EE 273
Digital Systems Engineering

Project Report

Prepared By:
Xuefeng Zhuang
Kai Yu
Ilya Katsnelson

March 2003

Table of Contents

List of Figures

List of Tables

1 Overall design approach

Our overall design strategy was to maximize the performance as much as possible in that short amount of time allocated for the project. As a result we choose the scheme for our system that would most likely work on the first iteration. We decided to use current mode, differential, bipolar signaling method due to the fact that is has very good noise immunity. The data transfer in only one direction has been used because in this case we didn't have to deal with the interference of the two signals on the same line as with the bidirectional signaling. Using bidirectional signaling would allow us to use smaller clock frequency to achieve the same data rate. However, it would also significantly increase the complexity of the receiver/transmitter circuits. For the signaling method we choose to use two-level signaling because this method can tolerate larger noise in the transmission than multi-level schemes. The overall block diagram of the router signal connection is shown on Figure 1.

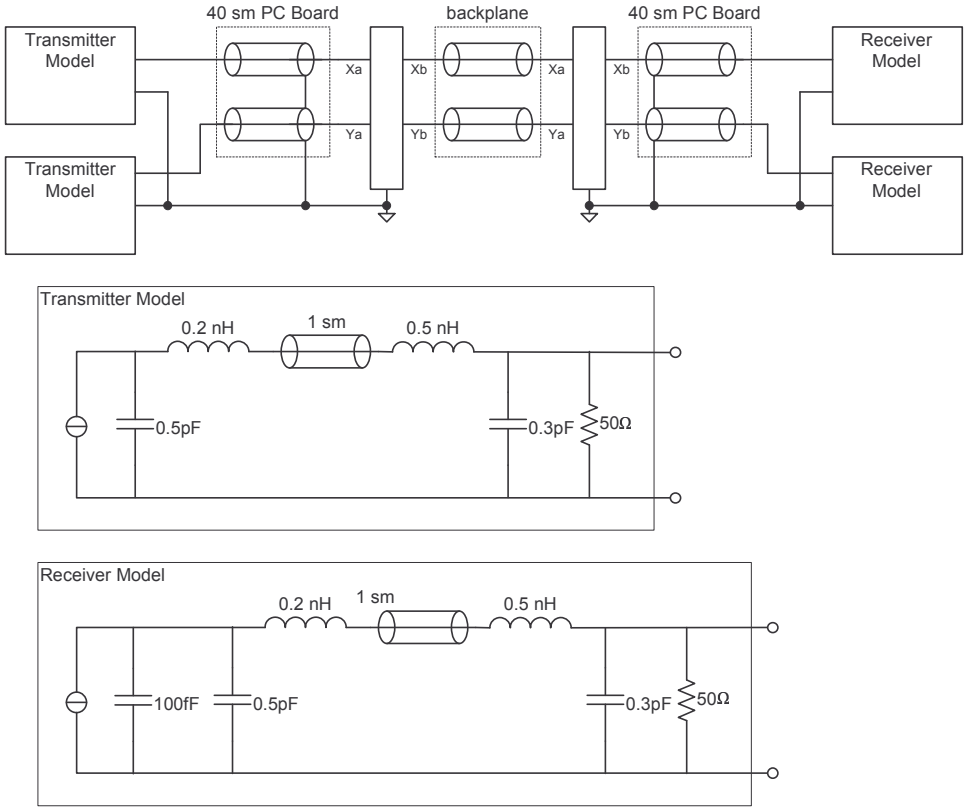

Figure 1. Block Diagram of the Router Signal Connection

Early in our design process we realized that sending the data directly would result in unreliable transmission due to the significant noise in the router backplane and the card connectors. In this case we decided to transmit compensated signals. Two compensation algorithms were used: transmission equalization and crosstalk compensation. The details of the algorithms are descried in parts 3.6 and 3.7.

In order to achieve the required 4Gb/s data rate we decided to use 2GHz clock but utilize both of its edges to sample the data signals. With smaller frequency many of the noise sources in the

system will be less pronounced. For the timing control we decided to use phase locked loop (PLL) clock generation with per-line closed loop timing. In this case the timing nose can be significantly reduced at the price of additional components, however. The details of the timing design are presented in part 4.

Table 1. Summary of Performance and Noise Results

Clock frequency	2GHz	Average Current, per line[1]	6.98 mA
Achieved data rate	4Gb/s	Average Power, per line	1.21 mW
Noise Margin	49 mv	Power Consumption, Switch card	619.52 mW
Timing Margin	96 ps	Power Consumption, Line card	77.44 mW
BER	1.39E-21	Total Power	2478.08 mW

Notes: 1 – The average current was measured from Hspice simulation run for a transmission of a random bit pattern.

Table 2. Summary of Cost Estimate

Crystal Oscillators	$90 (18, $5 each)	Switch card power	$12.38
Chip (flip-chip package)	$540 (18, $30 each)	Line card power	$1.54
Switch card chip	$65.6 (656 pins)[1]	Total power cost	$49.56
Line card chip	$13.6 (136 pins)[1]	Total switch card	$209.18
Switch card PC Board, 9 layers	$131.2 (16.4X40.0cm)	Total line card	$42.34
Line card PC Board, 8 layers	$27.2 (6.8X40.0cm)	Total router cost	$1725.8

Notes: 1 – We decided to use the same chip packaging to simplify the overall system design.

2 Description of 40 cm transmission lines on the line card and switch card

2.1 Board Specification

The specifications for the PCB boards were derived mainly from the constraints put on the parameters of the traces. The purpose of a board is to route signals from a connector on one side to a connector on the other. Therefore, the structure of the board traces is limited by how well it can route signals to and from a modern multi-pin connectors.

