IMPEDANCE MATCHING TECHNIQUES FOR ETHERNET COMMUNICATION SYSTEMS

A Thesis

by

RICHARD ALAN KAMPRATH

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

May 2007

Major Subject: Electrical Engineering

IMPEDANCE MATCHING TECHNIQUES FOR ETHERNET COMMUNICATION SYSTEMS

A Thesis

by

RICHARD ALAN KAMPRATH

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

Approved by:

Chair of Committee,	Jose Silva-Martinez
Committee Members,	Edgar Sanchez-Sinencio
	Jay Porter
	Jun Zou
Head of Department,	Costas Georghiades

May 2007

Major Subject: Electrical Engineering

ABSTRACT

Impedance Matching Techniques for Ethernet Communication Systems. (May 2007)

Richard Alan Kamprath, B.S., Texas A&M University

Chair of Advisory Committee: Dr. Jose Silva-Martinez

In modern local area networks, the communication signals sent from one computer to another across the lines of transmission are degraded because of reflection at the receiver. This reflection can be characterized through the impedances of the transmitter and the receiver, and is defined by the Institute of Electrical and Electronic Engineers (IEEE) as the S11 return loss. The specifications for S11 return loss in Gigabit Ethernet are given in terms of magnitude only in the IEEE 802.3 guidelines. This does not fully take into account, however, the effects of frequency dependant impedances within the bandwidth of interest. With a range of 30% error in the category 5, or CAT5, transmission line impedance used in this specification and no further requirements for individual components within the Gigabit Ethernet port, such as the RJ45 magjack or the physical layer, the system can easily be out of tolerance for return loss error. A simple impedance matching circuit could match the CAT5 cable to the physical layer such that the return loss is minimized and the S21 transmission is maximized.

The first part of the project was commissioned by Dell Computer to characterize the return loss of all of its platforms. This thesis goes further in the creation of a system that can balance these two impedances so that the IEEE specification failure rate is reduced with the lowest implementation cost, size, power and complexity. The return loss data were used in the second phase of the project as the basis for component ranges needed to balance the impedance seen at the front of the physical layer to the CAT5

transmission line. Using the ladder network theory, an impedance matching circuit was created that significantly reduced the S11 return loss in the passband of the equivalent ladder network. To manage this iterative process, a control loop was also designed. While this system does not produce the accuracy that a programmable finite impulse response (FIR) filter could, it does improve performance with relatively minimal cost, power, area and complexity.

DEDICATION

To my parents and my sister.

ACKNOWLEDGEMENTS

I began the journey that culminated in this thesis project six years ago as an undergraduate at Texas A&M and finish it as a masters student with the Analog and Mixed Signal Center. In that time, there has been no bigger influence on my development than Dr. Silva-Martinez. He has taught a handful of my classes, and more recently served as my academic advisor. Without his guidance and support I would not be where I am today. I am very thankful for everything he has done for me.

I would also like to thank the other professors I have worked with in our group. Dr. Sanchez-Sinecio has also been instrumental in my technical development, and I have learned a great deal about project planning and time management during the classes I have taken from him. I would also like to thank Dr. Karsilayan and Dr. Cilingiroglu for their contributions to my education through their classes and technical guidance. They both have the rare ability to portray very technical ideas in an easy to understand manner. I only hope that I emulated their abilities while teaching my own classes.

While my project has been a personal journey, it could not have happened without the support of my friends. Jake Tappan, Dr. Aaron Spiegel, Will Morrison, Reid Vidrine, Rory Dyck, Rob Rohlfs, Mike Hopkins, Steve Wake and Frank Thrash have been there for me for the entire time I have been in college, and they have encouraged me through all the challenges I have faced.

I would also like to thank all of those in the Analog and Mixed Signal Group for their continued support. Alberto Valdes-Garcia and Chinmaya Mishra have assisted me with numerous questions and problems, even when they were busy with their own studies.

I am eternally grateful for their help. Alfredo, Xiao, Lin and Tom have been good friends and have been a great part of my graduate experience. It would not have been the same without them. I would also like to thank Raghavendra Kulkarni. We have worked on many things together, and his insight has been invaluable to me. I have really enjoyed getting to know him and we have had too many great conversations to recount. He is a wonderful person, and I am honored to count him amongst my friends. Jason Wardlaw, Johnny Lee and Felix Fernandez deserve special recognition. They have been by my side during all of my battles in graduate school, more often than not in the trenches with me. Together we have overcome countless sleepless nights and very long weeks that turned into long months and semesters. I have redefined my definition of work ethic and I thank them for pushing me to excel. I can only hope that I work with such great people in the future.

There are some industry people I would like to thank as well. Bryan Kelly and Jon Lewis of Dell Computer have been helpful to me as supervisors and colleagues throughout my internship with Dell and beyond. Jon has become a wonderful friend, and they both have encouraged me to accomplish my goals, whatever they may be. I would also like to thank the Tyco, Pulse and Belfuse vendors that I have worked with during this project for their support and product help.

