to a larger bus, the number of code words available for use increases which provides the flexibility of choosing desired codewords and avoiding the ones that could cause problems. For example, if a 4 bit bus is encoded on to a 5 bit bus, the number of available code words is 32 while only 16 code words are required to transfer 4 bits of data. Hence, it is possible to choose 16 codes out of the total 32, leaving out the ones that are more unbalanced than the others.

This approach will be applied to balance the 16 bit wide bus in this work and the details about the codes used are provided in the next chapter.

2.5 Common-mode Rejection

Common-mode voltage is defined as a common voltage that appears on both lines of a differential signal with the same magnitude and phase. Since a differential signaling system makes decisions of a signal being high or low based on the difference of two signals, any common voltage that occurs on both lines gets canceled out. As a result, if common-mode noise occurs on a bus, it just gets canceled out and does not affect signal integrity very much.

However, in single-ended signaling, if common-mode noise occurs and the common voltage varies during transmission, it might result in errors. For example, if a single-ended signal is transmitting low but common-mode noise causes it to rise beyond the threshold that

distinguishes highs from lows, it will be received as a high signal. Even if the variation is not big enough to cause errors, it will still reduce noise margins of the system, which is undesirable.

Since differential signals have inherent common-mode noise rejection, a simple and reliable solution in single-ended signaling is using a differential channel along with other single-ended channels. This differential channel can be used to extract the common-mode voltage at the receiver end of the bus and distribute it to all the other single-ended channels in the bus. Also, since a 16-bit wide bus would be spread over a relatively large PCB area, it might be possible that channels that are far apart experience different variations in common-mode voltage. Hence, it was decided to have two differential channels placed at a distance from each other, which provide common-mode voltage for channels closest to them. This adds two extra channels to the bus as an overhead but this was an acceptable trade-off for achieving 'pseudo common-mode rejection'.

2.6 Channel Loss Compensation

2.6.1 Inter-Symbol Interference(ISI)

ISI occurs when the energy from a symbol spreads to other adjacent symbols and appears as noise in the other symbols. When a data stream containing a series of pulses is sent down a band-limited channel, the pulses spread out in time hence interfering with each other. This spreading of pulses can be explained by means of Fourier analysis as described below.

From Fourier analysis, it is a well known fact that any periodic signal can be obtained by adding a series of sinusoidal signals with proper amplitudes and phases. This series could include up to an infinite number of terms depending on the characteristics of the waveform to be constructed. The Fourier Series is given by eq. (2.5). :

$$f(t) = \frac{A_o}{2} + \sum_{n=1}^{\infty} (a_n \cos n\omega_o t + b_n \sin n\omega_o t) \qquad \text{eq. (2.5).}$$

where the A_0 term is the DC component while a_n and b_n are the Fourier coefficients for each harmonic. The fundamental frequency in radians is given by ω_0.

Hence, it can be shown that a square wave can be formed by combining a number of odd numbered harmonic sinusoids as shown in figure 6

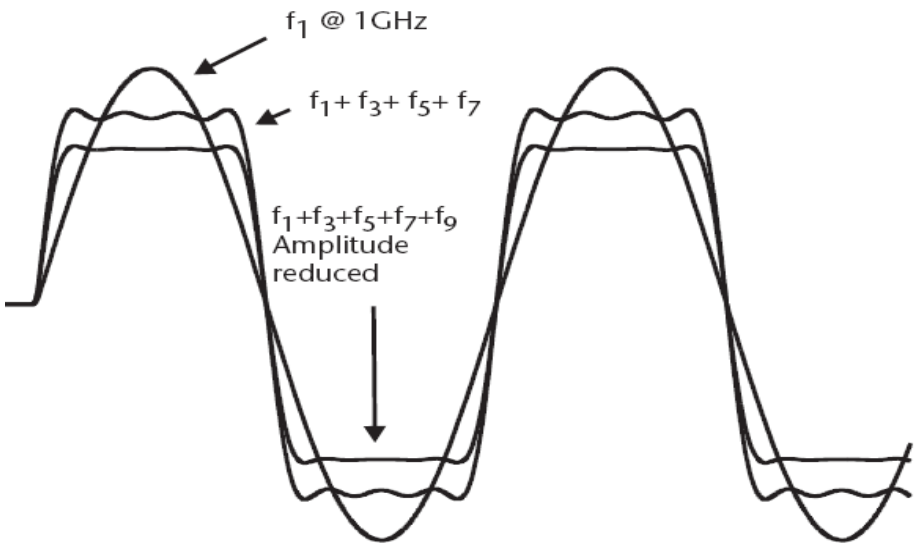

Figure 6 Effect of removing higher frequency harmonics on a square wave

The waveform gets squarer as more and more harmonics are added to it. Or conversely, the waveform gets rounded out and spreads out compared to the square wave when higher harmonics are removed. Although a lossy transmission channel does not completely remove high frequency harmonics but it does attenuate them. And although a data stream is not necessarily a square wave, it contains pulses of various widths and these get distorted in different ways due to attenuation and phase shifts of their harmonics.

This distortion is most commonly compensated for by using equalization techniques as described in the next section

2.6.2 Equalization and Pre-Emphasis Techniques

Equalization techniques are used to compensate for the lossy response of a transmission channel by providing a response which is approximately the inverse of the transmission response of the channel. As a result, when the signal is received at the receiver, it has a uniform or an 'equalized' response.

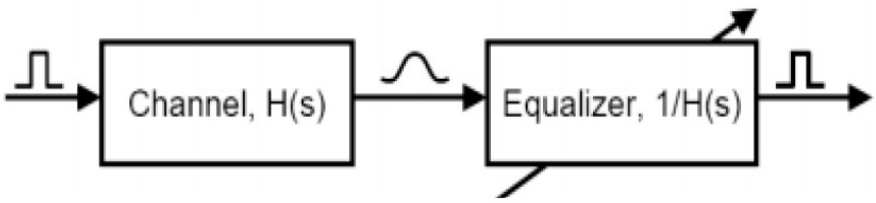

Figure 7 Basic principle behind equalization schemes

Such techniques can be applied at either the transmitter or the receiver end, however most often it is done at both ends. However, when equalization is applied at the transmitter, it is commonly referred to as pre-emphasis.

2.6.2.1 Pre-emphasis

Pre-emphasis is a common feature in multi-Gbps serial transmission designs. In these designs, the current is increased over the nominal current whenever a bit transitions. Similarly, the current is decreased when no transitions occur for several bit times. The result is that the higher frequency components of the signal are pre-emphasized as compared to the low frequency components.

2.6.2.1.1 Finite Impulse Response (FIR) Filters

The most common method of achieving pre-emphasis is by using Finite Impulse Response (FIR) filters with adjustable tap coefficients. This technique is also used for this work. Figure 8 below shows a general block diagram for a 3-tap FIR filter with adjustable coefficients, C_{-1}, C_0, and C_1. The blocks denoted by Δ are unit delay elements. This delay can be as long as the symbol period but it is usually smaller, in which case the equalization is called fractionally spaced equalization. These delay blocks are often realized using flip-flops or latches.

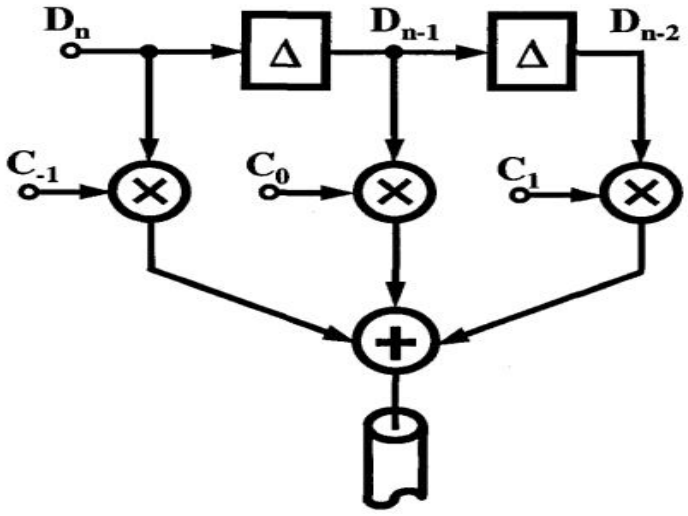

Figure 8 Block Diagram for FIR filter equalization scheme

The tap coefficients can be adjusted according to the characteristics of the transmission channel and hence this equalization approach can be applied to channels with different characteristics. In a time-invariant system, where the channel characteristics do not vary with time, the taps have to be adjusted just once and there would be no need to monitor the characteristics continuously. However, the channel characteristics could change by their usage

over long periods of time and some adaptive techniques have been proposed to continuously adjust the tap coefficients periodically. A very commonly used algorithm to achieve this is the LMS algorithm as described in [6] or its variations as described in [8]. However, these adaptive algorithms are quite complex and increase hardware complexity significantly. For this reason, adaptive equalization is not considered in this work and the channel characteristics are assumed to be constant.

2.6.2.1.2 Continuous Time Equalization

An alternative to discrete time FIR filters is the use of continuous time equalization which uses analog circuitry with tunable poles and zeros to obtain characteristics which are the inverse of the channel characteristics.

The equalization scheme presented in [9] uses this method and is based on a continuous time filter with a second order filter transfer function. The transfer function is given by equation eq. (2.6). and it has one zero and two poles.

$$F(s) = \frac{s + z}{s^2 + as + b} \qquad \text{eq. (2.6).}$$

The transfer characteristics of the filter can be tuned by changing the coefficient 'a' in the denominator. Also, the zero location given by 'z' can be adjusted. This idea can be further extended to include several high pass filters with different transfer functions and summing their outputs to get a more complex transfer function.

The advantages of this approach over the FIR filter approach is that it consumes less power and contributes less noise and jitter to the transmission system. However, it is much more

complex to realize a continuous time filter in practical applications because it is difficult to tune such circuits. For this reason, they also require adaptive tuning circuits based on algorithms similar to those in adaptive FIR filter based schemes which add overhead to the system. Considering the trade offs between the two approaches, the commonly used FIR filter based technique was chosen for this work.

2.6.2.2 Passive Equalization

Passive equalization is achieved by using filters to compensate for high frequency losses. A common method to do this is by placing passive high-pass filters on the transmission lines which would attenuate the low frequency components of a signal, making their amplitudes close to those of high frequency components. This is similar to de-emphasis since low frequency components are attenuated to achieve equalization.

This method can be used at both the transmitter and the receiver side and is sometimes used along with other equalization techniques. When more filtering action is required, second or third order high-pass filters can be used.

2.6.2.3 Receiver Equalization

Although transmitter-side equalization is implemented using linear FIR filters in almost all serial transmission designs, there are quite a few choices available when to comes to choosing the receiver-side equalization method. Some common approaches are described below.

2.6.2.3.1 Decision Feedback Equalization (DFE)

Decision feedback equalizers are non-linear equalizers which equalize a received symbol based on the decisions for the previous symbols. In simple terms, the ISI caused by previous

symbols is subtracted out from the symbol that is currently being received. The block diagram is shown below in figure 9.

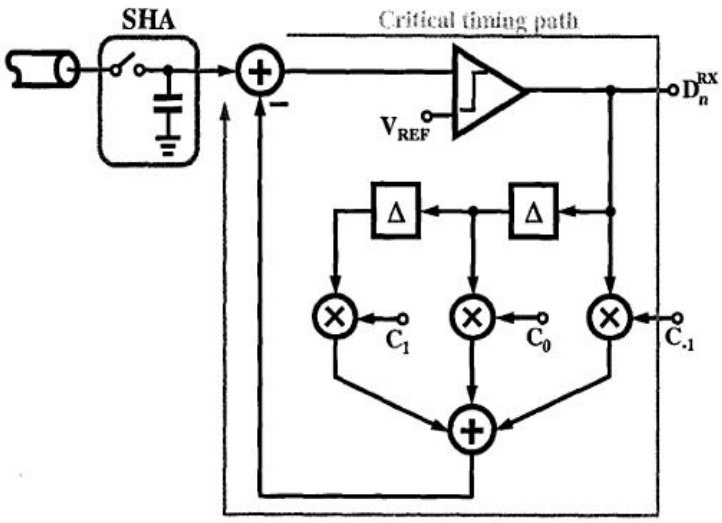

Figure 9 Block Diagram for a basic DFE equalizer

A sample and hold circuit is used to sample the incoming data bits. A comparator is used to decide if the current bit is a high or a low. This signal is then delayed and subtracted from the bits that follow after being scaled appropriately by the multiplication coefficients C_{-1}, C_0, and C_1. Also it is possible to use adaptive techniques to adjust the tap coefficients periodically similar to FIR filters.

Non-linear equalization schemes like DFE do not amplify crosstalk noise while linear equalizers do because they amplify all high frequencies and crosstalk energy lies in the higher frequency bands. However, they consume more area and power due to the presence of comparators and sample and hold circuits. Also, another problem with the DFE is that if a wrong decision is made in the first place, the ISI could get added to the bits that follow which could lead to more wrong decisions. The error would then keep propagating for several bits. This

approach was not studied further in this work as much simpler equalization techniques are available that provide power and area savings.

2.6.2.3.2 Continuous Time Linear Equalization (CTLE)

These equalization schemes are very similar to those used for continuous time equalization for the transmitter side. They are based on the use of passive elements to obtain high-pass filter

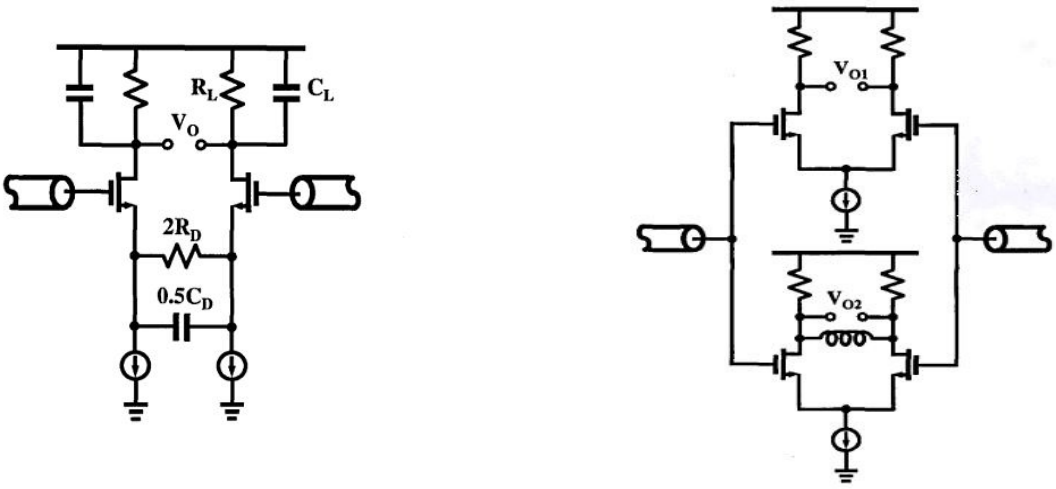

Figure 10 Continuous Time Equalizers used for the receiver

characteristics. However, there are some instances where active CTLE techniques are used which provide gain in addition to equalization.

Some CTLE circuits are shown in figure 10 and figure 11 as examples. Their working is similar to the continuous time equalizer described earlier and they basically have a filter response which allows for appropriate placement of poles and zeros through the adjustment of the passive components. When the zeros are placed before the dominant poles, gain peaking occurs for the frequency range in between the poles and the zeros. The circuit on the left of figure 10 shows an

implementation that uses a parallel RC circuit to add a zero to the transfer function. The circuit on the right of figure 10 achieves the same purpose but through load inductors. Figure 11 shows an active implementation which can provide gain by using an operational amplifier.

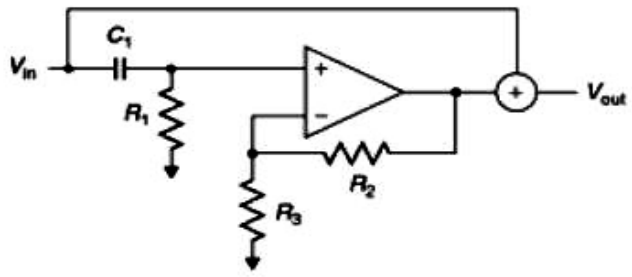

Figure 11 Example of an active CTLE

These equalizers do not provide as much gain as the DFE and are not as easy to tune for the required channel. However, they are very simple to implement and require much lower power.

2.7 Signaling Schemes

So far it has been considered that signals use two levels to convey whether a high (bit 1) or a low (bit 0) is being transmitted. This type of signaling scheme is called non-return-to-zero (NRZ). However, there are other signaling schemes which use more than two levels for communication. The commonly used multi-level schemes are duobinary and PAM-4 signaling schemes which have three and four levels respectively.

The advantage of having more levels is that the fundamental frequency of the data stream is decreased while keeping the same data rate. Since the fundamental frequency is decreased, the insertion loss for multi-level signals is lower than NRZ. However, this comes at the cost of reduced eye heights. The same voltage range gets divided into several levels and this makes the

detection more difficult at the receiver, which is essentially the reduction in noise margins. Also, more complex and accurate receivers are required to handle these small signal levels.

At low data rates and with no equalization techniques, multi-level signaling sometimes proves beneficial. However at high data rates and when equalization techniques are used, the trade off between insertion loss and signal levels is not practical since equalization already compensates for the insertion losses. Hence, multi-level signaling is not as efficient at these speeds and NRZ signaling is almost invariably used in high speed serial communication.

Chapter 3

3. System Study

The goal of this thesis was to design a high speed single-ended data bus for chip-to-chip communication with data rates close to that for conventional differential busses and investigate techniques that could make single-ended signaling possible. The final simulation results show that the design should achieve data rates up to 8 Gbps and beyond depending on the channel characteristic and is comparable to what differential buses can achieve for a similar length channel on the same PCB material. The bus is designed to be 16 bits wide which is a practical size for data busses and the bus provides a total data throughput of 128 Gbps. The bus is designed for standard FR-4 PC boards which are most commonly used for board designs. The bus is designed to accommodate different lengths for each channel ranging from 5 inches long to 7 or 8 inches long. It employs equalization schemes at both the transmitter and receiver ends which are tunable independently for each channel.

The problems commonly associated with single-ended signaling, namely significant crosstalk, the absence of common-mode rejection like differential busses and power supply bounce, were addressed to within permissible limits. A crosstalk reduction scheme is implemented which uses FIR filters to reduce crosstalk noise. Two differential channels are placed in between the single-ended ones to provide pseudo common-mode noise rejection. And finally, the 16 bits are encoded to 19 lines to eliminate power supply bounce problems that arise due to imbalance in the single-ended bus. This brings the total number of channels to 21 although only 16 bits are transmitted on them. However, this is still a 35% improvement over a 16 bit differential bus which would require 32 channels. This translates to savings in I/O pins but

not much in board area as guard traces were used in the final design to isolate the channels adjacent ones. Although, this could still be less than the area required by a differential bus since the channels were placed close together, an exact measure of board area requirements is not known yet. A differential bus will also be included on the same chip in the final design for comparison and hence more information might be available after characterization of the busses.

Also, since the bus is designed to be source synchronous, a phase adjust circuit was included at the receiver to avoid the need for a phase lock loop (PLL) or similar techniques which are very complex to design and are not really a concern for this work. A simple delay line is used to adjust the phase of the clock to center the clock edges at the correct positions.

Considering the overhead lines, the final bus structure looks like the one shown in figure 12.