The design of the board trace routing to the chips is shown on Figure 2. We took the dimensions for the pads for BGA package from [1].

Our main consideration in developing the chip routing scheme was to maximize the distance between transmission lines of different channels. Due to the fact that we chose the differential signaling method, the separation between two signals of the same channel can be made very small without creating significant crosstalk. In this case we have decided to use 8 mils wide traces with 3 mils separation between traces of the same channel. Two lines of the same channels are connected to the two consecutive rows on the chip. In this case the separation between two adjacent channels can be made as wide as 22.6 mils. This value represents the common trace separation on the entire area of the PC board.

In order to achieve the above specified worst case separation between channels we have decided to connect every two rows of the chip (two rows specify connection to a single channel) to a different layer of the PC board. Putting ground or power layer between signal layers on the PC board will create complete isolation between different channels. The diagram of the board layer layout is shown on Figure 3. The board utilizes 9 layers with 4 signal layers, 2 power and 3 ground between them. The electrical parameters of the board are represented in Table 3. Despite that fact that the diagram on Figure 2 shows that the distance between traces is smaller then specified 22.6 mils when a trace leaves the chip, we didn't take this into account because the length of a trace affected by this effect is very small in comparison to the entire length of the trace on the board.

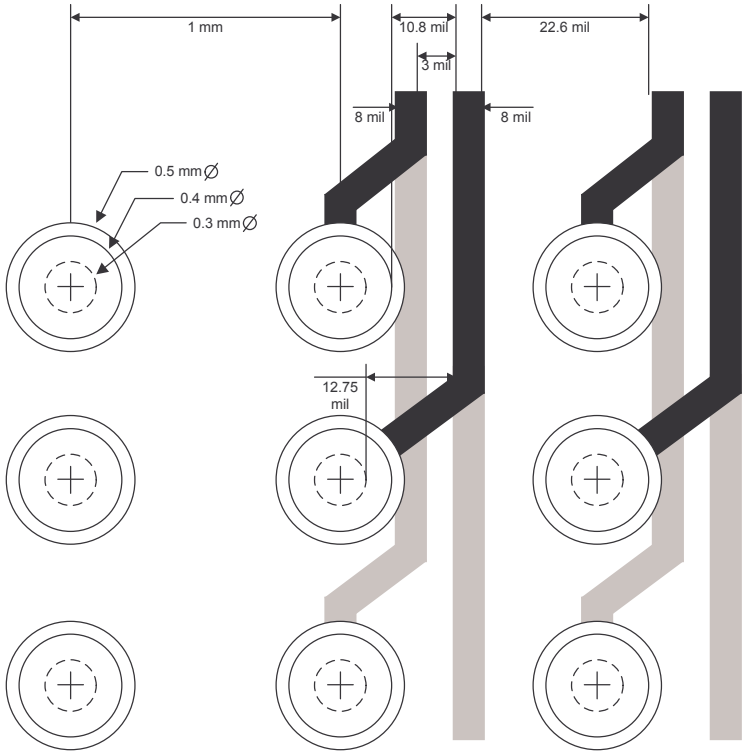

Figure 2. Chip Connection Layout

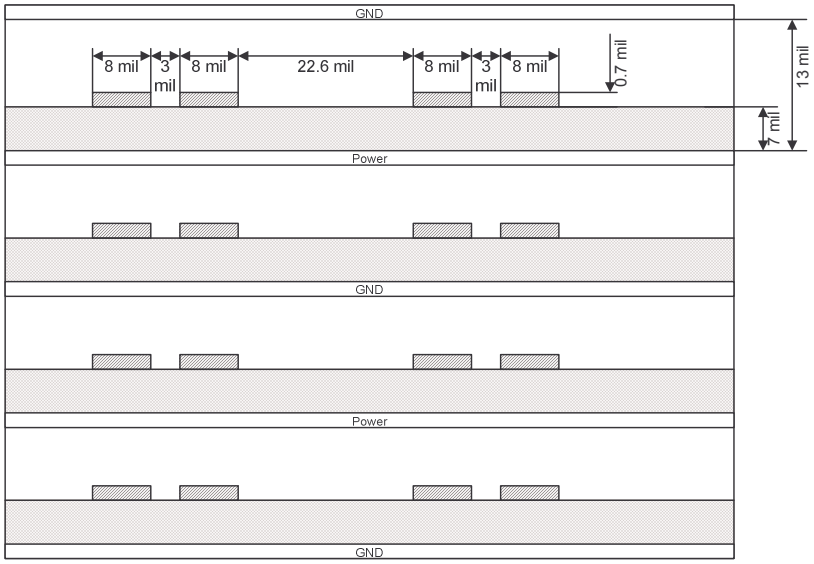

Figure 3. PC Board Layout

Each PC board is connected to the router back plane with Teradyne VHDM connector. These connectors allow for low noise high density signal interconnect. For our design we decided to use 6 row Teradyne VHDM-HSD configuration [2]. The routing diagram is shown on Figure 4. With this scheme we were able to use each two columns of the connector to connect to 3 differential channels with maximum of 128 columns required to connect the switch card. Each signal channel of a column connects directly to a separate layer of the PC board.