Lastly, I would like to thank my parents for their never ending love and support. They have continued to push me to do my best, and without them I would not be who I am today. I owe everything that I have accomplished to them. My sister has been there for me as well, and I am very thankful that she has been a part of my experience at A&M.

TABLE OF CONTENTS

Page

CHAPTER Page

LIST OF FIGURES

FIGURE Page

FIGURE Page

FIGURE Page

FIGURE Page

LIST OF TABLES

TABLE Page

CHAPTER I

INTRODUCTION

Large computer manufacturers spend millions of dollars each year on IEEE compliance testing and fault definition. One area of this testing involves the Ethernet port found on all systems produced such as laptops, servers and PCs. System failures are found only during post-production testing, and this step is costly in terms of time and money. Other parts fail in the field, which causes headaches for consumers. The first part of this project was initiated at the request of Dell Computer to address their Gigabit Ethernet interface yield. The circuit design and proof of concept were required to show the feasibility of a low cost, easy to implement on board solution for reduced return loss.

In the IEEE 802.3 specification for Gigabit Ethernet, signal return loss created by impedance mismatch between the medium dependant interface (MDI) and the transmission line should be minimized as much as possible. Without an on chip/board solution, this minimization can only be achieved by optimizing the differential impedance of the physical layer (PHY), to match the category 5 transmission line impedance (CAT5), before fabrication. However, with process variations, mismatch and the Ethernet cable impedance tolerances allowed for in the specification, this specification cannot be met across all corners. This translates into lower yield and higher costs for the manufacturer and ultimately the consumer.

This thesis follows the style of *IEEE Transactions on Microwave Theory and Techniques.*

Each product line, such as servers, PCs and network interface cards, must order fabricated magnetic coupling transformers to match each individual product lines, which again increases cost and lowers yield.

For high volume manufacturers, an on chip/board solution for PHY impedance matching could save hundreds of thousands of dollars over the lifetime of a single product line and millions of dollars in all [1]. Thus, the first step in this project was the characterization of the return loss across all product lines, and then a test suite was created so that engineers could test their Ethernet boards for compatibility under the worst case corner conditions. From that point on, all work has concentrated on an active control loop for the matching of the PHY and line impedance based on the data collected at Dell. This matching circuit would be run once at system turn on, and the resulting solution would be used until shut off. The characterization data and IEEE compliant I/O Silicon structure data are Dell Computer proprietary information. While the on chip solution presented is based on the collected data, this design can be modified to match any line impedance. By creating a circuit that can match the PHY to the line impedance, one set of coupling magnetics can be used across all product lines in the case of Gigabit Ethernet. The result is twofold; the cost of magnetics will go down due to the high volume required of a single magnetic and the product yield will go up because a wider range of PHY impedances can be matched.

The biggest constraint on the implemented solution should be cost. This means that the simplest circuit solution should be realized in the most inexpensive technology. One recent solution using a finite impulse response (FIR) filter has been proposed [2]. However, the complexity of this circuit prohibits its use in cost constrained situations,

such as the Gigabit Ethernet problem. Thus, a control loop using an active impedance matching device is implemented with an on chip solution or discrete passive components capable of being put directly on the motherboard. The voltage across the PHY is compared to a reference while components are swept, and this information is digitally used to control the variable resistance and capacitance banks. While this proposed solution cannot produce the same resolution as an FIR approach, it can be implemented with reduced cost, area and power consumption.

The ladder network matching implementation also has other major applications, such as bond wire impedance matching. Included in the Appendix is more information on this issue and some Cadence plots to show this bondwire solution. The technique presented is shown to reduce return loss in RF bondwire applications and increase the frequency bandwidth of operation.

The goal of this project was to create a proof of concept circuit and control loop to reduce the return loss at the Gigabit Ethernet port and provide theoretical and experimental results that could later be used to create a mass production solution. This goal becomes all the more important with 10 Gbps and higher systems due out in the near future. The next chapter, Chapter II, goes through the process used to collect and characterize the return loss data. New empirical specifications for the system and the magnetics are created with the return loss data. This data is then used in Chapter III to create a circuit model of the PHY. From this, a network is then constructed to match the created model to the CAT5 transmission line. Chapter IV shows the Cadence simulation plots of the PHY model and the ladder network created. These plots verify that the proposed solution does indeed reduce the S11 return loss seen at the input of the PHY.

Chapter V shows the created experimental model results of the PHY and ladder network. The S11 return loss is lowered with the solution in place, confirming the simulation results. At the same time, it is seen that the bandwidth is increased, the jitter is reduced and the eye diagram opens. Chapter VI gives a brief conclusion of the work done in this project and the results obtained.