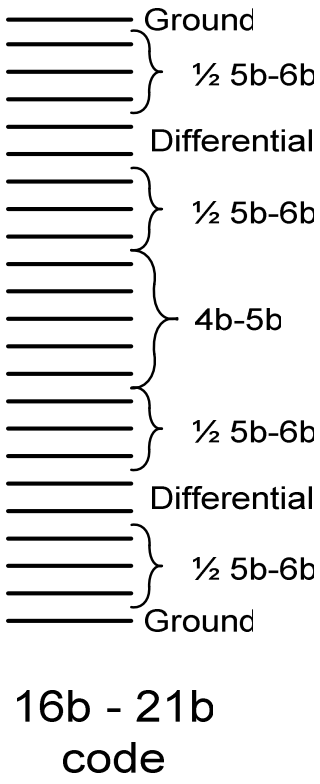

Figure 12 Single-Ended Bus Structure

Two ground traces for shielding are placed on both ends of the bus to shield the entire bus from any other signals on the same PC boards. The system level transmitter and receiver designs are described next in this chapter.

3.1 Transmitter Design

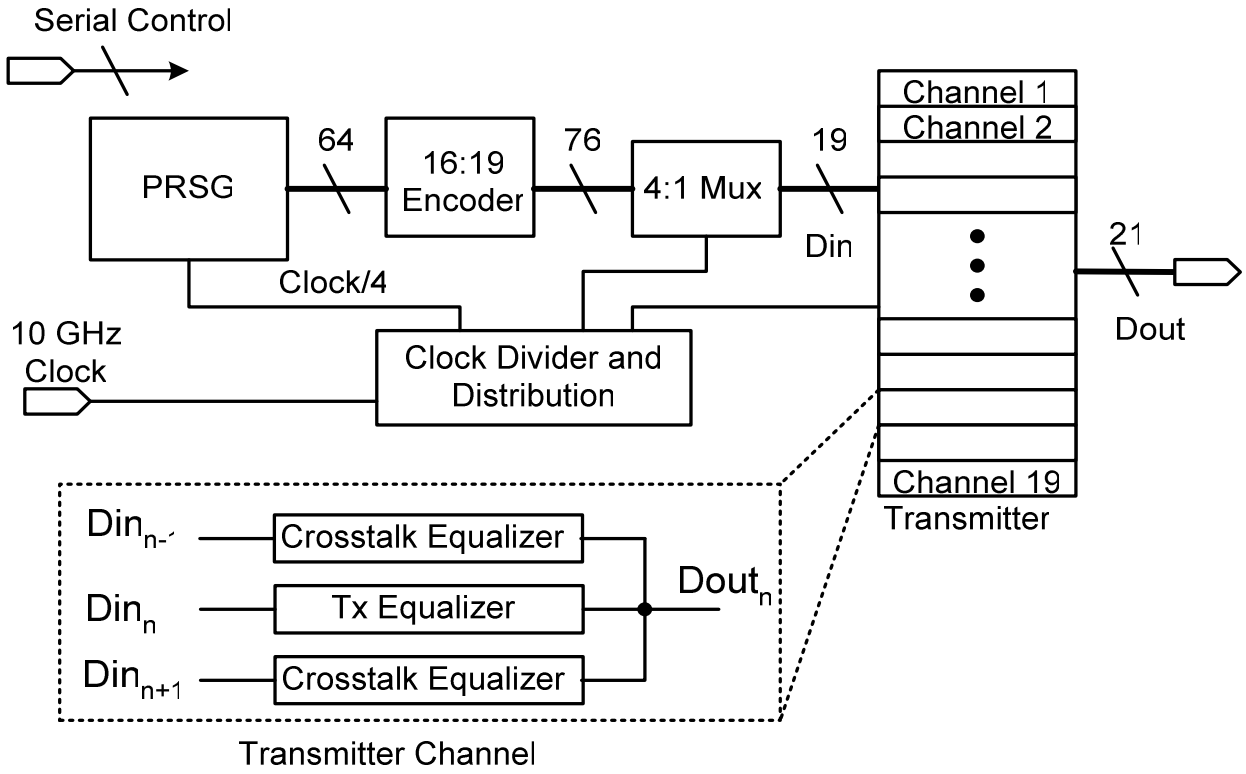

Figure 13 Transmitter Block Diagram

A pseudo random sequence generator (PRSG) is used to develop random data sequences for testing the transmission channel. The PRSG generates 64 bits of random data every clock cycle and operates at only one-fourth of the system clock rate which can be as high as 10 GHz because the operation of the bus is to be tested for 10 Gbps as well. The PSRG and the multiplexing is

not a part of this thesis but is included here completeness since it will be required later to test the system. This data from the PRSG is then encoded using a 16 bit to 19 bit coding scheme to remove power supply bounce. This slower data is then multiplexed onto the 19 channels of the bus to obtain a data rate as high as 10 Gbps. Each of the channels is made up of three basic blocks: two crosstalk equalizer blocks to cancel out crosstalk from the two adjacent channels and the main transmission equalizer filter block. The output of these blocks is summed and then transmitted over the physical channel on the PCB.

3.1.1 Coding Scheme

The coding scheme for power supply bounce reduction is focused on having the single-ended bus balanced at all times, which means that the bus should have an equal number of high and low channels at all times. As described earlier, this is done by encoding the channels on to more output channels which increases the number of available code words. Hence, it becomes possible to choose the code words according to our requirements and in this case, the code words were chosen so as to minimize imbalance of highs and lows. Also, instead of encoding the whole 16 bit bus, the bus was divided into three sections and each section was encoded separately. Two channels are differential and hence are inherently balanced, leaving 14 single-ended channels that need to be balanced. The 14 channels are divided into 3 groups: two groups of 5 channels each and one group of 4 channels. The groups containing 5 channels are encoded on to 6 channels and this is called 5b-6b encoding. This means that 32 code words are required for each group for communication while 64 code words are possible. The other group containing 4 channels is encoded using a 4b-5b encoding scheme which encodes the channels on to 5 output channels. In this case, 16 code words are required, while 32 are available. Hence the code words with least imbalance were chosen and they are shown in table 1.

5b6b Codes		
Balanced	Difference of -2	Difference of +2
000111	000101	111010
001011	000110	111001
001101	001001	110110
001110	001010	110101
010011	001100	110011
010101	010001	101110
010110	010010	101101
011001	010100	101011
011010	011000	100111
011100	100001	011110
100011	100010	011101
100101	100100	011011
100110	101000	010111
101001		
101010		
101100		
110001		
110010		
110100		

4b5b Codes	
Difference of -1	Difference of +1
00101	01011
00110	01101
01001	01110
01010	10011
10001	10101
01100	10110
10010	11001
10100	11010

Table 2 Code words selected for the two coding schemes

However, although the available code words increase by a factor of 2 by these coding schemes, they are still not enough to have all the code words balanced. Hence, even after encoding we are left with some imbalance. The maximum imbalance as can be seen from the table, is +/-2 for 5b-6b encoded channels and +/-1 for 4b-5b encoded channels. The codes words for 5b-6b encoding can be perfectly balanced at some times and have an imbalance of +/-2 at other times. If both 5b-6b groups are unbalanced at +/-2 then one can be chosen as the opposite of the other so that overall the balance is determined by the 4b-5b group, which has an imbalance of +/-1. If they are each balanced, then overall the balance is also determined by the 4b-5b group. If one is balanced and the other is unbalanced, the sign of the balance can be chosen to be the opposite of the imbalance in the 4b-5b group and the balance is the opposite of the 4b-5b group.

In all cases the imbalance is +/-1. To cancel this out, a dummy driver is included on the chip which is like all the other channel drivers but is not connected to an I/O pad. Thus the 16 bit bus is balanced overall. The implementation was again not a part of this thesis however some details have been provided in the next chapter that describes the design of system components.

3.2 Receiver Design

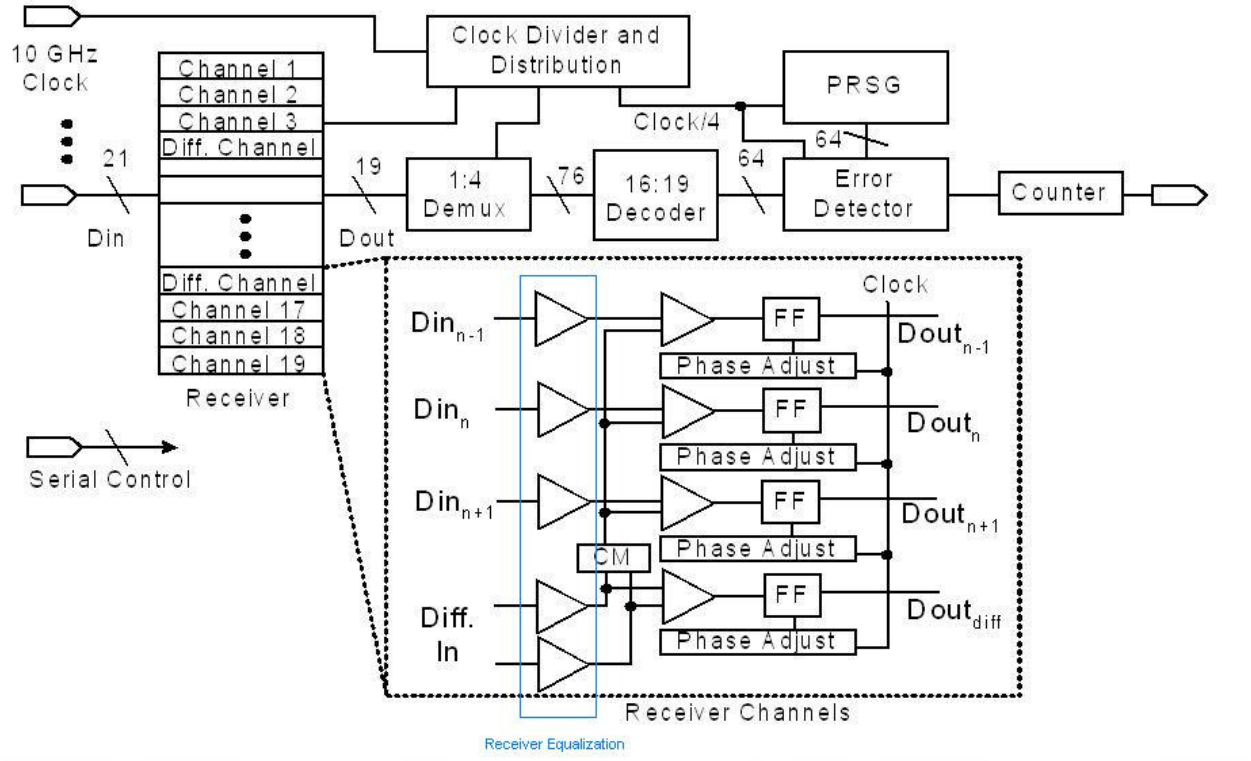

Figure 14 Receiver Block Diagram

The signals coming in to the receiver over the channels are received by an equalizer which is integrated into a common-gate amplifier circuit. This common-gate amplifier provides 50 ohms termination to the channel and also a method to adjust the common-mode voltage of each channel individually. However, once the common-mode has been set, a method is required to monitor the common-mode and adjust the common-mode if a change occurs. This monitoring

and adjustment is done by extracting the common-mode information from the two differential lines and distributing it to the single-ended channels. Assuming that a common-mode noise signal affects all the channels equally, then the common-mode voltage on the differential lines changes, then the common-mode voltage on the single-ended lines change as well.

The received data signals, after being equalized are amplified by using an n-MOSFET differential pair. This differential pair takes the received data on the single-ended channel and the common-mode signal from the differential channel as the two inputs and amplifies the received data signal. This differential pair also provides a differential output, essentially converting the single-ended data signals to differential data signals. In the case of the differential channel, it takes the two differential channels as the inputs, like any conventional differential receiver scheme.

The differential data signals from the differential amplifier are then clocked in to the flip flop which decides whether the signal bit is a high or a low and stores it for the current clock cycle. This is followed by a 4 to 1 demultiplexer which demultiplexes the 19 data signals on to 76 channels to again bring the data rate down to one fourth the original data rate. This is the data that was encoded at the transmitter end to remove power supply bounce problems. Hence, the data must be decoded to reproduce the original 16 bit data that the PRSG generated at the very beginning. The same PRSG, as the one in the transmitter, is also present at the receiver and it generates the same data that is used to compare the received data. Any differences between the two would mean that an error occurred and the error detector keeps track of the errors. A counter is used to count the number of errors from which the bit error rate can be determined.

Chapter 4

4. Design of Components

4.1 Transmitter Equalizer

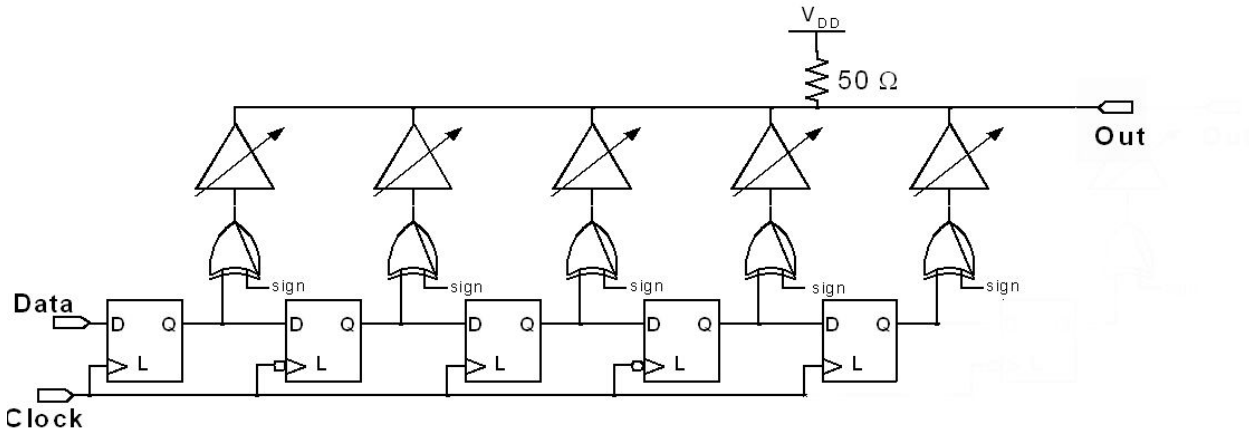

Figure 15 Transmitter Equalizer Filter

The circuit level implementation for the FIR filter equalization scheme that was described earlier is shown above in figure 15. The delay elements are implemented using D-latches here although other methods are also possible. Some designs use flip-flops instead of latches but the delays in that case are one full clock cycle while they are only half clock cycles if D-latches are used. Another method is to use analog circuits for implementing delays in which case even smaller delays would be possible, but an analog implementation for delay would use capacitors and inductors which would require more area and also the delays would be more susceptible to process variations. The coefficients of the filter can be set by using adjustable taps which are driven by separate control signals. The sign of the coefficients is set by using a multiplexer which basically inverts the signal polarity to change the sign of the coefficient that it controls.

All these components are designed in current mode logic (CML) because of speed advantages over most other logic styles.

4.1.1 D-Latches

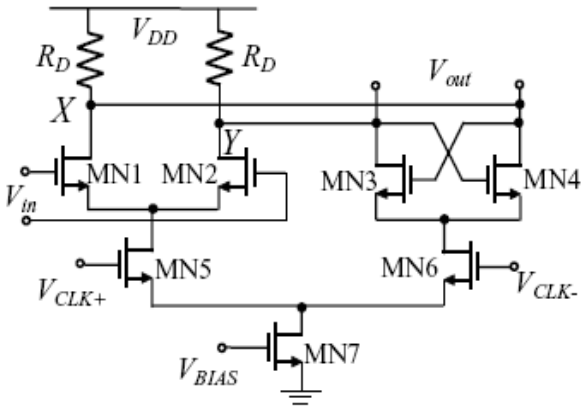

Figure 16 Conventional CML D-latch

CML is a differential logic family and CML gates are frequently used in high-speed designs. It contains a constant current source MN7, which draws input-independent current at all times.

The operation of the latch can be explained in two stages, track and hold. The transistors on the left of the schematic, including MN5, MN1 and MN2 make up the tracking circuit while those on the right, including MN6, MN3 and MN4 form the hold circuit. The tracking stage occurs when the clock goes high, turning the transistor MN5. The output resistors are then driven by the differential pair made up of MN1 and MN2. Depending on the input signals, one of them gets turned on while the other is turned off, providing a differential voltage output V_{out} across the resistors as shown. The hold stage occurs next, when the clock goes low, turning on transistor MN6. This then turns on the cross-coupled differential pair, made up of MN3 and MN4, which holds the value that was available at the output at the end of tracking stage. When the clock goes

high again afterwards, the next input is tracked by the track circuit and the previous data is overwritten.

The CML D-latch provides high-speed operation but at the cost of high power dissipation low output signal swing. The output swing depends on the resistor values times the current. Larger resistor sizes provide higher swing but also slow down the speed. Low resistor values require larger current to obtain reasonable output swings and increases power dissipation. Optimizing speed and power dissipation is a challenge. The D-latch designed for this thesis consumes around 1.8mW of power at 8 GHz. Another drawback with CML gates is that they consume relatively more area than CMOS.

4.1.2 CML Multiplexer

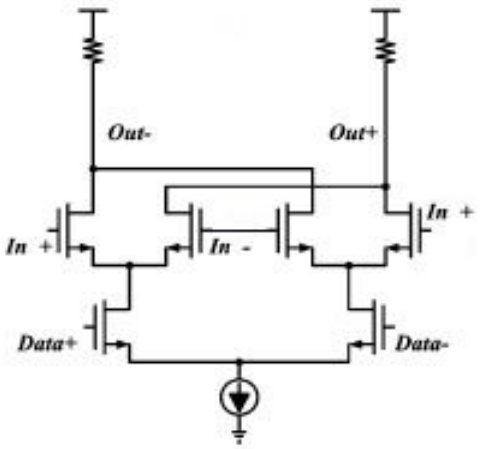

Figure 17 Conventional CML Multiplexer

The CML multiplexer shown in figure 17 is used to control the sign of the filter coefficients of the equalizer. Similar to the latch, a constant current source is present and two transistors control the steering path for this current. The output from the latch is connected to the positive and negative inputs as shown. The sign bit is shown as Data in the figure and hence, if the sign bit is high, the differential pair on the left side of the circuit is active and the differential inputs

appear reversed at the output. If the sign bit is low, the differential pair on the right is active and the input appears as the output of the multiplexer without any changes.

The CML multiplexer does not have to go as fast because the sign of the coefficients is set once at initialization. Once the coefficients and the signs are set, they are assumed to be constant for the channel since the channel characteristics should only vary over long periods of use. Hence, the current in the circuit can be reduced compared to the other CML gates. The multiplexer in this design consumes 0.77mW.

4.1.3 Variable Gain Tap

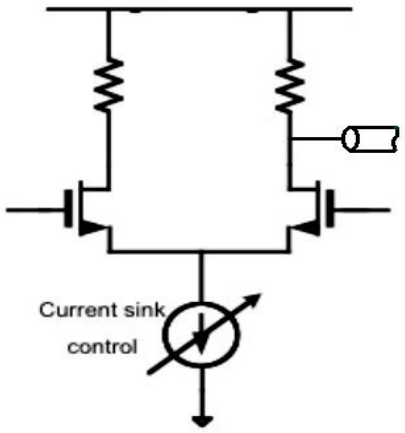

Figure 18 Variable Gain Tap

The function of the variable delay tap is to transmit signals on the transmission line where the amplitude can be controlled. This is done by using a simple differential pair driver which is driven by a tunable current source. The resistors are 50 ohms and they provide termination for the transmission line. Each tap does not have these terminating resistances. They are common for all the taps. The tap provides differential output but only one of the outputs is transmitted since

the channel is single-ended. For the two differential channels in the bus, both the outputs are connected to transmission channels.