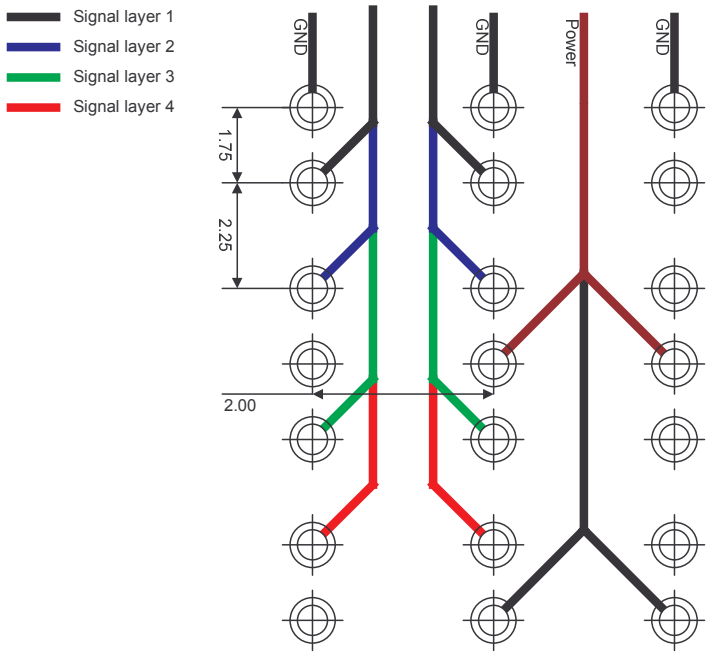

Figure 4. Teradyne VHDL-HSD Connector Routing

2.2 Electrical model of the line

Electrical parameters of the PC board are shown in the table below. In order to calculate the parameters we have used LINPAR for Windows 1.0 software package.

Table 3. PC Board Electrical Parameters

Base Parameters			
Strip Width Strip Thickness Separation Width Substrate Thickness	8.0 mil 0.7 mil 22.6 mil 7.0 mil	Cover Height Permittivity (ε_{ru}) Loss Tangent (tanδ)	13.0 mil 2.55 0.005
Electrical Characteristics			
Frequency	2.000E+09 Hz	Conductivity	5.600E+07 S/m
Inductance Matrix [L] (H/m)		Capacitance Matrix [C] (F/m)	
2.656E-07 3.592E-10	0 2.656E-07	1.068E-10 -1.442E-13	0 1.068E-10
Skin Effect Matrix [R_S] (Ohm)		Characteristic impedance	50.2
0.00085 0	0 0.00085		

For our model we didn't take into account R_0 and G_0 parameters. This is done in order to separate AC and DC analysis of the system. Without R_0 and G_0 we can assume the board to be a lossless line and concentrate more on the frequency dependent effects.

3 Signaling Analysis

3.1 Unidirectional vs. Bidirectional signaling

We decided to use unidirectional signaling because we would not have to deal with the increased noise due to bidirectional signaling. The increase in noise is due to the reverse crosstalk, which is not present in unidirectional signaling. We decided to first connect a unidirectional circuit, and if the noise margin were large enough, we would try to implement bidirectional signaling. However, after hours of tweaking we had only achieved a 100mV eye with unidirectional signaling, so we focused all of our energy on making the unidirectional signaling work. While a bidirectional system is more elegant, we could not handle the increased noise.

3.2 Single-ended vs. Differential signaling

As with most decisions, the option that reduced noise the most was chosen, and differential signaling had better noise performance. Differential signaling cancels common mode noise, and also doubles the signal swing, which would help our noise margin. Single-ended would require fewer pins, and also operate at a higher frequency, but would not help our noise margin at all. Differential signaling also serves as its own reference, and would make the receiver circuitry simpler.

3.3 Number of Signal Levels

We decided to stick with 2-level signaling mainly for the same reason we chose unidirectional signaling: lack of noise margin. More levels would have let us operate at a smaller frequency, but even at 2Gb/s there was a small noise margin, so adding more levels would have eliminated the existing noise margin. It would also require more power, and it would be nice to have a reasonably priced board. It turned out part of our original problem was due to the skin effect, so a decreased frequency would have helped. However, after fixing it, it turned out almost all the noise was due to crosstalk, so operating at 2-levels was chosen.

3.4 Signaling Rate

Increasing the clock frequency from 1GHz to 2GHz doubled the transmission rate. We sampled on both clock edges, so with a period of 500ps, we had a bit cell width of 250ps. Other signaling methods would have meant an increase in frequency, which would have led to increased noise.

3.5 Signaling Convention

Bipolar differential current-mode was chosen because it cancelled the most noise. Using a bipolar scheme means the noise margin for zeros and ones are the same, while in unipolar mode the threshold would need to be decreased to balance the noise margins. We chose to encode a “1” as +20 mA and a “-1” as -20 mA. Note that after equalization and compensation, the peak current increased so we just scaled the total current down so that the maximum current didn’t exceed 24 mA. We also decided to use NRZ encoding to maintain as small of a frequency as possible while still maintaining 4GB/s. Using any return-to-zero encoding would have required a much larger frequency. The reason for choosing differential signaling is answered in section 3.2. Current mode was chosen because it has a large input impedance, which means greater isolation from any power supply noise.

3.6 Equalization

This turned out to be one of the key features of our design. Due to slow rise and fall times, we needed an FIR filter to make sure the transitions occurred without the sampling frequency. Ideally, the impulse response convolved with the FIR filter would generate a pulse with a fast rise and fall time (see picture below).

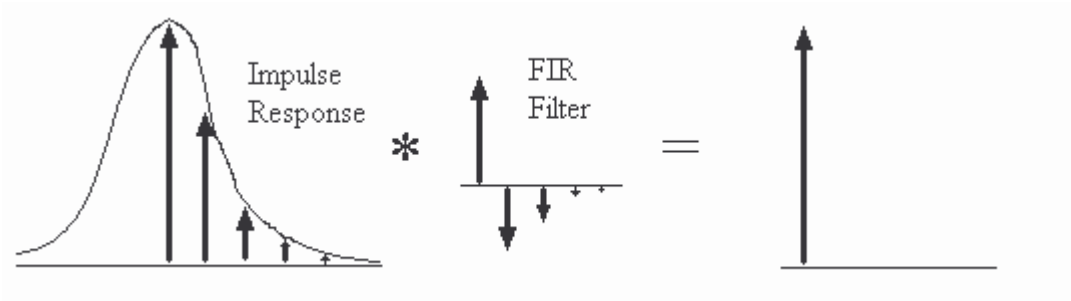

Figure 5. General Equalization Method

To determine the tap values, we sent an impulse response through the line, and used the maximum value of the impulse response as the time of the first tap. We then took the amplitudes spaced 250ps apart, and used the attenuations to determine the tap values.