The number of taps in the equalization filter was varied to determine an optimal number. Initially 4 taps were used but adding a fifth tap improved performance considerably. Adding a sixth tap did not improve the equalization much and would consume additional power and area. Hence, 5 variable gain taps were chosen for the equalizer and each one is sized according to the amount of gain that it would approximately require. For example, the first tap is for canceling the precursors and does not require a lot of gain since the precursors are small in magnitude, but the second tap is the main driving tap and it should have a large gain to produce the required amount of voltage drop across the 50 ohm resistor. Hence the second tap has larger transistors to support larger currents while the first tap is designed with smaller transistors. Hence, although the taps have adjustable currents, the range for each tap is different with the first and the last taps having a smaller range, the second tap having a large range and the third and fourth taps with slightly less gain than the second tap.

The variable current sources are controlled digitally using separate digital-to-analog converters (DACs) which are described later.

4.2 Crosstalk Cancellation

The crosstalk reduction filter is very similar to the equalization filter and also has the same number of taps. It is based on the design described in [7]. Similar to the equalizer, different number of taps were tested in simulation and using 5 taps again seemed to be the best configuration. The only difference with the equalizer is that the taps are not required to have large gains since they are only supposed to cancel out small crosstalk signals. Hence the

transistors in the taps are smaller than those used in the equalizer; however the topology of the taps remains the same in both cases.

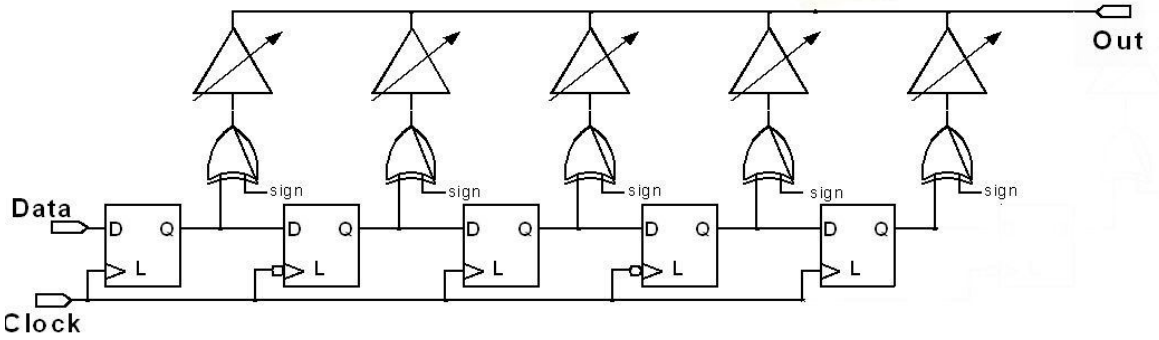

Figure 19 Crosstalk Reduction Filter

Also, the 50 ohm termination resistors are absent since they are already included in the transmitter equalizer and the data coming in to the crosstalk reduction filter is the data from the adjacent channel and not the main data channel. Each channel has two of these filters, one for each channel adjacent to it.

4.3 Receiver Equalization

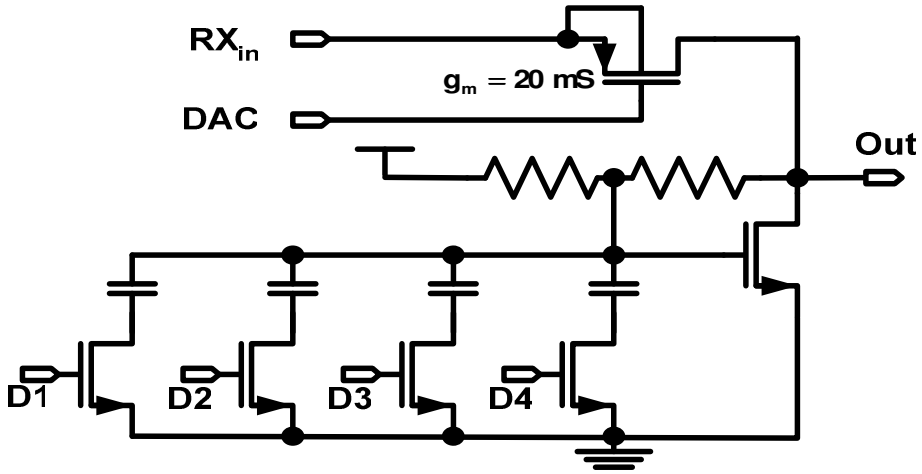

Figure 20 Receiver equalization

The receiver equalizer is basically a common-gate amplifier as shown in figure 20, with equalization added to the circuit. Common-gate amplifier circuits are sometimes used for RF receivers due to the ease of impedance matching and lower noise. In this configuration, the source terminal of a transistor serves as the input, the drain is the output and the gate is biased at a fixed but variable voltage. It is characterized by low input resistance and high output resistance.

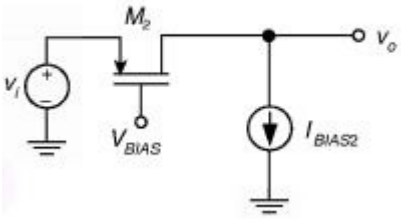

Figure 21 Basic Common-gate Amplifier

The input resistance is given by eq. (2.7).

$$R_{in} \cong \frac{1}{g_m + g_{mb}} \approx \frac{1}{g_m} \qquad \text{eq. (2.7).}$$

Hence, if g_m for the PMOS transistor is biased at 20 mS, the input resistance is 50 ohms. This eliminates the requirement for 50 ohm terminating resistors for the receiving end of the channels.

The biasing voltage at the gate of the PMOS transistor is used to control the common-mode output voltage and is controlled by a DAC. Changing this biasing voltage also changes the input resistance but it is not a large variation and hence can be tolerated. The resistors and the capacitors provide the peaking at higher frequencies which provides equalization. The magnitude of peaking depends on the capacitance value and hence, the capacitance value needs to be adjustable to have tunable equalization. This is implemented using four binary weighted

capacitors which can be added or removed from the circuit by using transistors as switches to connect them or disconnect them. For example, if D1 and D2 are high, the first two capacitors will be connected in the circuit while the other two will be disconnected and so on. The bits D1-D4 are set using digital registers.

4.4 Differential Amplifier

The output of the common-gate receiver is fed to a differential amplifier which amplifies the received single-ended signal and also converts it into a differential signal. The gain of the amplifier is dependent on g_m of the transistors and the load resistance, R_L as seen in eq. (2.8).

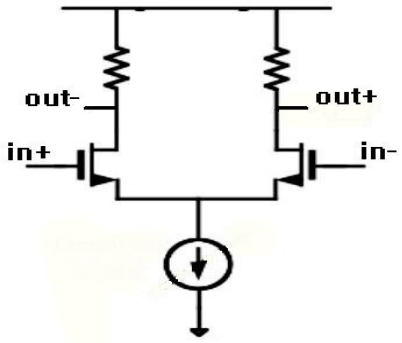

Figure 22 Differential Amplifier

$$\text{Differential gain} = g_m.R_L/2 \qquad \text{eq. (2.8).}$$

Hence, the gain be increased by either increasing the g_m or the load resistance, however increasing g_m requires increasing the size of the transistors or the current and both approaches have limitations. Another way is to increase the load resistors but that also decreases bandwidth and hence the gain cannot be increased too much with this approach.

The output of the common-gate amplifier stage is one of the inputs to the differential pair while the common-mode voltage that is extracted from differential lines is the other input. The final design after considering these trade-offs has a gain of 10dB and consumes 2.16mW.

4.5 CML Flip-Flop

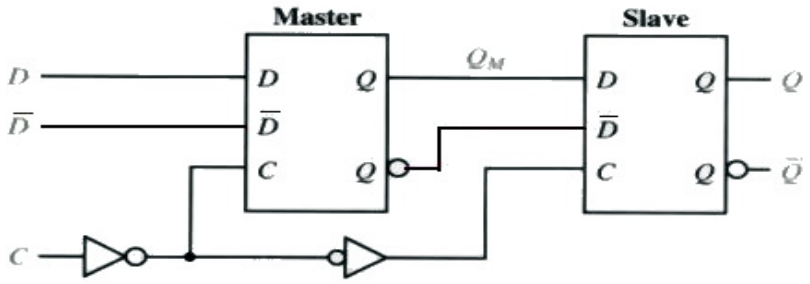

Figure 23 Master Slave D Flip-flop

The amplified output from the differential pair is fed to a decision flip flop which decides if the current symbol is high or a low and also stores the data for one full clock cycle. A common type of flip-flop topology is the master-slave configuration which is created by connecting two D-latches in series. The CML latches have already been discussed earlier in this chapter. The latches are clocked with complementary clock signals and the clock connections can be interchanged to obtain positive edge triggered or negative edge triggered flip-flops.

The master latch will be transparent when the clock is low and the data appears at the master latch output, denoted by Q_m here. However, the slave latch is closed at this time and the data at the output is not affected. When the clock goes high, the master latch goes into hold mode and does not propagate any change at the input but holds the data that was present during track mode. Also, the slave latch gets enabled at this time and the data at Q_m now appears at the output. Hence, the data appears at the output when the clock makes the transition from low to high and

the output does not change at any other time. If the clock connections are interchanged for the master and the slave, it would work the same way but the data would appear at the falling edges of the clock and hence it will be a negative-edge triggered flip-flop. Also, in the case of CML D-latches, differential clocks are used and hence the inverters shown in figure 21 are not required.

4.6 Pseudo Common-mode Rejection

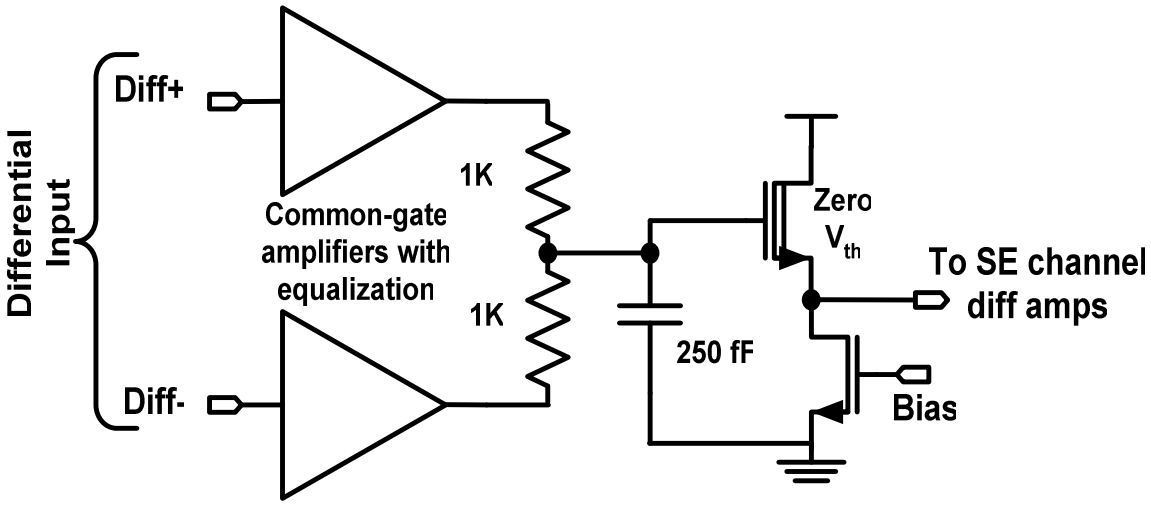

Figure 24 Pseudo Common-mode Rejection Scheme

Figure 24 shows the schematic for extracting the common-mode voltage from each of the two differential lines present in the bus. The differential signals are received using a couple of common-gate receivers that were described earlier. The common-mode voltage is then obtained from the output of the common-gate receivers by taking the average of the differential signals by means of two 1 KΩ resistors as shown. This common-mode voltage needs to be distributed to seven single-ended channels. However, the common-gate amplifier is not able to drive the high output capacitances at the inputs of the seven differential amplifiers with high enough

bandwidth. As a result, a source-follower amplifier is used as a buffer to distribute this common-mode voltage.

A source follower has a gain near unity and level shifts the output. The output voltage is related to the input gate voltage as described by eq. (2.9).

$$V_{GS} = V_{in} - V_o = V_{TO} + \sqrt{\frac{2I_{DS}}{\beta}} \qquad \text{eq. (2.9).}$$

where $\beta = K.(W/L)$ and V_{TO} is the threshold voltage of the transistor. Also since the current is low and W/L is high, the second term is small compared to V_{TO} and the offset can roughly assumed to be equal to the threshold of the transistor. However, the threshold voltage is significantly high for a standard 90 nm transistor compared to the supply voltage and it is difficult to compensate for this drop. Hence a zero-threshold transistor is used, which has a very low threshold (close to 50mV), and has a manageable level shift. Another addition to the circuit was the extra 250fF capacitor at the input of the source follower. This was added to filter out any spurious common-mode variations that occur at very high frequencies. The main reason for the occurrence of this high frequency common-mode variation could be a mismatch in the two differential channels which would cause the common-mode for the differential channels to vary at high frequencies while the common-mode on the single-ended channels is unaffected. Hence, this variation needs to be filtered out and the capacitor provides simple RC filtering along with the 1K resistors and the input capacitances for the source follower. Again, the assumption made here is that any common-mode noise will affect all the channels on the bus and will not be at very high frequencies.

4.7 Control Circuits

4.7.1 Digital to Analog Converters (DACs)

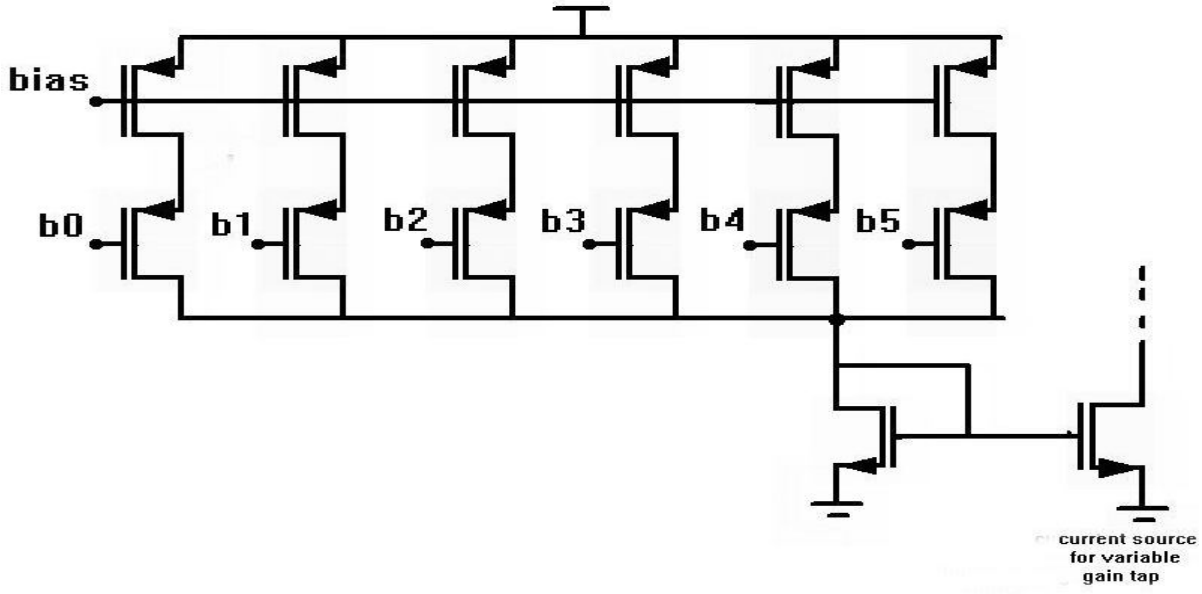

Figure 25 DAC for Transmitter Equalization and Crosstalk Cancellation Modules

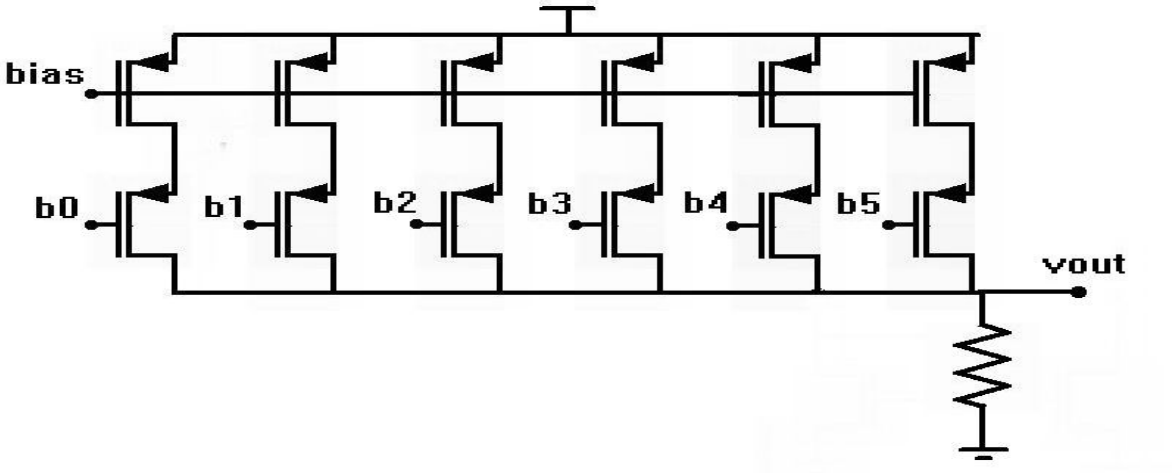

Figure 26 DAC for Common-mode Control at the Receiver

Several DACs are needed for controlling and setting the various parameters throughout the design. Each variable gain tap in the equalization filter is driven by a variable current sink and a DAC is required to set the value of this current. The same setup is required for the crosstalk cancellation filter taps. As each channel has 5 taps for the equalization filter and 5 taps each for the 2 crosstalk cancellation filters, a total of 15 DACs are required for each channel and these DACs are like the one shown in figure 25.

The current source driving the tap is in the form of a current mirror where the current through the transistor on the right of the current mirror can be controlled. A series of six PMOS current sources are connected in parallel so that the current from all of them add. These sources are binary weighted, meaning each source has twice as much current as the previous one and can be connected or disconnected from the circuit by means of PMOS switches. The PMOS switches are further controlled by a 6-bit digital register. The NMOS current mirror transistors can be sized appropriately to get the required range of current drive.

Only one DAC is needed at the receiver for controlling the common-mode voltage for the common-gate amplifier receiver. The receiver DAC is very similar to the transmitter DACs but since it should have a voltage output instead of a current output, a resistor is used instead of the current mirror arrangement. Hence the sum of all the currents from the PMOS sources is summed and a resistor is used to create a voltage. The range of the DAC can again be controlled by changing the value of the resistor. Another thing to note here is that the speed of the DAC is not of concern because once set, the DAC values do not need to be changed during normal operation.