3.7 Crosstalk Control

The major proportional noise source in the router is crosstalk. Since we have homogenous medium, the forward crosstalk coefficient is zero. There will be no signal return crosstalk due to the differential signaling scheme that we used. Furthermore, we placed the neighboring wires far apart (22 mils) to minimize crosstalk. We simulated impulse responses for the backplane, the connector, the chip and the PCB trace and came into realization that the connector is the single most important contributor for crosstalk. The peak to peak value of crosstalk on the victim line is about the same as the magnitude of the impulse response on the aggressor. Therefore, crosstalk cancellation is compulsory.

The cross impulse response exhibits a peak as well as a dip, so the filter must be operating at a frequency at least double of the aggressor line frequency, and it must start earlier than the aggressor bit. Since the crosstalk is caused by transition edges on the aggressor, we analyzed the cross step response on the victim line. Tap values were first calculated in MatLab then hand tuned to minimize the effect of the aggressor transition edge (after equalization) on the victim line. Note that this method does not consider the second order effect that is the transition edge will be different after itself is adjusted for crosstalk and hence will induce a slightly different crosstalk pattern on the victim and so on. Fortunately these second order effects are sufficiently small and our simulation showed that neglecting them will still result in a robust system. Our final crosstalk cancellation scheme was an eight tap filter with 62.5 psec resolution and 0 hold time, as illustrated on Figure 6.

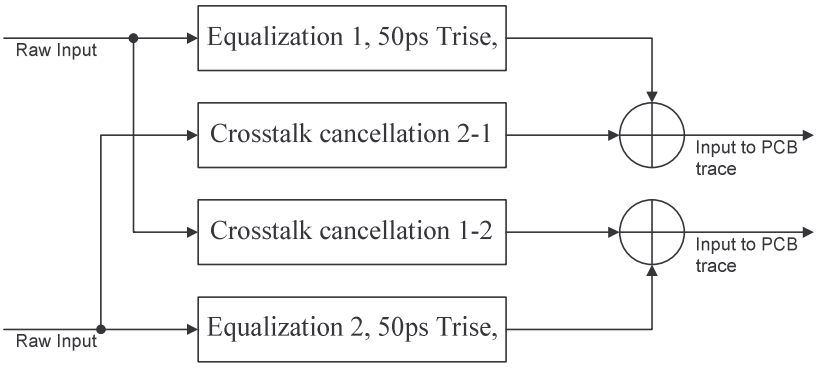

Figure 6. Equalization and crosstalk cancellation

The worst case equalized crosstalk impulse response after compensation is shown on the figure below:

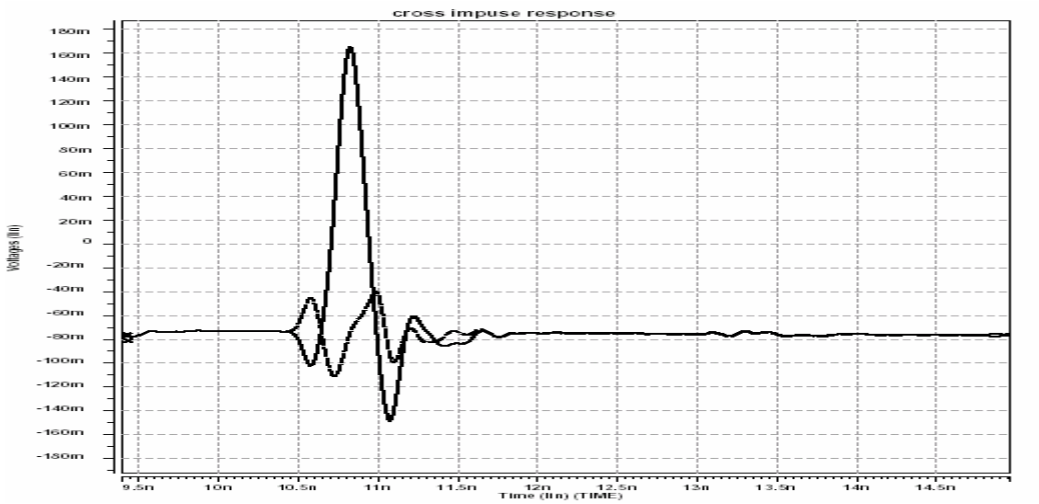

Figure 7. Worst case equalized crosstalk impulse response after compensation

3.8 Signaling details

We used two separate current source drivers per channel. One of them has a rise time of 50 psec, and is used to transmit bits without considering crosstalk; the other current driver is dedicated for crosstalk cancellation and has a rise time of 62.5 psec. Because the current drivers are integrated in the transmitter chip the additional current source basically comes for free. By considering them as two separate sources simplifies our crosstalk cancellation algorithm tremendously.

Since the technology we are using is 1.2V, the voltage swing on any one line can not exceed 1.2V. Since we have a termination resistance of 50 Ohm in parallel with a 50 Ohm transmission line, the maximum current swing can not be larger than 48mA. In our system, the overall worst case peak current is ±24mA, and the average current is well below the 20mA mark.

3.9 Details of your driver, receiver, and termination circuits

Because of our tight noise budget, the 40mv receiver offset is formidably excessive. Lee et al show a method to reduce such offset values [5]. The method is to put binary weighted PMOS capacitors at the integrating nodes of the differential sense amplifier, as illustrated in Figure 8. Lee et al demonstrate that the offset cancellation network will reduce a ±120mv offset to 8mv. Accordingly, the ±40mv offset can be reduced to 3mv (assuming direct proportionality). To be conservative, we took the offset value of 5mv after cancellation. We will not elaborate on the details of the scheme due to page limitations.

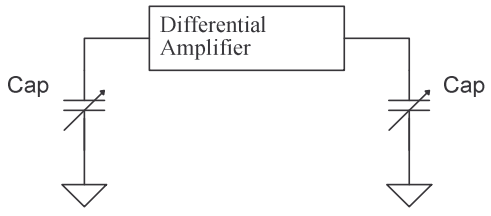

Figure 8. Receiver offset cancellation

Both ends of the system are terminated to minimize reflections. The programmable termination resistor is used to minimize dissemination and hence ISI. Dally and Poulton explain the operation of such devices [4]. In this scheme, an external reference resistor equal to the characteristic impedance of the transmission line is used to set the resistance of the internal terminators. Both external reference and internal replica terminator share a common node to voltage supply. Identical current sources are connected

to the two resistors and voltage difference across the two blocks is measured by a clocked comparator and the counter value will be increased or decreased accordingly. This creates a feedback loop that adjusts the termination resistance within 5% of the nominal value. A block diagram of the controller is shown on Figure 9

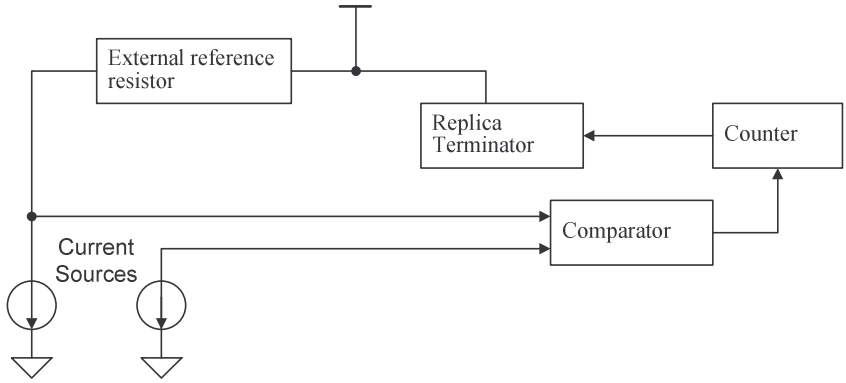

Figure 9. Automatic terminator resistance adjustment

4 Timing

4.1 Overall Timing Concept

To reduce noise as much as possible, we decided to use PLL clock generation with per-line closed loop timing for our timing system. Generating a clock signal of 2GHz with the oscillators (maximum frequency = 300MHz) given required some frequency multiplication method, and we decided to use on-chip PLL clock multiplication, implemented with a charge pump filter with an integrating loop filter (4 figure 12-81, page 624). We later found VCO jitter was the greatest source of timing noise, so to decrease it we ran it at 4GHz and divided the clock down to 2GHz with a counter, ultimately cutting the jitter by 2. Two clock signals with a phase difference of 180^{o} would be multiplexed and sent to the transmit flip-flop and through the backplane to the receiver. For per-line closed loop timing we needed multiple phases of the clock, so we added a differential ring oscillator to the VCO to generate four phases. We actually only need to send 2 signals with a 90 degrees phase difference because the flip-flops trigger on both edges. Per-line closed loop timing cancelled out most types of skew and some jitter.

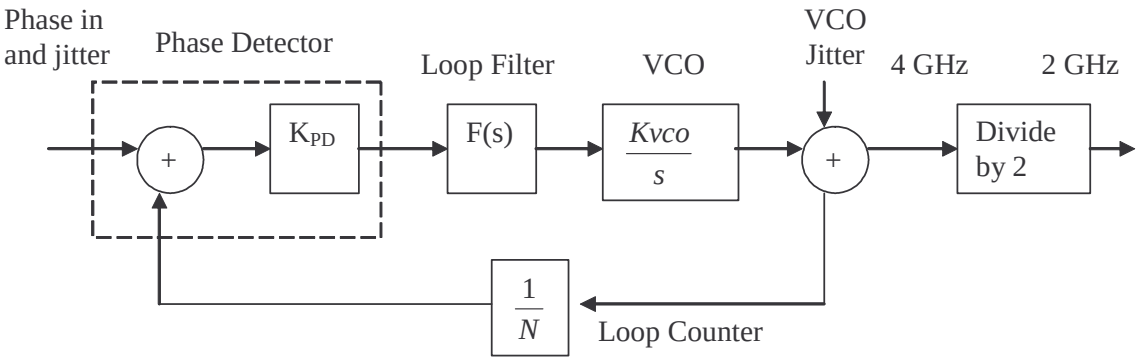

Figure 10. Phase Lock Loop Continuous Time Model

4.2 Number and Location of Independent Oscillators

We have 18 independent oscillators, one used for every transmitter and receiver circuit per line card. The oscillators will be off chip, but the PLL will be on chip. One of the clock phases is sent to the transmitter, and 4 to the receiver. We could have used fewer oscillators and distributed the clock signals, but to create a robust design we decided to have more oscillators at a minimal cost.