4.7.2 Phase Adjust Delay Lines

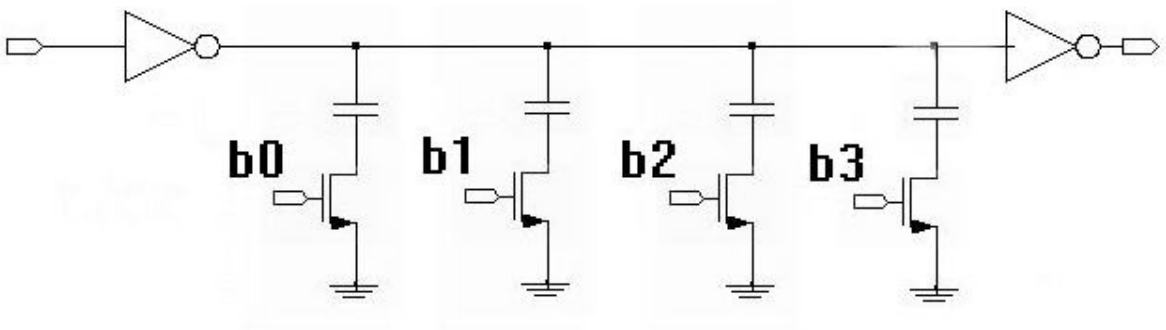

Figure 27 Basic configuration for the delay line

The design is source synchronous meaning the same clock source is used for both the transmitter and the receiver. However, some delay is always associated with the transmission channels and hence the data reaches the receiver out of synchronization. In a communication system that is not source synchronous, this issue is resolved by using clock and data recovery circuits like the Phase Lock Loop (PLL) to extract the clock from the data. However, these techniques are very complex and are not needed with a source synchronous system. All that is needed is to provide a variable clock delay at the receiver. A delay line with adjustable delay is used for each individual channel to delay the clock and place the edges exactly at the center of the data eye.

The basic idea behind the delay line is to let the clock signals pass through a series of capacitive loaded inverters or buffers. The amount of capacitive loading can be changed to adjust the total delay through the delay line as shown in figure 27 where an example of a 4 bit delay line is shown. The capacitors are sized so that when they turn on, they add binary weighted delays to the line. The capacitors are turned on or off by using NMOS switches which are controlled by digital registers. The delay line used for this thesis is a 6-bit delay line with 6ps

resolution and 188ps of maximum delay range, which is larger than one bit time of 125 ps at 8 Gbps.

Chapter 5

5. Simulation and Layout

The functionality of the components was verified using Cadence simulation tools and Synopsys Hspice. Another tool from Synopsys called Cosmoscope was used for viewing the simulation results, most often for its capability to display eye diagrams. The Hspice simulator also supports multi-port S-parameter models which was required since the channel characteristics were available in the form of S-parameter matrices.

Cadence Virtuoso XL was used for layouts for its capability to pick and place from Cadence schematic capture tools. Parasitic extraction was performed on the layouts and the functionality was verified again.

5.1 CML D-latches

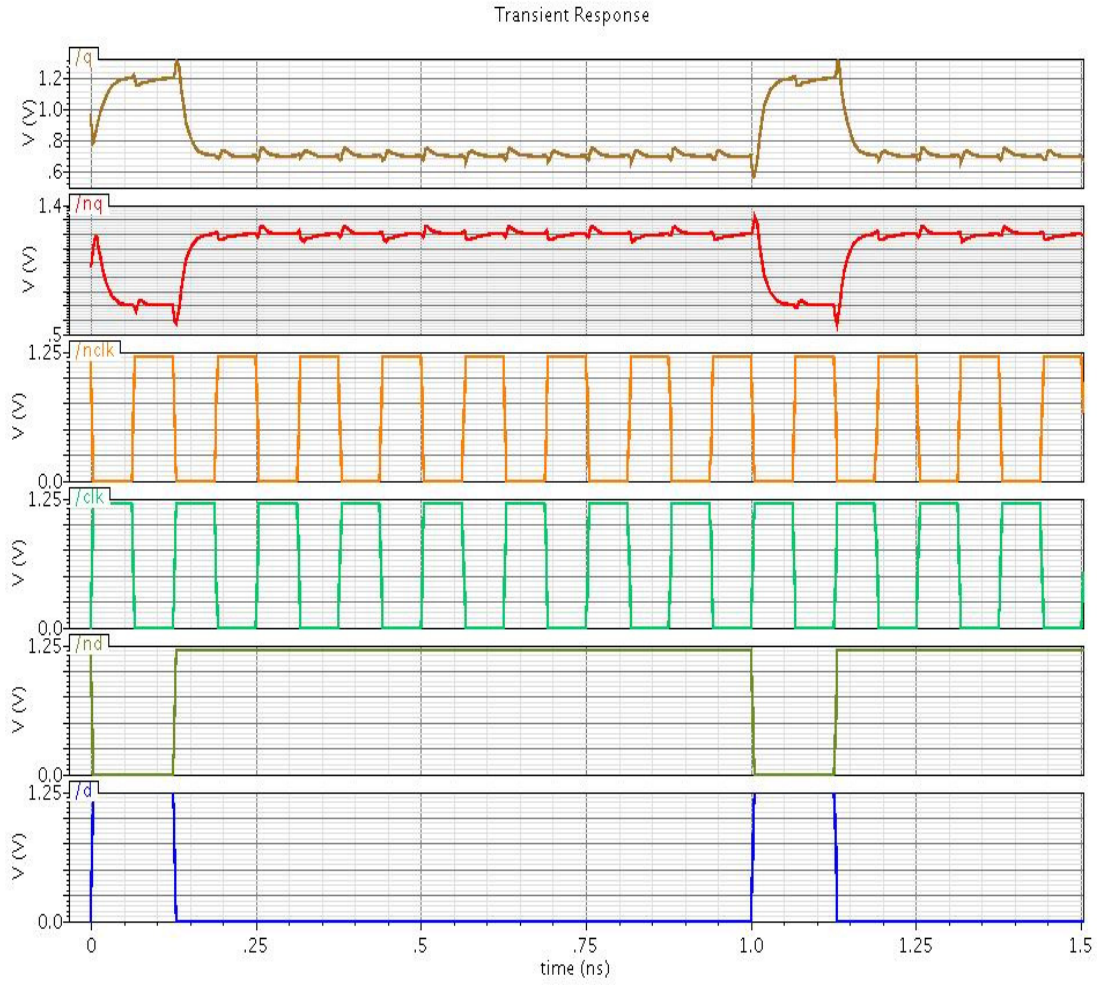

Figure 28 Single D-latch simulation

Figure 28 shows d-latch simulations where the inputs to the d-latch are 125ps, 1.2V data pulses and the output shows the data pulses latched and held correctly. The outputs swing from the latch is around 500mV.

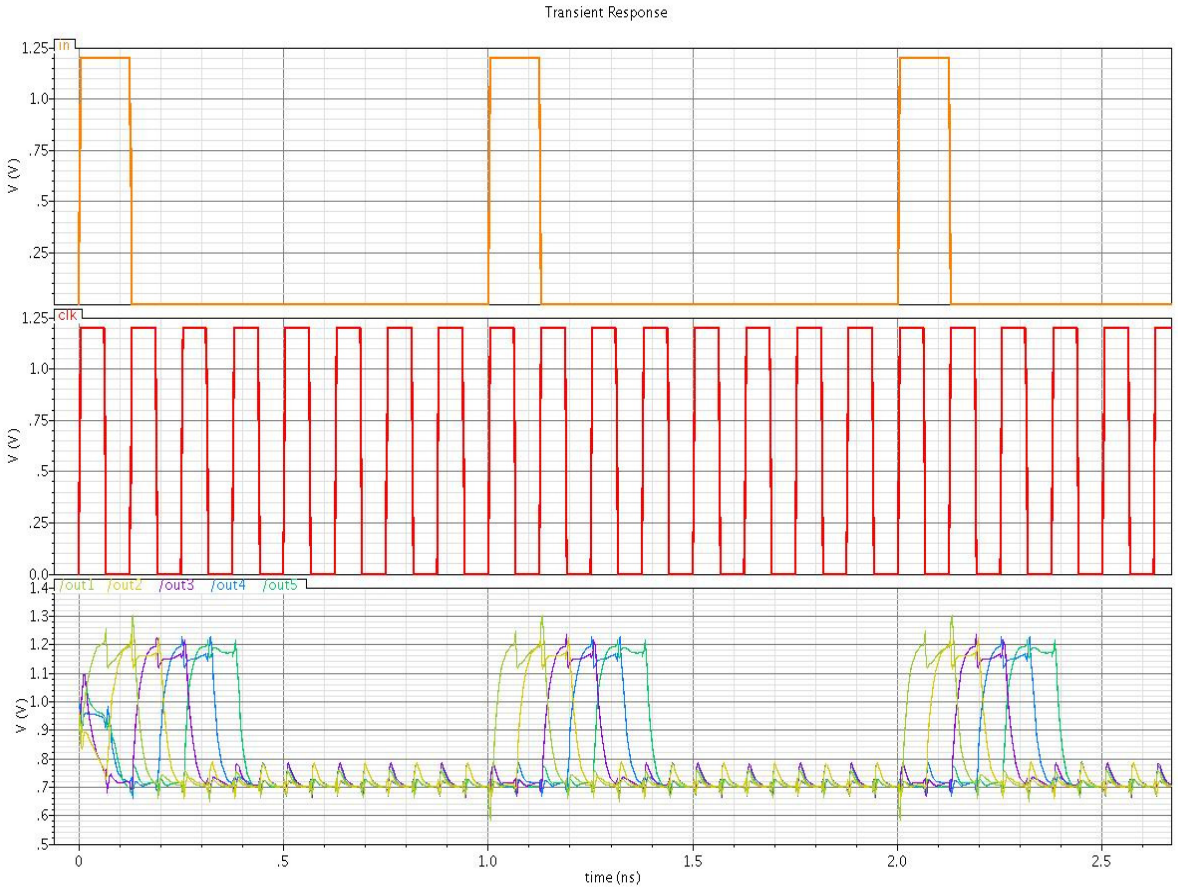

Figure 29 Series of pulses through 5 cascaded D-latches

Figure 29 shows the same data pulse going through 5 latches connected in a cascaded configuration. This configuration is used in the FIR equalization filter. The pulse can be seen going through the five latches and the output from the latches are separated from each other by half a clock cycle.

5.2 CML Multiplexer

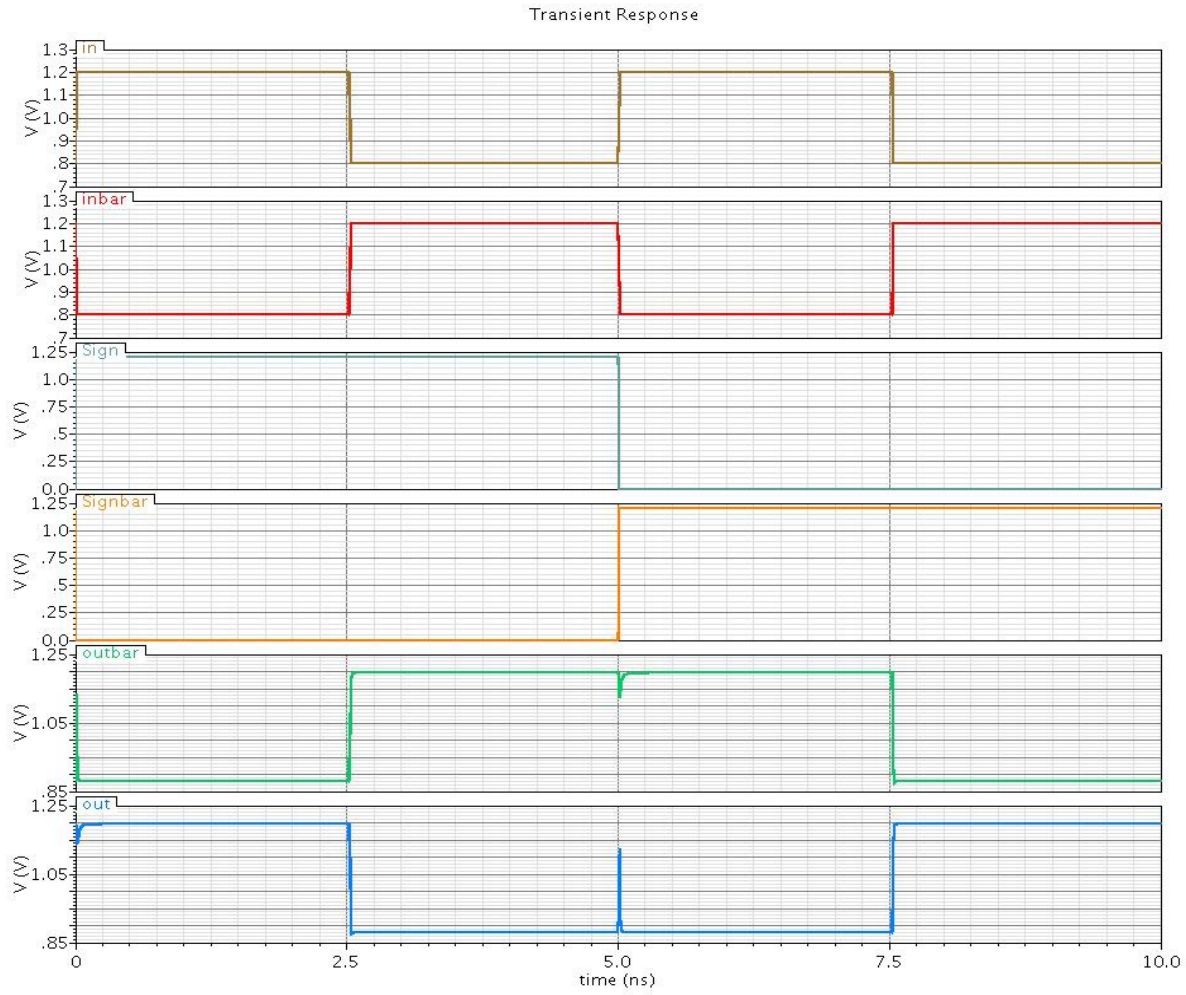

Figure 30 Simulation results for multiplexer gate

The multiplexer is used to set the signs of the taps. The simulation results in figure 30 shows this operation for all possible cases. When the sign bit is high, the input pattern appears as it is at the multiplexer output. However, when the sign bit goes low, the input appears inverted at the output.

5.3 Transmitter Equalization

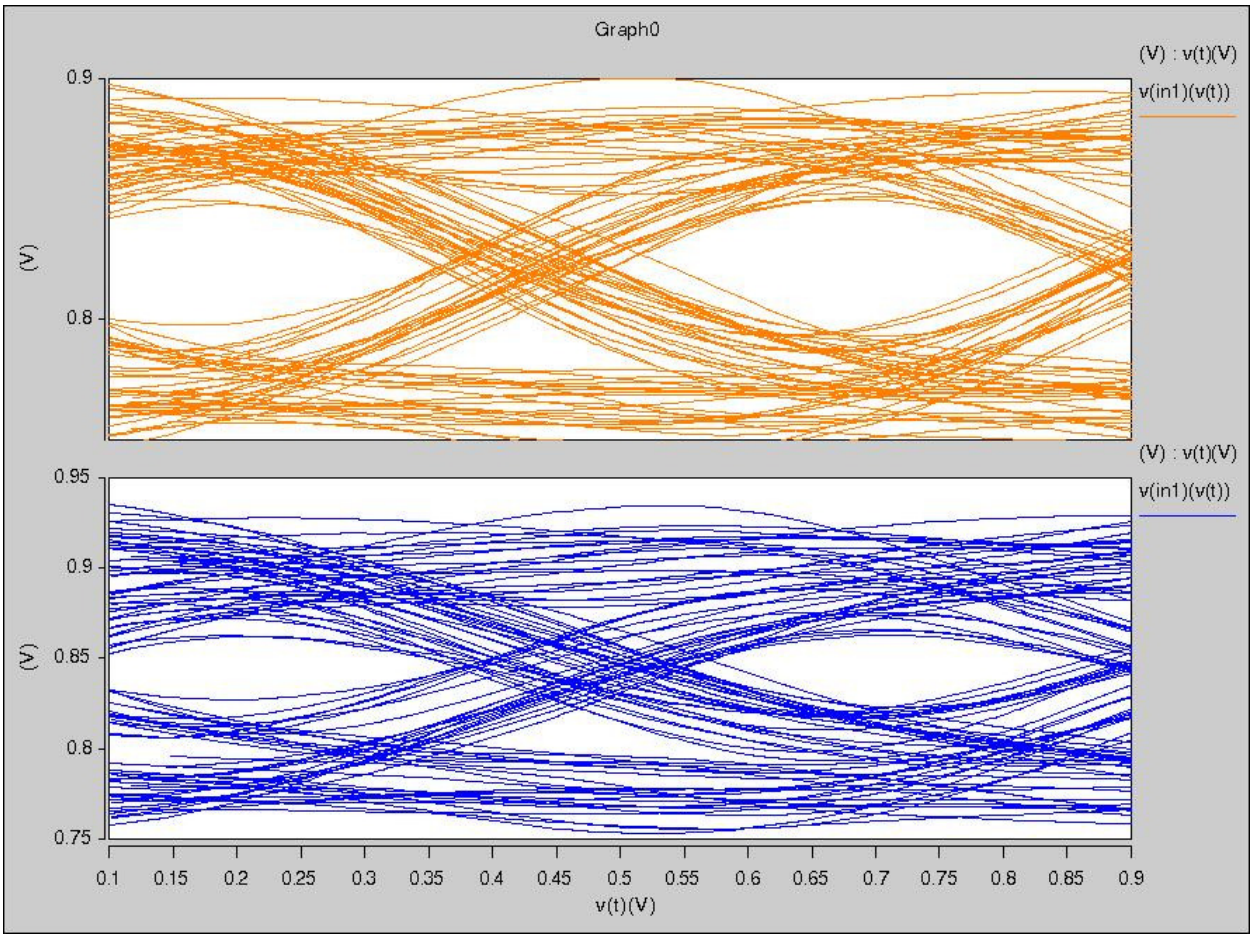

Figure 31 Improvement in Eye Pattern after Transmitter Equalization at 8 Gbps

The eye pattern at the bottom of figure 31 shows the eye pattern received at the receiver with the transmitter and receiver equalization turned off. The eye pattern on the top is obtained when the transmitter equalizer is turned on. An improvement in both the eye width and height can be seen.