4.3 Description of Control Loops

We used per-line closed loop signaling to eliminate almost all skew and some jitter. Another method to help clock recovery would have been oversampling, but that would require running the clock at a much larger frequency than 2GHz. Although we were generating a 4GHz signal already, to oversample we would have needed at least 6GHz to sample each bit multiple times, but the VCO can only center up to 5GHz, so it was not an option. Thus, we just used the per-line closed loop timing as in our notes (HO #20, panel 22). Here is our implementation:

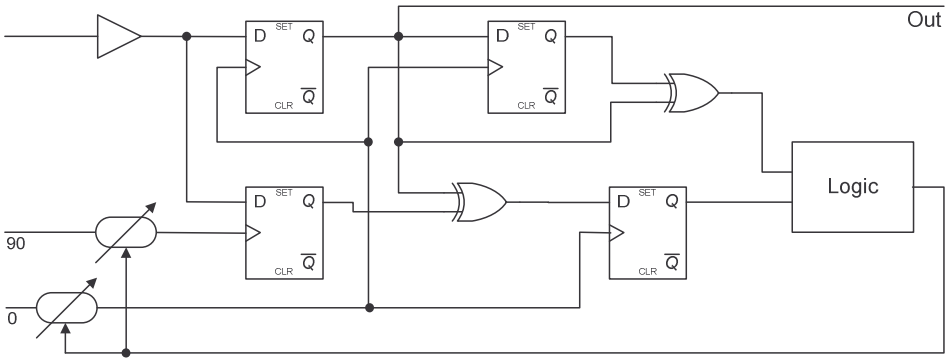

Figure 11. Per-Line Closed-Loop Timing

Our diagram is slightly different because we have generated our own phase-shifted signal. In the blocks given there was no phase-shifter, so we generated our own. We have already generated the phased clock signals, so we can cut down slightly on the timing jitter. The logic will just adjust the delays such that the middle of the eye is continuously sampled. It will take a total of 4 gate delays to get the proper output (2 XOR and 1 AND gate for detection, 1 gate to send proper signal to the adjust the sampling time).

4.4 Timing Circuit Details

The PLL for clock multiplication uses a charge pump and integrating filter. The loop frequency should be set to the highest value (rule of thumb is $\leq .1\omega_{ref}$ from pg. 625), or 25 MHz, though this is not the optimum value. From the equations (12-22, page 625) with a transfer function of F(s) = 1 + 1/sRC (integrating) gives us:

$$\omega_N = \sqrt{\frac{I_P K_{VCO}}{2\pi R_2 C N}}\ ;\ \xi = \omega_N \frac{R_2 C}{2}$$

As ω_N increases, I_p increases, which means greater power dissipation. Thus, to conserve power (we are already sacrificing power for noise margin, so every little bit helps), we decided to set ω_N as low as possible, or 2π*1MHz. The current sources

available cannot drive large amounts of current either, so if the required current is too large, the PLL cannot be implemented on-chip. Assuming VCO has a tuning gain of 1V and 50% variation, and the output is centered around 4GHz, $K_{VCO} = 2\pi*(6GHz-2GHz)/1V = 8\pi$ rad GHz/V. The damping ratio is set to 1 to optimize noise and lock acquisition (pg. 623), so $R_2C = 3.18e-7$. We decided to set R = 100kΩ, and C = 3.183 pF, making I_p = 50.2 mA. If we had used 25MHz for the loop frequency, I_p would have equaled 1.26 A, 25 times the amount of current needed to achieve a damping ratio of 1 for a loop frequency of 1MHz. It would have taken too many on-chip current sources to generate this 1.26A, so decreasing the loop frequency was a good decision. The PLL will even low-pass filter some of the jitter and noise because of the loop filter, but that cannot be measured.

Table 4. Summary of PLL parameters

ω_N	2π*1MHz	ξ	1
N	16	I_p	50.2 mA
VCO range	2-6 GHz	R_2	100kΩ
K_{VCO}	8π rad GHz/V	C	3.183 pF

For closed loop, a multiphase clock signal is needed. By using a differential ring oscillator, we are able to produce the 4 phases needed. The ring oscillator will be connected to the output of the PLL, and a diagram is shown below [3]:

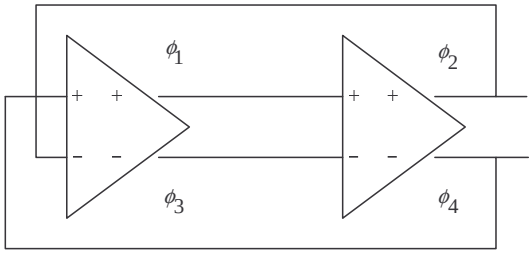

Figure 12. Differential Ring Oscillator

This way, the first driver will transmit when ϕ_4 is high, and the second driver will transmit when ϕ_2 is high (the drivers alternate). An LC based oscillator can also be used, which would have less jitter because it has a high Q. However, it was not chosen due to lack a familiarity, and it also requires transistor level design.

5 Signaling Convention, Noise Budget, BER

5.1 Transmitter Offset

The current source at the transmitter has a ±5% worst case offset. This has been taken into consideration in the eye diagram simulation.

5.2 Receiver Offset and Sensitivity

As mentioned earlier in 3.9, the receiver offset can be reduced to ±5mv. This plus the receiver sensitivity of 10mv makes the total fixed noise in the router to be 15mv.

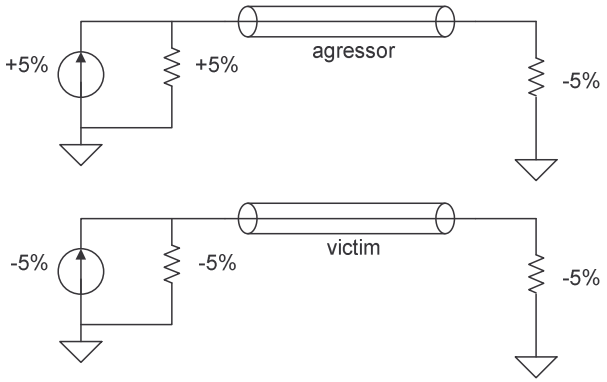

Figure 13. Worst case setup

5.3 Crosstalk

After crosstalk compensation, the residue crosstalk on the victim line is about 25%. This value is obtained by measuring the impulse response of the victim line after crosstalk cancellation.

5.4 Inter-symbol Inference

Our unidirectional signaling implementation eliminates the adverse effects of reverse crosstalk. The worst case ISI occurs when we have the setup illustrated on Figure 13.

If we consider only the first reflection (since the attenuation is large at frequency of 2 GHz), the worst case ISI (big current drive, big transmitter termination, small receiver termination) is (1.05)*(1.05//1)*[(1-0.95)/(1+0.95)]/0.5=2.75%. Due to the fact that there are mismatches on the backplane-connector connection and the connector-PCB trace connection, the ISI will be larger than this value. An evaluation of the TDR response of the whole transmission line reveals that the mismatches are less than 5%. Based on these facts, a very conservative estimate of the total ISI will be 4%.

5.5 Gaussian Noises and Other Noises

5mv Gaussian noise is from the backplane. The termination resistor thermal noise is V_n= (4KTR(delB))^.5=(1.3807E-23*300*50*0.5*4E9)^.5=0.02mv. We can ignore this small value for all practical purposes. Perpendicular crosstalk is eliminated since we shield different layers with power and ground planes.

5.6 Power Supply Noise

Since ideal current sources have infinite impedance, we can safely ignore the power supply noise in our analysis.

5.7 BER Calculation

In calculating the bit error rate, we use the worst case setup and the following noises are accounted in the eye diagram: transmitter offset, crosstalk, inter-symbol interference and receiver offset. The thermal noise, power supply noise and perpendicular crosstalk can be ignored since they all really small. We still need to account for the receiver offset and sensitivity of 15mv. The calculated values of various noise sources are summarized in Table 1. Table 2 shows the final BER calculation. The height of the eye opening (i.e., the height of the rectangle drawn on the eye, which is NOT the maximum height) is 64mv. The bit error rate obtained is 1.39E-21. Note that the sampling window is 96 ps, which is more than sufficient to account for all the timing uncertainties.

Table 5. Noise Sources

V_S	Signal swing (after attenuation)	425	mv	Included in simulation?
V_{NI}	Receiver offset and sensitivity	15	mv	no
	Power supply noise	0	mv	no
	Perpendicular crosstalk	0	mv	no
K_N	Crosstalk	25	%	yes
	Transmitter offset	5	%	yes
	Inter-symbol interference	4	%	yes

Table 6. BER Calculation

	Calculated	Simulated	
Gross margin	212.5	212.5	mv
Eye Opening	68	64	mv
Net Margin	53	49	mv
Total Gaussian noise	5	5	mv
Timing margin	53.75	96 (measured eye opening)	ps
VSNR	11.6	9.8	
BER (measured)	6.03E-30	1.39E-21	

6 Timing Analysis

6.1 Bounded Sources of Jitter

The bounded sources of jitter are:

- On-chip VCO (10% of period) 25ps
- Data-Dependent Jitter 25ps
- Phase Comparator (PLL) 20ps
- Delay Line (Receiver End) 6.25ps
- Transmit flip-flop 50ps

6.2 Statistical Sources of Jitter

- Clock Jitter 20ps

6.3 Cancelled and Uncancelled Skew

Due to closed loop timing, we cancelled virtually all the skew. We cancelled the skew: between clock and data line, fixed differences in transmitter and receiver circuits, transmit clock between flip-flops, aperture offset in receive flip-flop, and the offset in the 90 degree delay line.

6.4 Conclusion

The bit cell width is 250ps, and to find the gross margin we need to subtract the aperture (20ps) and rise time (20-80% of 50ps = 30ps). This, our gross margin is 200ps. We ignored the jitter in the receive flip-flop because that does not affect our receive circuitry, it affects the circuits that use the data and we will let them deal with it. Thus, our net margin is the uncancelled jitter subtracted from our gross margin, or **53.75ps**. Actually, the low-frequency jitter of the transmitter will be slightly filtered by the loop filter, but that is not included.