5.4 Crosstalk Cancellation

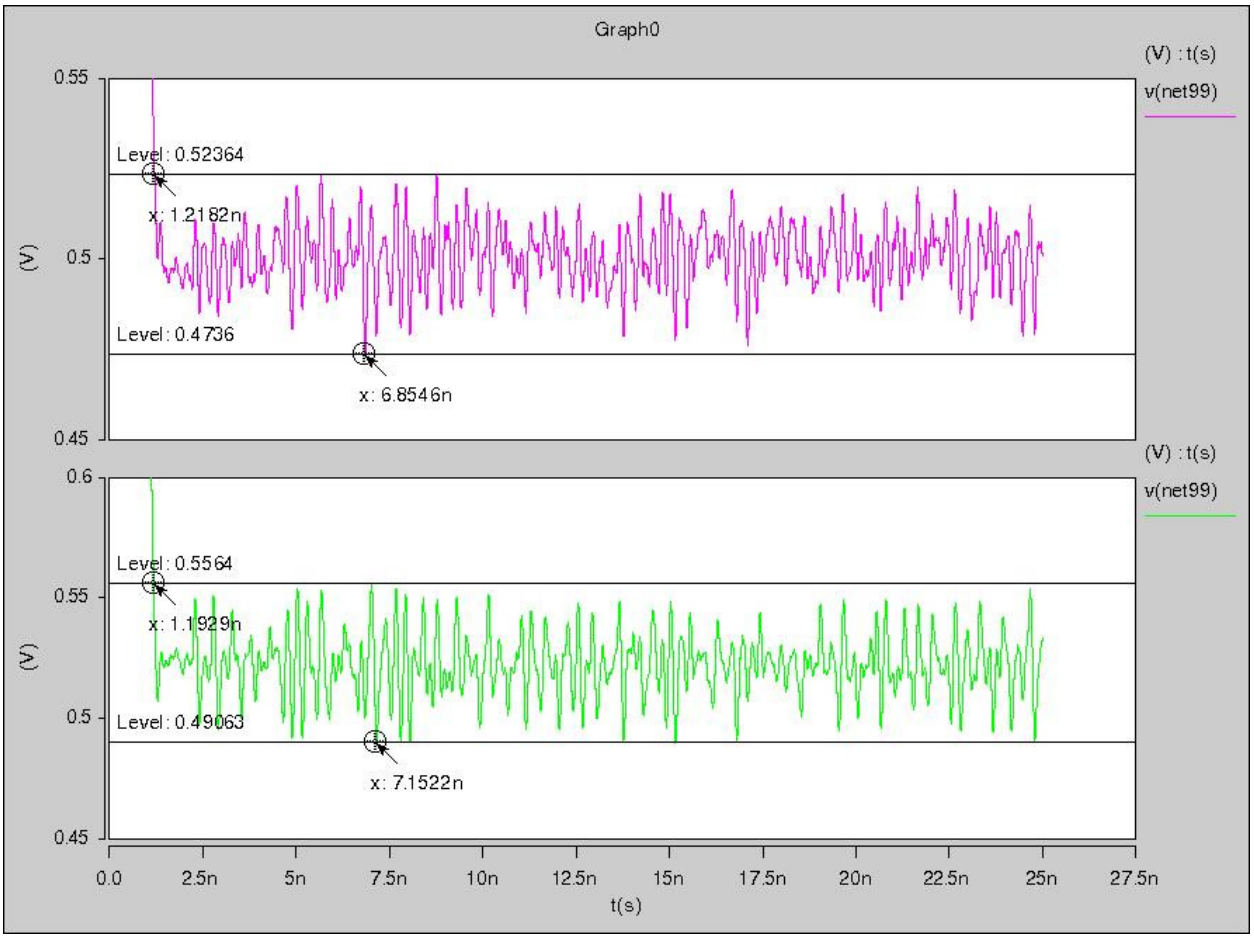

Figure 32 Crosstalk Cancellation

Figure 32 shows the crosstalk on a quiet channel when its 2 nearest neighbors and 2 next nearest neighbors are transmitting pseudo-random data patterns. The bottom waveform is obtained without any crosstalk cancellation and the one on top shows the effect of crosstalk cancellation. The crosstalk levels are reduced by about 30% by using crosstalk cancellation.

5.5 Receiver Equalization

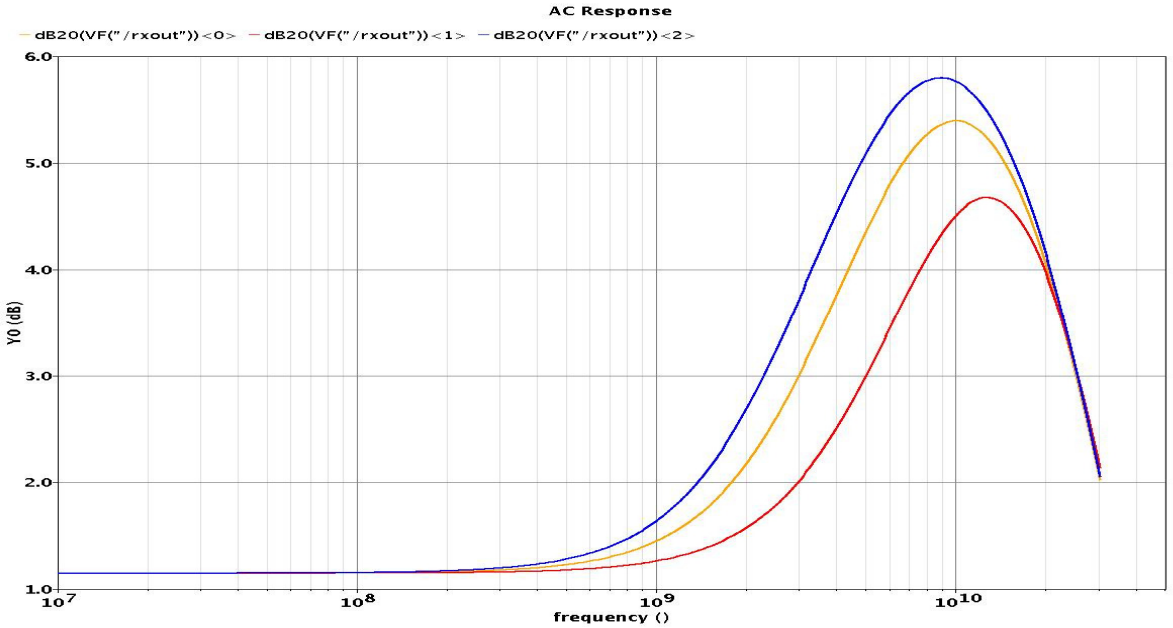

Figure 33 Receiver Equalization Gain Adjustment

Figure 33 shows the gain adjustment for the receiver equalization for 3 different gain settings. A maximum gain of about 5.8dB is possible.

5.6 Common-mode Control

Figure 34 Common-mode Control Simulation

Figure 34 shows the output common-mode voltage variation with changing common-mode control voltage. For a 400mV variation in control voltage, about 350mV change can be achieved in the output common-mode voltage which is sufficient. The common mode output was set at 550mV for the data transmission simulations.

5.7 Receiver Equalization

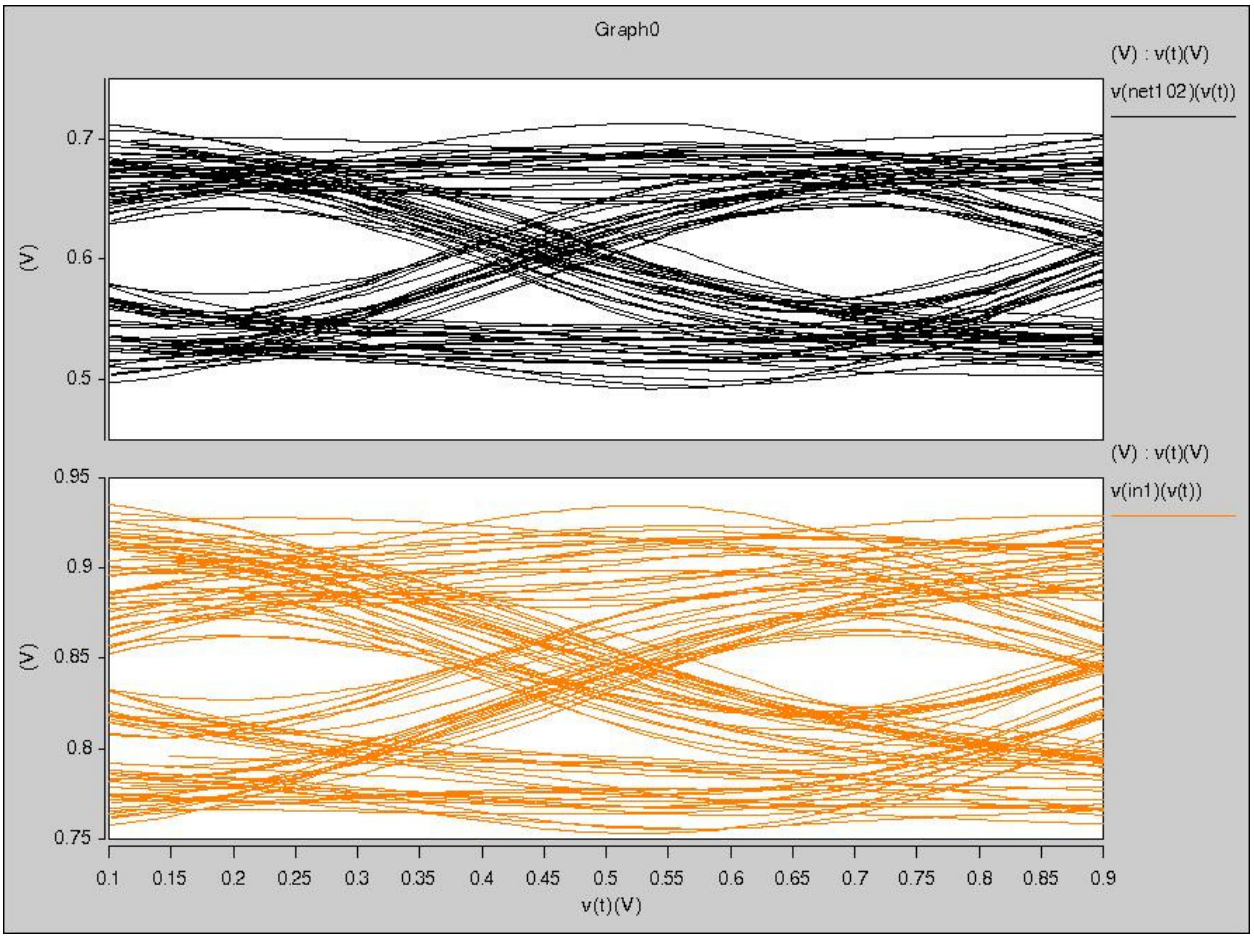

Figure 35 Improvement in Eye Pattern after Receiver Equalization

Figure 35 shows the eye patterns at the receiver when the transmitter equalization is turned off. The eye pattern on the bottom shows the received eye pattern with no receiver equalization, while the one on top shows the improvement with the receiver equalization turned on.

5.8 Common-mode Rejection

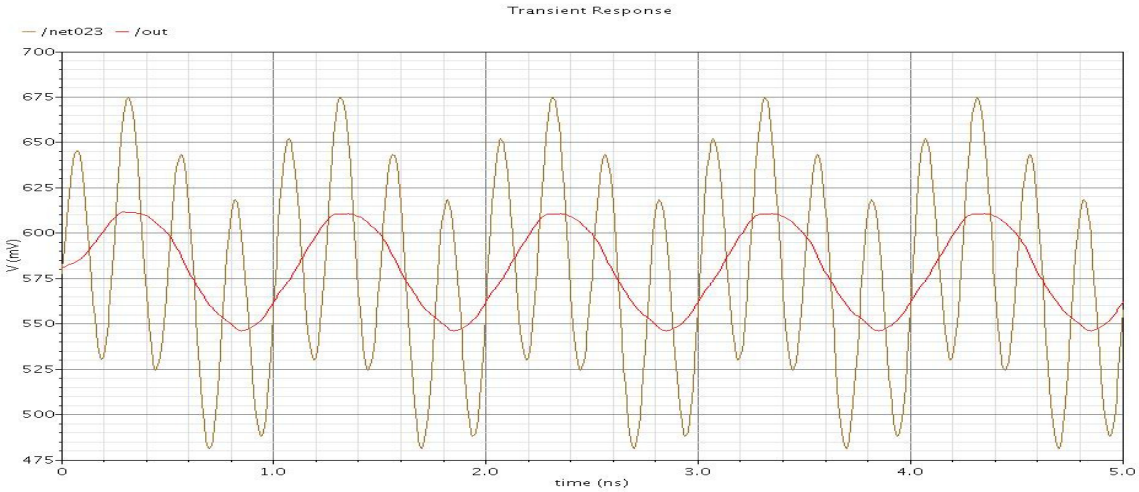

Figure 36 Common-mode Tracking for 1 GHz Common-mode Noise

For testing the common-mode tracking circuit, a continuously switching data pattern was used and a common-mode noise signal was added with 70mV peak-to-peak amplitude at 1 GHz frequency.. The output from the common-mode tracking circuit can be seen to follow the varying common-mode voltage and also varies by almost the same amplitude and frequency as the noise.

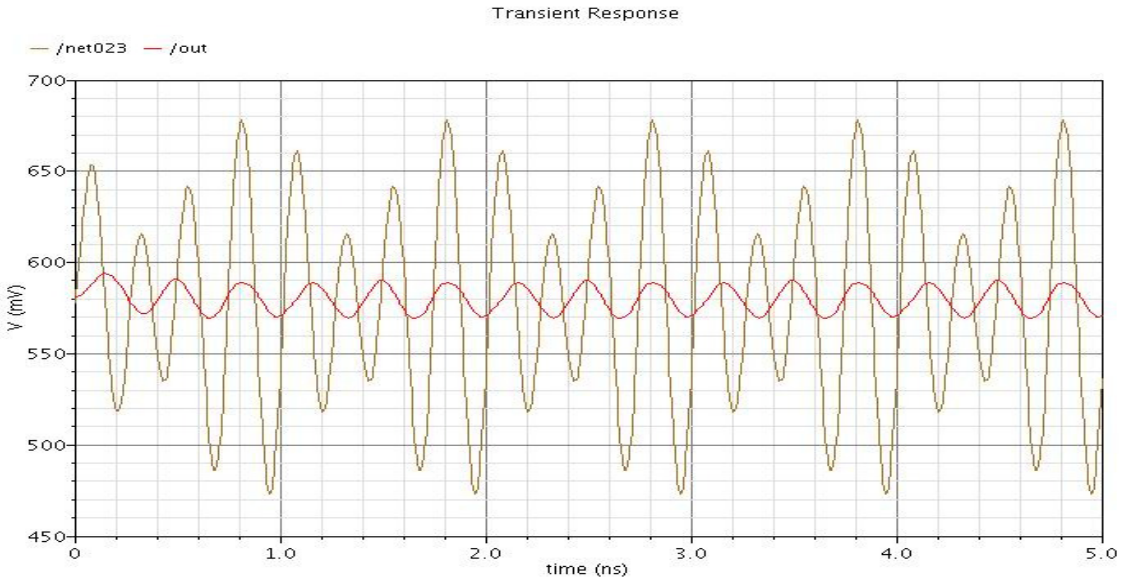

Figure 37 Common-mode Tracking for 4GHz Common-mode Noise

Figure 37 has the same setup for the common-mode tracking circuit as figure 36 but the noise frequency has been increased to 4 GHz. This is done to demonstrate the effect of filtering of common-mode variations at higher frequencies. The input common-mode voltage noise has amplitude of 70 mV peak to peak but this variation is filtered out and the output is only 20mV peak to peak.

5.9 Differential Pair Amplifier

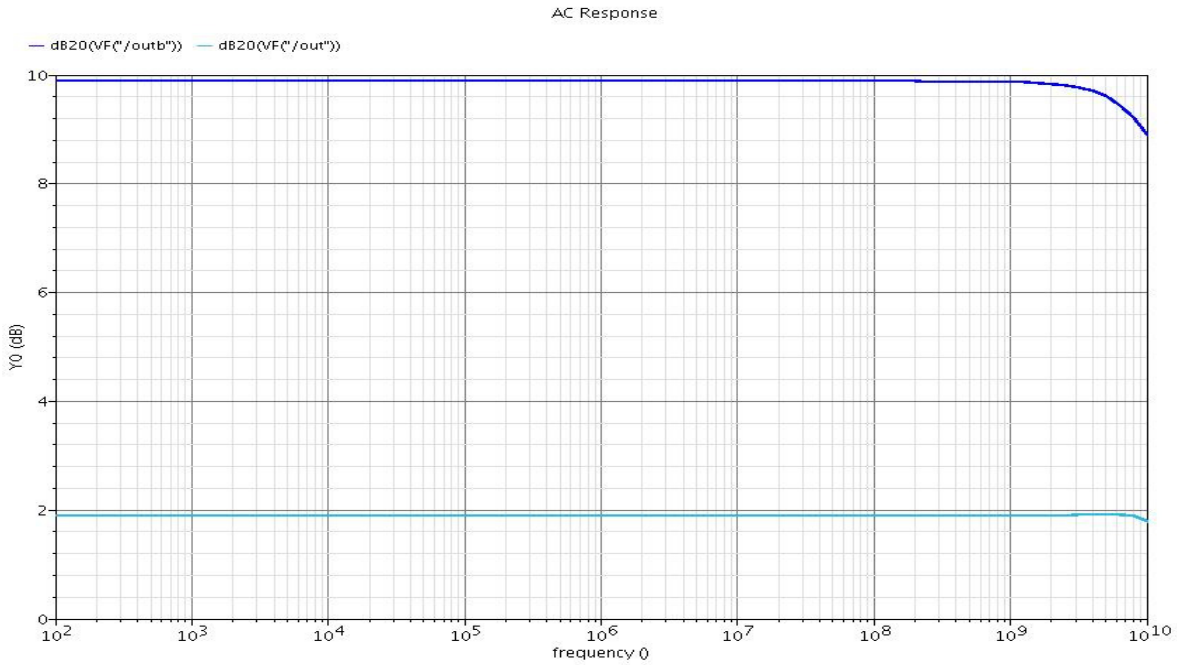

Figure 38 Differential Pair Amplifier

Figure 38 shows the ac response for the differential pair amplifier. The differential gain is nearly flat at 10 dB up to 2 GHz and only drops 1 dB at 10GHz.

5.10 Master-Slave Flip-Flop

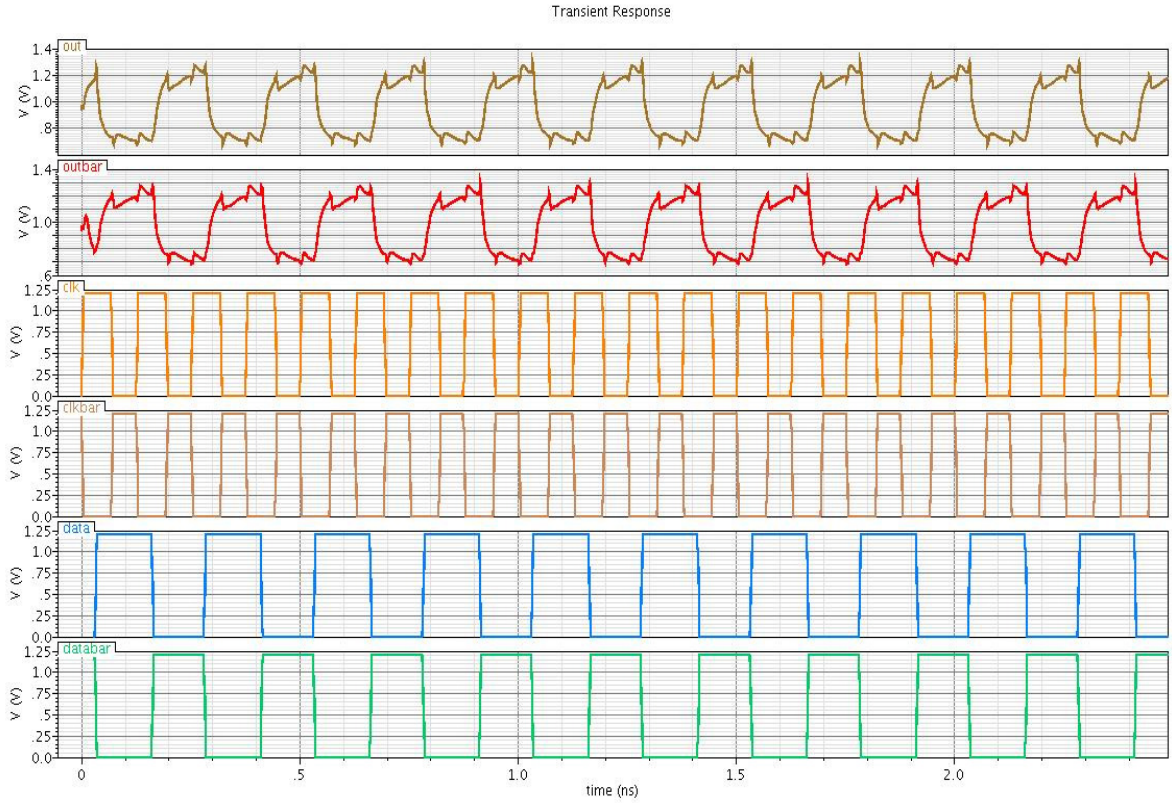

Figure 39 Flip-flop Simulation at 4GHz

Figure 39 shows the simulation for the master slave D flip-flops with a 8GHz clock and a continuously changing data stream at 8 Gbps. This verifies the operation of the flip-flop at the maximum data rate.

5.11 Verification of Data Transmission

The final system simulations were done with 10-port S-parameter models for 5 transmission lines in Hspice. The 5 lines were used to transmit data for three single-ended channels and one differential channel. Lines 1, 2 and 5 carry single-ended data, while lines 3 and 4 are used for differential data. The common-mode voltage information is extracted from the differential

channel and is distributed to the single-ended channels. Equalization is turned on at both the transmitter and the receiver and crosstalk cancellation is also turned on.

Figure 41 shows the eye patterns after receiver equalization when a 200 bit pseudo random data sequence was transmitted over the 4 channels. The same data patterns were sent over channels 1 and 3, and also 2 and 4. This is done to simulate the worst case crosstalk scenario when both the nearest neighbors transmit the same data pattern and the crosstalk signals from both the channels always gets added. The signal on channel 2 is of most interest here since it has crosstalk coupling from both its neighbors while the first and last channels have crosstalk coupled from only one channel each. The signals in1, in2 and in4 are single-ended channel eye patterns in figure 41 while in3 is the differential channel eye pattern. As expected, the differential channel eye is the cleanest due to its tolerance to crosstalk. Figure 40 shows the eye patterns for the same test setup but with equalization and crosstalk reduction turned off for comparison with the eye patterns in figure 41.

Figure 42 shows the transmitted data patterns generated using pseudo-random sequence generators along with data patterns received at the final output of the flip flops for each channel. It should be noted that the data patterns are received with a 1.3 ns delay due to the delay present in the transmission lines. However, the transmitted and received data patterns are identical indicating no errors on any of the channels.

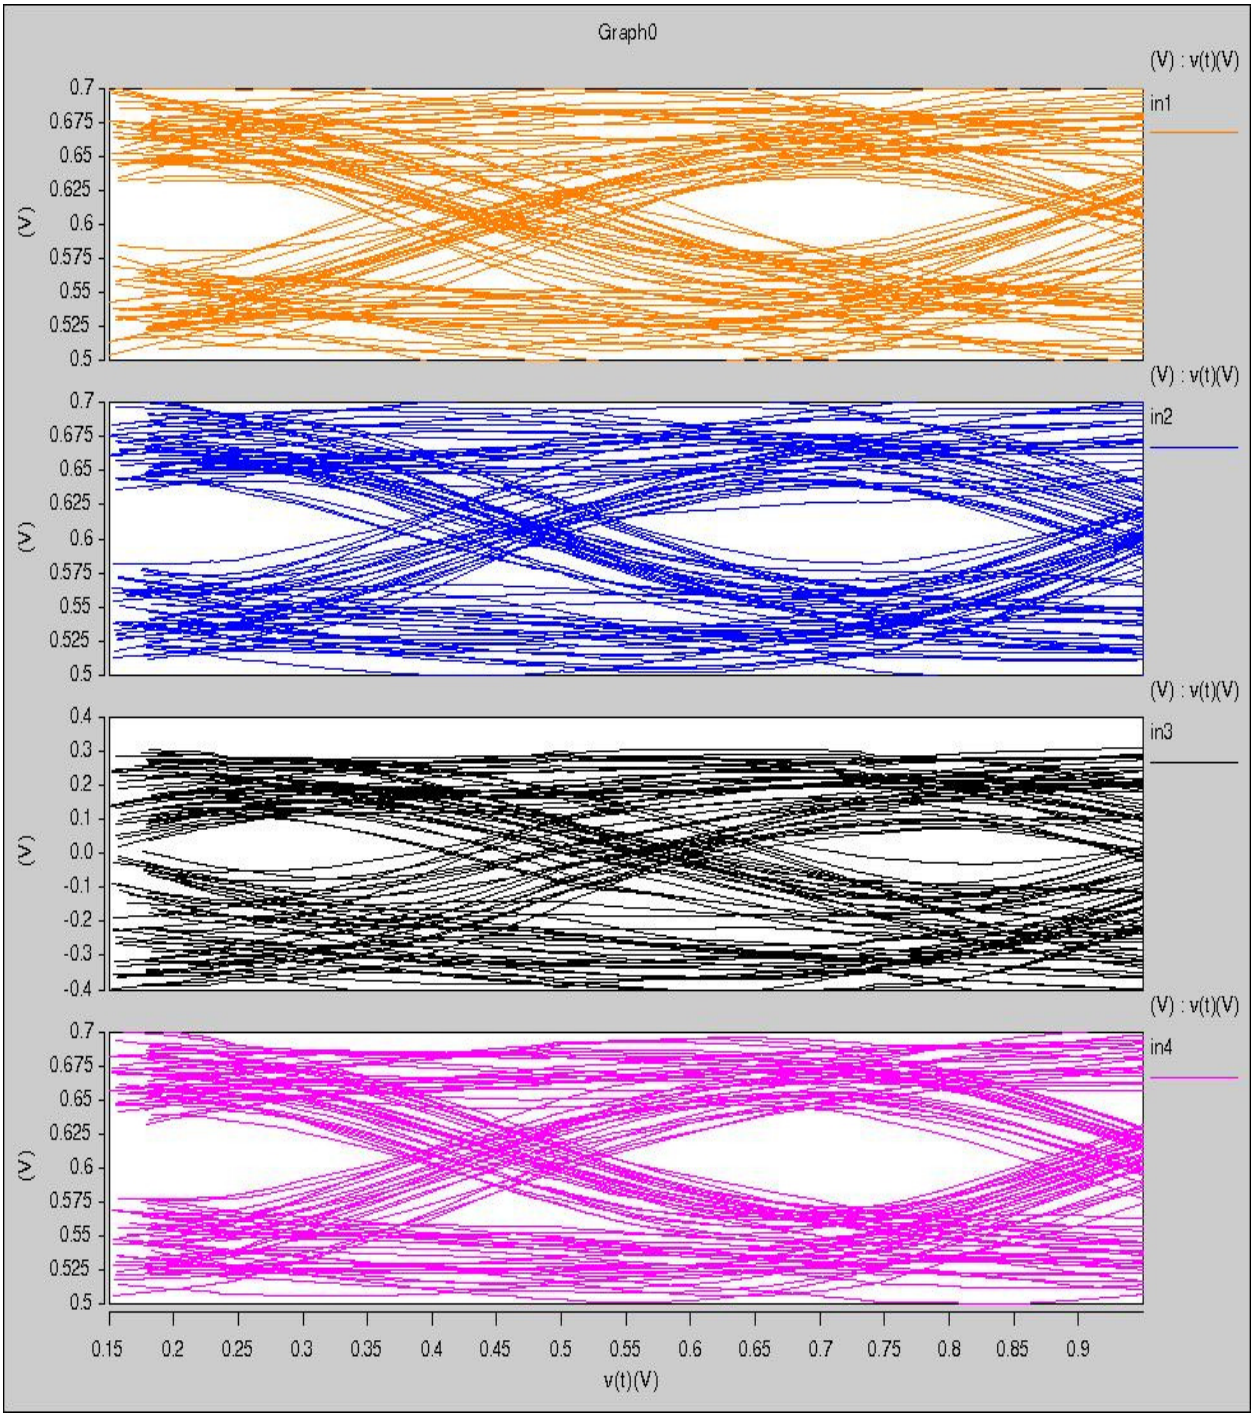

Figure 40 Received eye patterns with equalization and crosstalk cancellation turned off

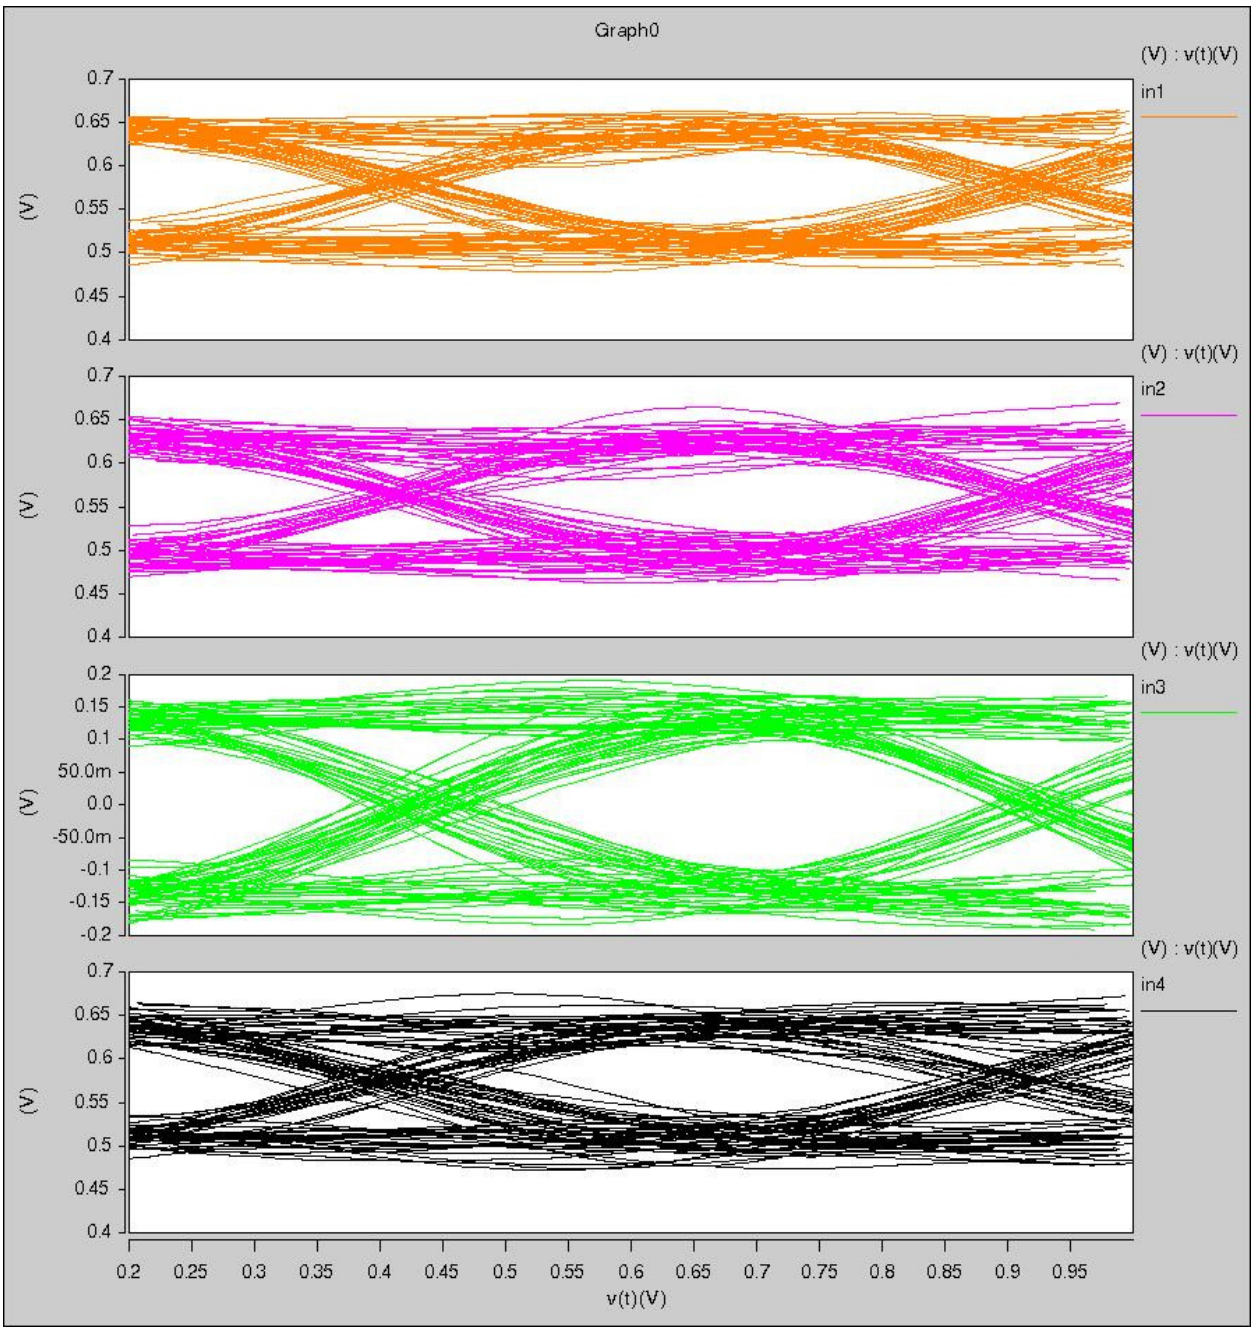

Figure 41 Received eye patterns with equalization and crosstalk cancellation turned on