7 SPICE Simulations

7.1 Eye Diagram with 500 random bit (with lone one and lone zero)

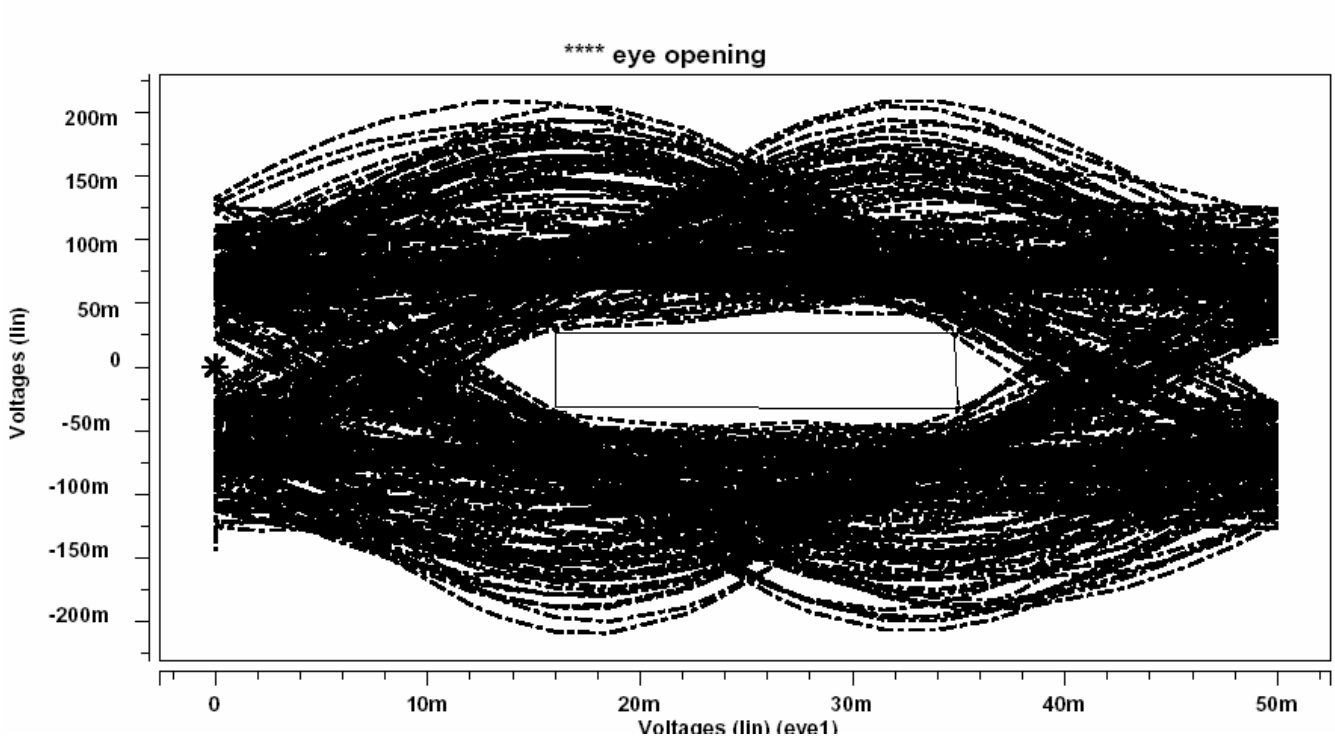

Figure 14. Eye diagram

7.2 TDR and TDT simulations of your transmission line including parasitics

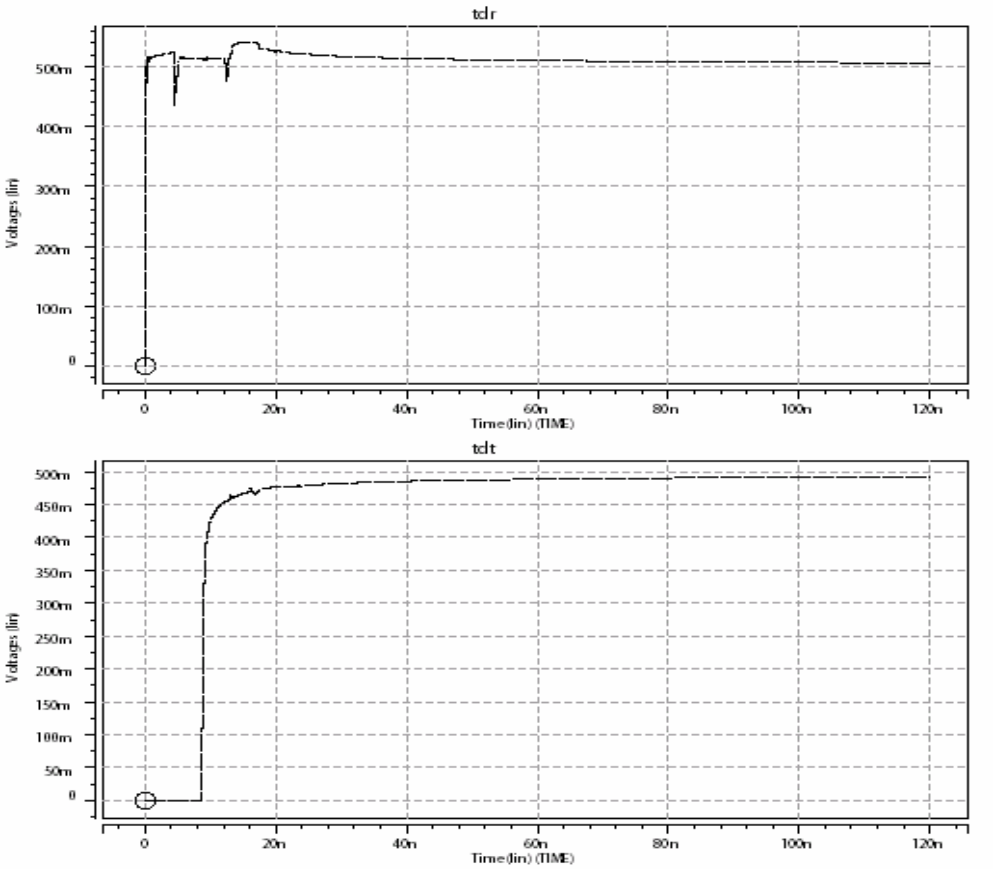

Figure 15. TDR and TDT Response of the System to a 50ps step

References:

1. Board Routability Guidelines with Xilinx Fine-Pitch BGA Packages. Abhay Maheshwari and Soon-Shin Chee. XAPP157 (v1.2) November 14, 2002
2. Teradyne Connection Systems, Inc. Power Module Backplane, VHDM. C-437-5050-500.
3. High-Speed Serial Transceivers for Data Communication Systems. John M. Khoury and Kadaba R. Lakshmikumar. Multilink Technology Corp. http://web.doe.carleton.ca/courses/97578/topic7&8/Multilink_SerialTransceiver.pdf. March 5, 2003.
4. Digital systems engineering / William J. Dally, John W. Poulton. Cambridge, U.K. : New York, NY, USA : Cambridge University Press, 1998.
5. M. E. Lee, W. J. Dally, P. Chiang, Low-Power Area-efficient High-speed IO Circuit Techniques, IEEE Journal of Solid State Circuits, Vol. 35, No. 11, Nov 2000.