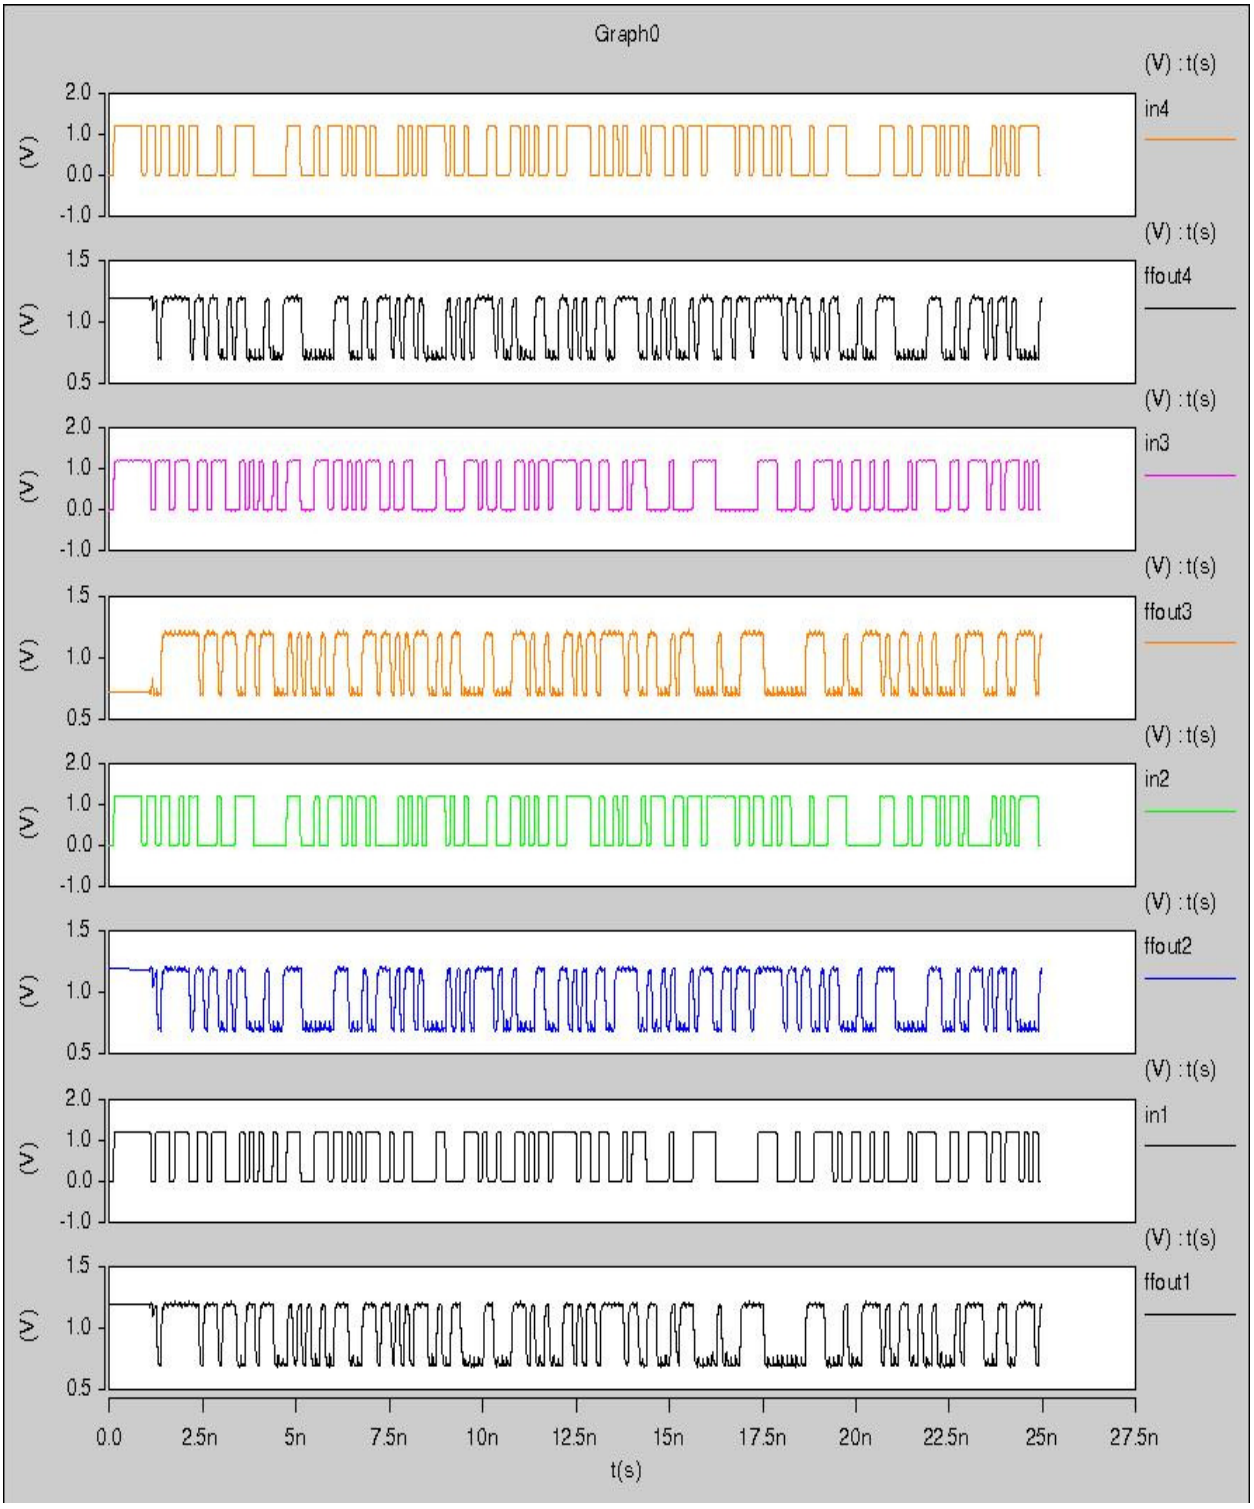

Figure 42 Transmitted and received data streams

5.12 Transmitter DAC

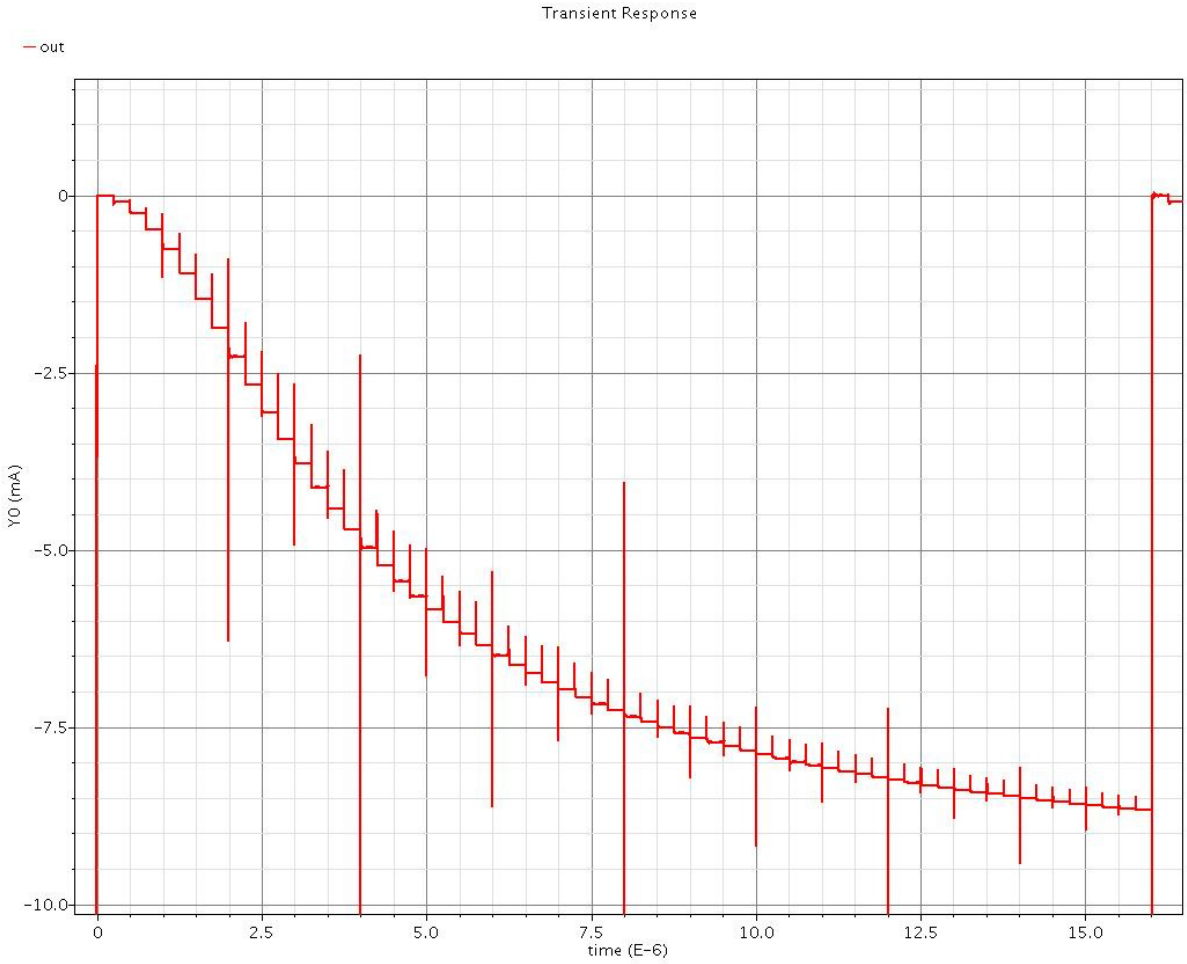

Figure 43 Transmitter DAC outputs over the full 8.7mA range

Figure 43 shows the DAC output steps for all possible 6-bit inputs for the main tap of the transmitter equalizer. The range is 0-8.7 mA and output steps get smaller when the current gets large. However, since the main tap always draws at least 5mA of current, the bigger steps at the beginning are not of concern. The DACs for the other taps in the equalizer and the crosstalk equalization schemes have a maximum current output of 2 mA or less and are very linear.

5.13 Receiver DAC

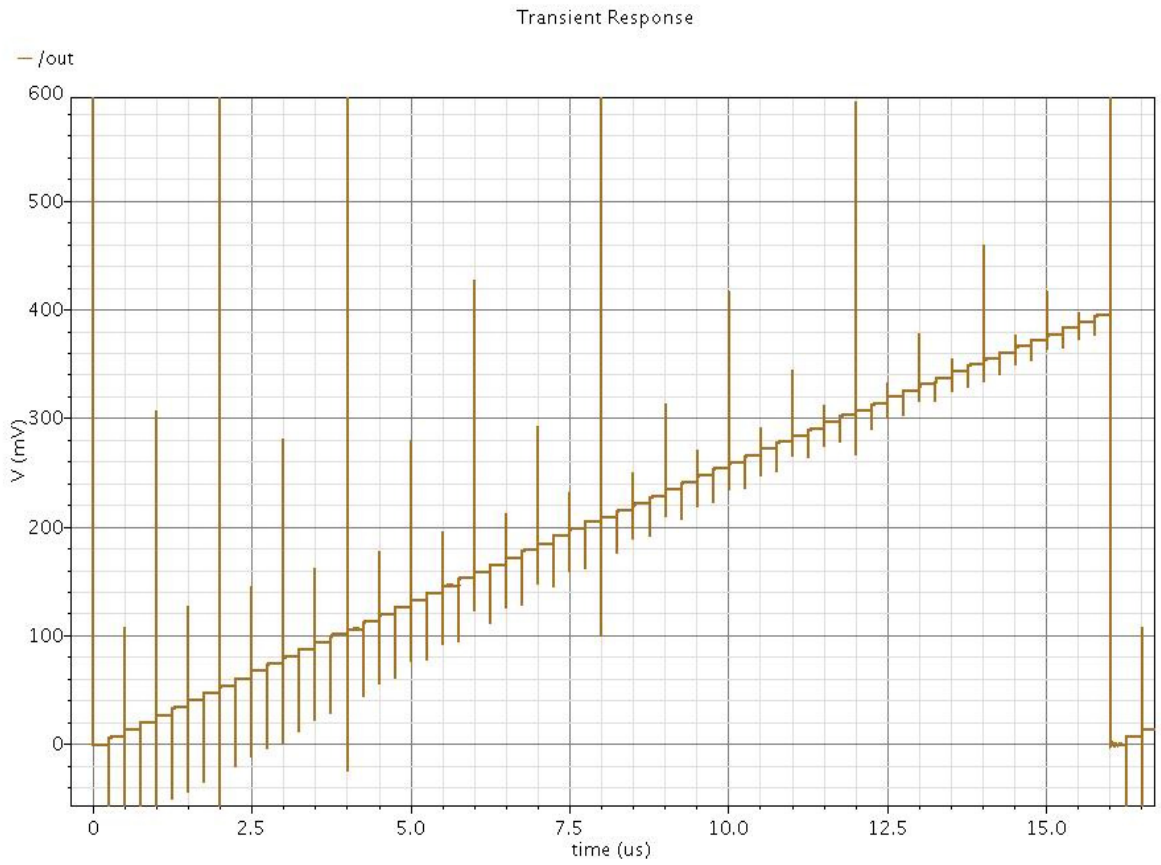

Figure 44 Receiver DAC output steps

The output steps for all possible 6-bit inputs for the receiver common-mode voltage control DAC are shown in figure 44. The range is selected to be around 400mV with step sizes of 6.3mV.

5.14 Phase Adjust Delay Lines

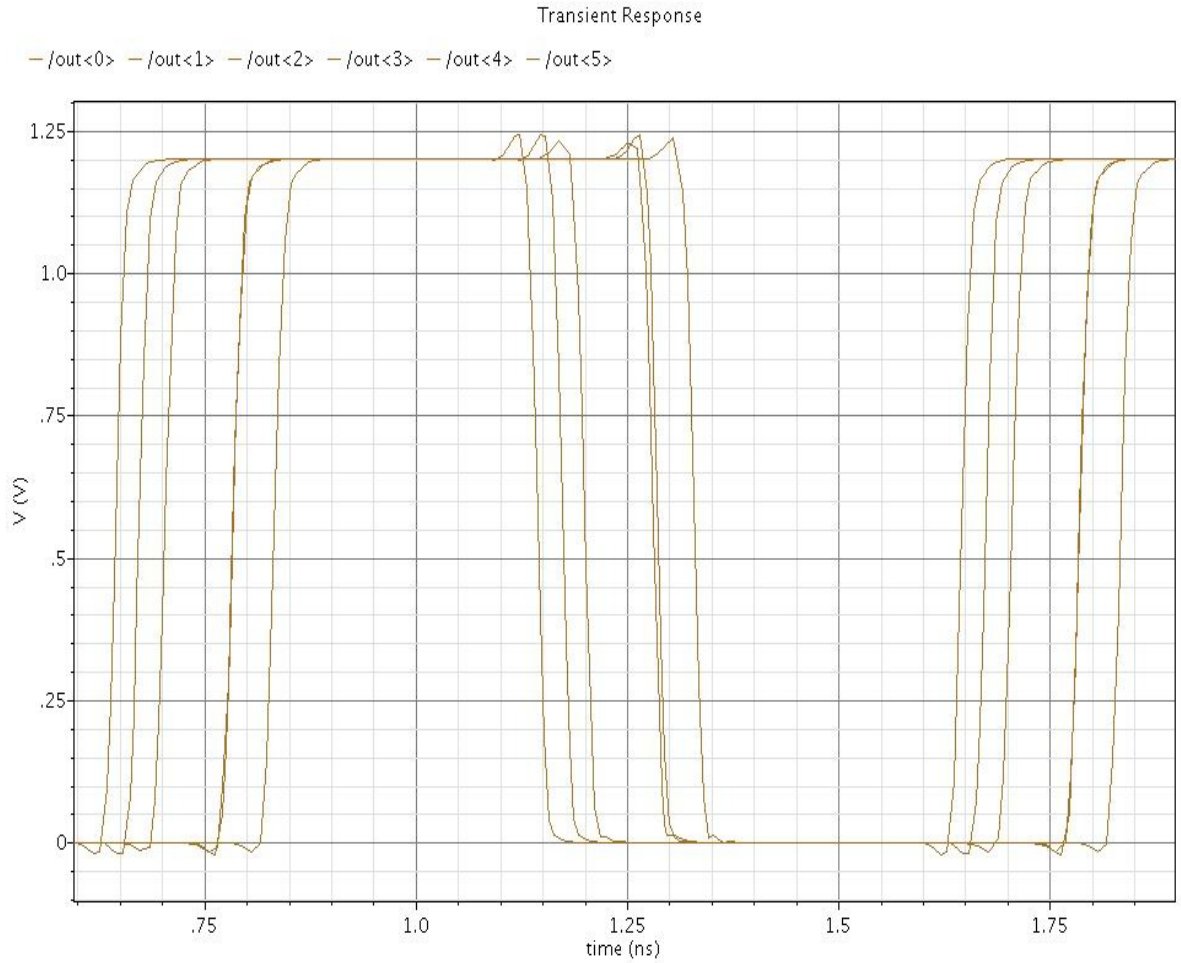

Figure 45 Delayed clock outputs for different control words

Figure 45 shows the effects of delaying the clock at the receiver by using the phase adjust delay lines. The clock outputs for different control words as inputs to the delay line can be seen. The output for a few random input words are shown, including 000000 (minimum delay) and 111111 (maximum delay). The range of the delay line is determined to be about 187 ps and the resolution or step size is about 6 ps.

5.15 Layout of components

5.15.1 CML D-latch

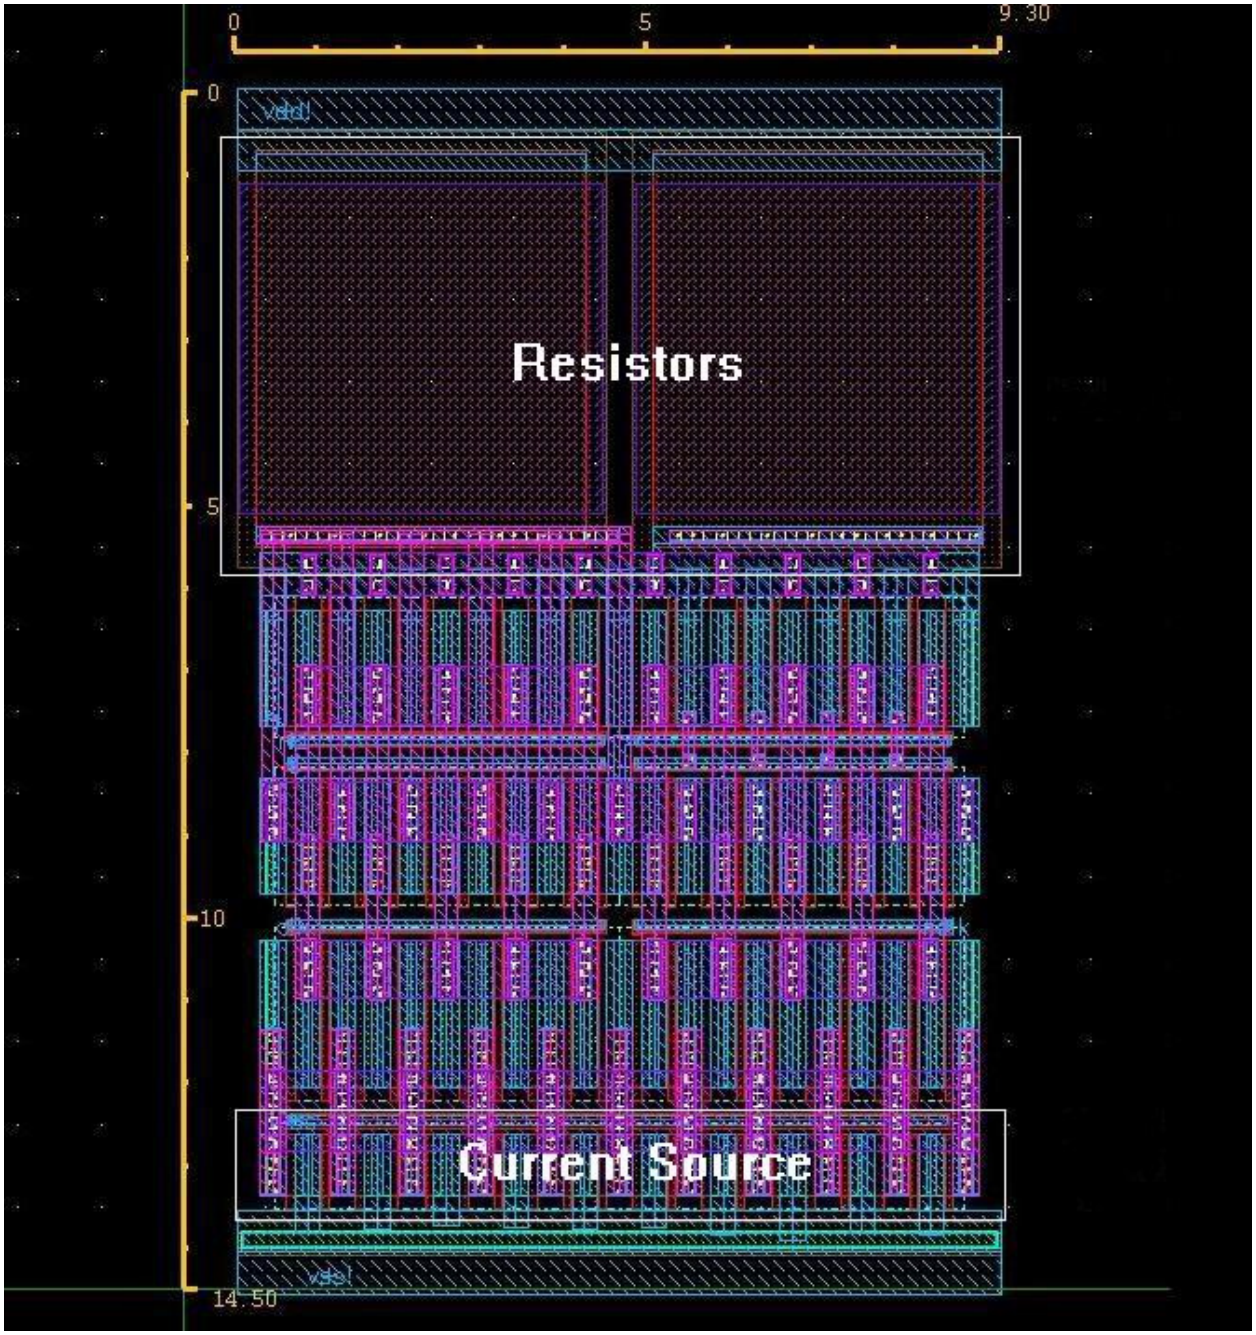

Figure 46 Layout for CML D-latch

5.15.2 Transmitter Equalization FIR Filter

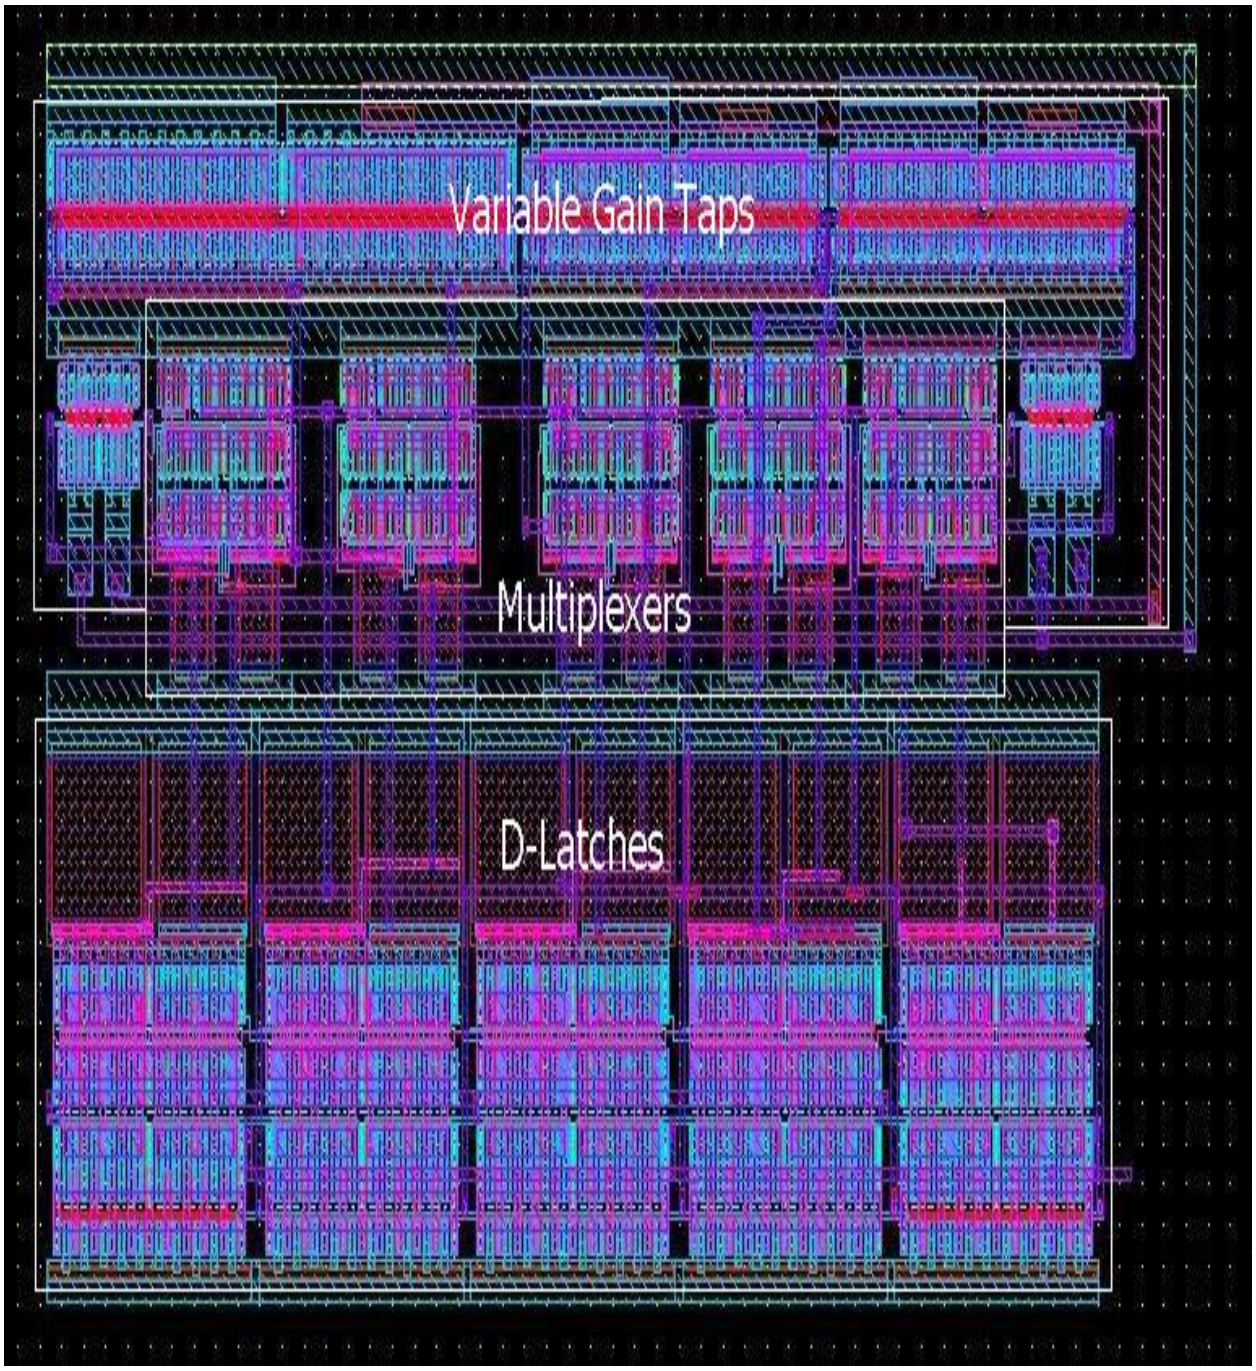

Figure 47 Layout for transmitter equalizer with 5 tunable taps

5.15.3 Complete transmitter for single channel

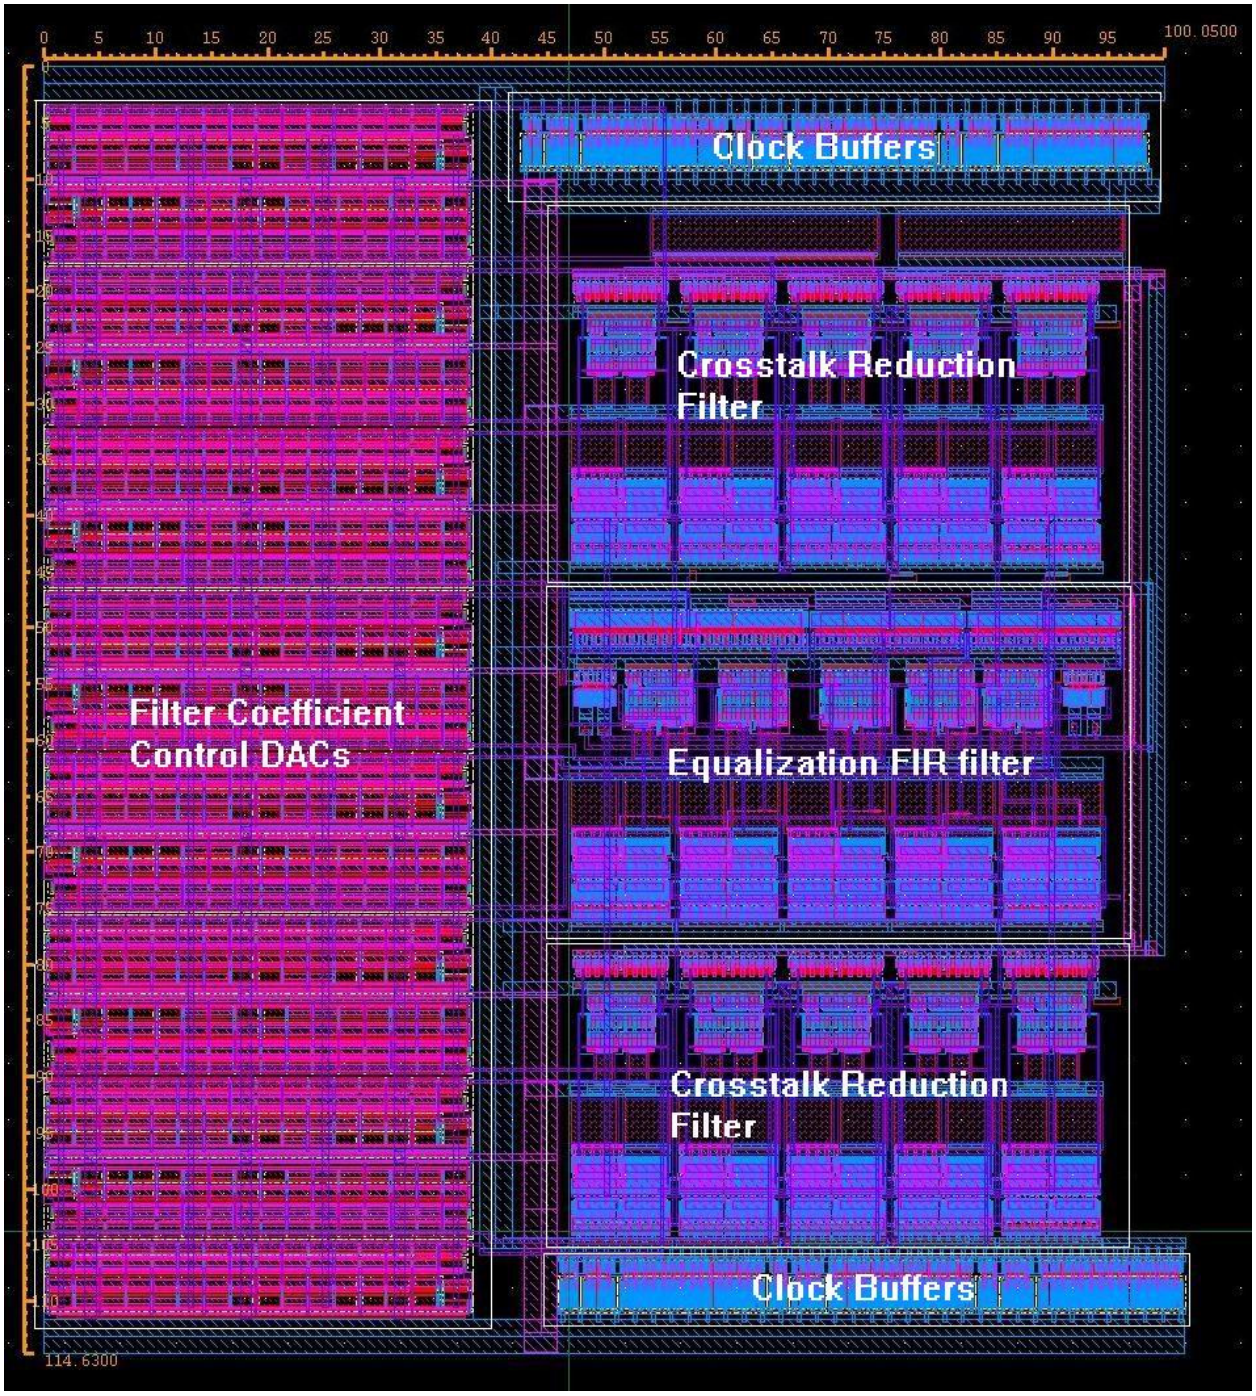

Figure 48 Complete single channel transmitter layout with equalization, crosstalk cancellation and control DACs (Layout Area is 100x115um)

5.15.4 Complete receiver for single channel

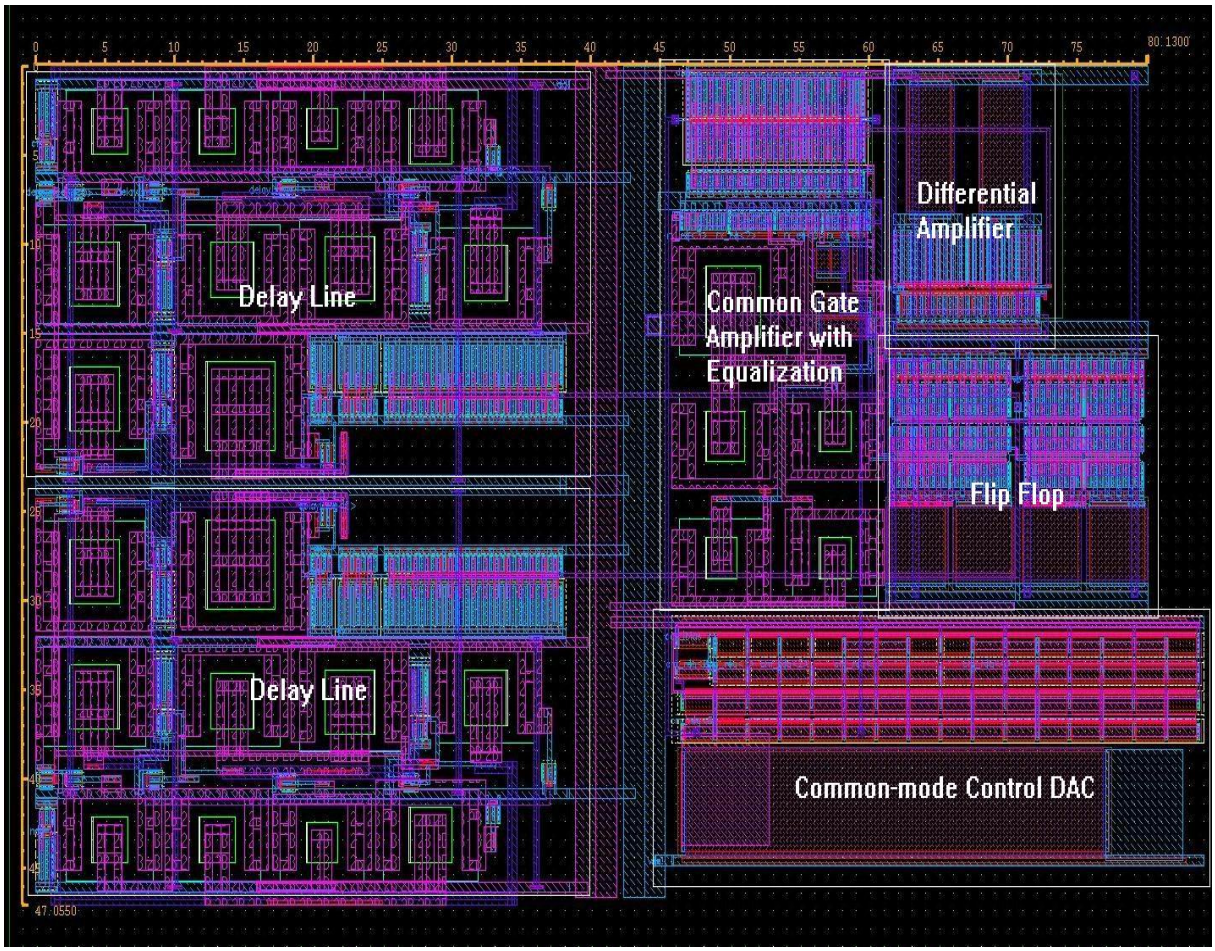

Figure 49 Complete single channel receiver layout with control DAC and delay lines for phase adjustment (Layout Area is 80x47um)

Chapter 6

6. Conclusion

The aim of this project was to design a 16-bit single-ended bus with data speeds of 8 Gbps on each channel. Some PCB's were fabricated to select the appropriate arrangement for the channels and the S-parameters were measured. 10-port S-parameter models for simulating 5 cross coupled channels were generated based on the measurements. A 5 bit FIR filter was used to equalize the losses at the transmitter and similar filters were used to cancel crosstalk from the two nearest neighbors. A common-gate receiver amplifier with added equalization was used at the receiver to achieve better overall equalization. Two differential channels were included to provide common-mode voltage information for the single-ended channels at the receiver. A coding scheme was developed to eliminate power supply bounce due to single-ended signaling. Control circuits including DACs and delay lines were designed to perform several control operations.

The design and layout for a single-ended channel was completed. Simulations were done for 5 channels transmitting pseudo-random data. The results show that data transmission is reliable and error free at 8 Gbps. The simulations were done with extracted layout netlists including parasitics. Problems associated with single-ended buses like crosstalk, common-mode rejection and power supply bounce were mitigated and the data speed of 8 Gbps is comparable to that for differential busses. The equalization schemes have the capability to achieve higher data rates on these channels but crosstalk becomes substantially higher beyond 8 Gbps. Hence, crosstalk is a major limiting factor for single-ended signaling at higher speeds and better crosstalk cancellation schemes would be required.

6.1 Future Work

The layout for the 16 channel transmitter and receiver chips needs to be completed. Also, the test structures including the pseudo-random sequence generators and the data multiplexers need to be designed. A differential bus with similar equalization and crosstalk cancellation schemes is to be included on the same chip for comparison with the single-ended bus. The chip layouts are expected to be completed this fall and the chips will be sent out for fabrication. Characterization of the chips will then follow fabrication. The single-ended bus needs to be tested for variable lengths of lines and compared with differential busses on the same chip. Any future work on single-ended busses needs to focus more on crosstalk cancellation techniques. Also, adaptive equalization schemes should be considered to eliminate the process of setting the equalization coefficients manually for each channel.

Bibliography

[1] Stefanos, S., et al., “High-Speed Inter-Chip Signaling,” in Design of High-Performance Microprocessor Circuits, A. Chandrakasan, W. Bowhill, and F. Fox, (eds.), New York:IEEE Press, 2001, pp. 397–42.

[2] Farjad-Rad, R., et al., “A 0.3m CMOS 8 Gb/s 4-PAM Serial Link Transceiver,” Journal of Solid State Circuits, Vol. 35, No. 5, May 2000, pp. 757–764.

[3] R.E. Matick, Transmission Lines for Digital and Communication Networks, IEEE Press, New York, 1995.

[4] Buckwalter, James F. and Hajimiri Ali, “Cancellation of crosstalk-induced jitter”. IEEE Journal of Solid-State Circuits, 41 (3). pp. 621-632, 2006.

[5] Zerbe, J., et al., “Equalization and Clock Recovery for a 2.5–10 Gb/s 2-PAM/4-PAM Backplane Transceiver Cell,” Paper 4.6, International Solid State Circuits Conference, Santa Clara, CA, February 9–13, 2003.

[6] S. Shankar Narayan and A.M. Peterson, "Frequency Domain LMS algorithm," Proc. IEEE, Vol. 69, No. 1, pp. 124-126, January 1981.

[7] Kin-Joe Sham, Ahmadi, M.R., Talbot, S.B.G., Harjani, R, “FEXT crosstalk cancellation for high-speed serial link design,” Proc. IEEE Custom. Integrated Circuits Conf., pp. 405–408, Sept. 2006.

[8] A.A.A. El-Fattah, F.A.N. Mohamed, A.M. Arafa, M.M. Ahmed, D.R.A. El-Hay, M.O.A. El-Aziz, “Equalizer implementation for 10 Gbps serial data link in 90 nm CMOS technology”, ICM 2007.

[9] Hollis, T.M.; Comer, D.J.; Comer, D.T., “Mitigating ISI through Self-Calibrating Continuous-Time Equalization”, IEEE Transactions on Circuits and Systems I: Regular

Papers, Volume 53, Issue 10, Oct. 2006 Page(s):2234 – 2245.

[10] Mumby, S. J., "Dielectric Properties of FR-4 Laminates as a Function of Thickness and the Electrical Frequency of Measurement," paper IPC-TP-749, IPC Fall Meeting, Anaheim, CA, October 24–28, 1988.

[11] Faraji-Dana, R., and Y. L. Chow, "The Current Distribution and AC Resistance of a Microstrip Structure, "IEEE Journal of Microwave Theory and Techniques, Vol. 38, No. 9, September 1990, p. 1273.

[12] Horowitz, M., et al., "High-Speed Electrical Signaling: Overview and Limitations," IEEE Micro, Vol. 18, No. 1, January/February 1998, pp. 12–24.

[13] Widmer, A. X., and P. A. Franaszek, "A DC-Balanced, Partitioned-Block, 8B/10B Transmission Code," IBM J. Res. Development, Vol. 27, No. 5, September 1983, pp. 440–451.

[14] Burns, D., et al," Design Techniques for High-Speed Source Synchronous Busses," Design-Con2002, Santa Clara, CA January 28–31, 2002.

[15] Cohn, S. B., "Shielded Coupled-Strip Transmission Line," IRE Transactions on Microwave Theory and Techniques, October 1955, pp. 29–38.

[16] Connolly, J. B., "Cross Coupling in High Speed Digital Systems," IEEE Trans on Electronic Computers, EC-15, No. 3, June 1966, pp. 323–327.

[17] Voranantakul, S., et al., "Crosstalk Analysis for High-Speed Pulse Propagation in Lossy Electrical Interconnections," IEEE Trans. Components, Hybrids and Manufacturing Technology, Vol. 16, No. 1, February 1993, pp. 127–136.

[18] Walker, C. S., Capacitance, Inductance and Crosstalk Analysis, Norwood, MA: Artech House, 1990.

[19] Jorgenson, C., “Is FR4 Running Out of Gas?” Printed Circuit Design, September 2000, p. 10.

[20] Weiss, F. Kehrer, D. Scholtz, A.L. . Transmitter and receiver circuits for serial data transmission over lossy copper channels for 10Gb/s in 0.13 um CMOS, Radio Frequency Integrated Circuits (RFIC) Symposium, 2006 IEEE.

[21] Stephen H. Hall, Howard L. Heck. Advanced Signal Integrity for High-Speed Digital Designs, Wiley-IEEE Press, 2009.

[22] Ramesh Harjani, Design of High-Speed Communication Circuits, World Scientific Publishing Company, 2006.

[23] F. Krummenacher and N. Joehl, “A 4-Mhz CMOS Continuous-Time Filter with On-Chip Automatic Tuning,” IEEE J. Solid-State Circuits, vol. 23, pp. 750-758, June 1988.

[24] R. Payne et al., “A 6.25 Gb/s Binary Adaptive DFE with First Post-Cursor Tap Cancellation for Serial Backplane Communications,” in IEEE Int. Solid-State Circuits Conf. Feb. 2005.

[25] R. Gitlin and S. Weinstein, “Fractionally-spaced equalization: An improved digital transversal equalizer,” Bell syst. Tech. J., vol. 60, no. 2, pp. 275-296, Feb.1981.