CHAPTER II

ETHERNET RETURN LOSS CHARACTERIZATION

For local area networks (LANs), the preferred mode of communication is Ethernet due to its high data rate and reliable transmissions. The following figure 1 shows the common components of the system: the category 5 (CAT5) transmission cable and the computer interface.

Figure 1: Ethernet Communication Port and CAT5 Cable

The CAT5 cable consists of four twisted pairs of wire transmitting and receiving at a rate of up to 125 MHz, giving a total bandwidth of 1 Gbps. Included in the Ethernet port is the magnetic jack, the physical layer (PHY) and the termination load. The port structure is illustrated in the following figure 2.

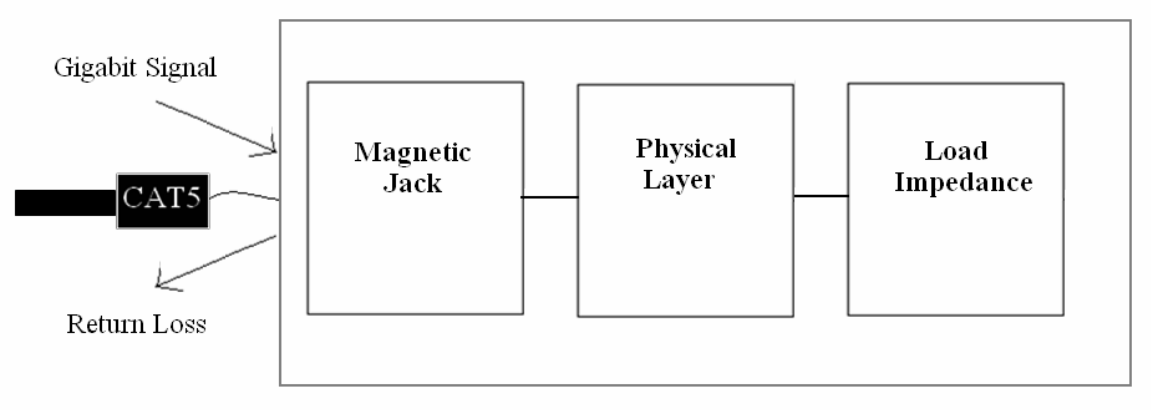

Figure 2: Ethernet Port Structure

The magnetic jack is a coupling transformer used to electrically isolate the CAT5 transmission line from the Ethernet port. It is mainly a reactive impedance, and it is relatively small compared to the PHY and the load. The physical layer is the trace on the motherboard that connects the magnetic jack to the load impedance which terminates at the receiver input. As seen in the above figure 2, the return loss occurs at the input of the Ethernet port.

In 1000 Base-T Gigabit Ethernet, which is the specification for this type of communication, the return loss of the data signal is bounded by the IEEE 802.3.40.7.2.3 Medium Dependant Interface (MDI) Return Loss Specification. This is stated as follows:

> *The differential impedance at the MDI for each transmit/receive channel shall be such that any reflection due to differential signals incident upon the MDI from a balanced cabling having an impedance of 100 Ohms ±15% is attenuated, relative to the signal, at least 16 dB over the frequency range of 1.0 MHz to 40 MHz and at least 10-20log(f/80) dB over the frequency range 40 MHz to 100 MHz (f in MHz). This return loss shall be maintained at all times when the PHY is transmitting data or control symbols.*
>
> [3]

The following table 1 is created from the above MDI return loss definition.

Table 1: IEEE 802.3 Gigabit Ethernet Return Loss Magnitude Specification

Frequency (MHz)	S11 Specification (Magnitude)	S11 dB
1	0.158489319	-16
10	0.158489319	-16
20	0.158489319	-16
30	0.158489319	-16
40	0.158489319	-16
50	0.197642354	-14.0824
60	0.237170825	-12.498775
70	0.276699295	-11.159839
80	0.316227766	-10
90	0.355756237	-8.9769496
100	0.395284708	-8.0617997

It is seen that the specification is for the return loss magnitude only, and says nothing of the individual real and imaginary parts or, conversely, magnitude and phase return loss. However, more information can be found by measuring the real and imaginary return loss of different systems. With both of these data sets for a system, the engineer may have a better understanding of the system response. To translate from real and imaginary return loss to the magnitude, also known as the reflection coefficient, the following formula is used.

$$Magnitude = \sqrt{(\mathrm{Re})^2 + (\mathrm{Im})^2} \qquad (1)$$

For Gigabit Ethernet, there are four fully balanced twisted pairs of wire, each running at 250 Gbps. Thus, each cable can handle data signals of up to 125MHz. Conversely, the symbol period can be as short as 8ns. According to the specification for the transmission line, this must be true for CAT5 cables up to 100 meters long [3]. Figure 3 is a graph of the Gigabit magnitude return loss specification versus frequency. This specification is for a voltage signal within the bounds of the IEEE 802.3 signal specification.

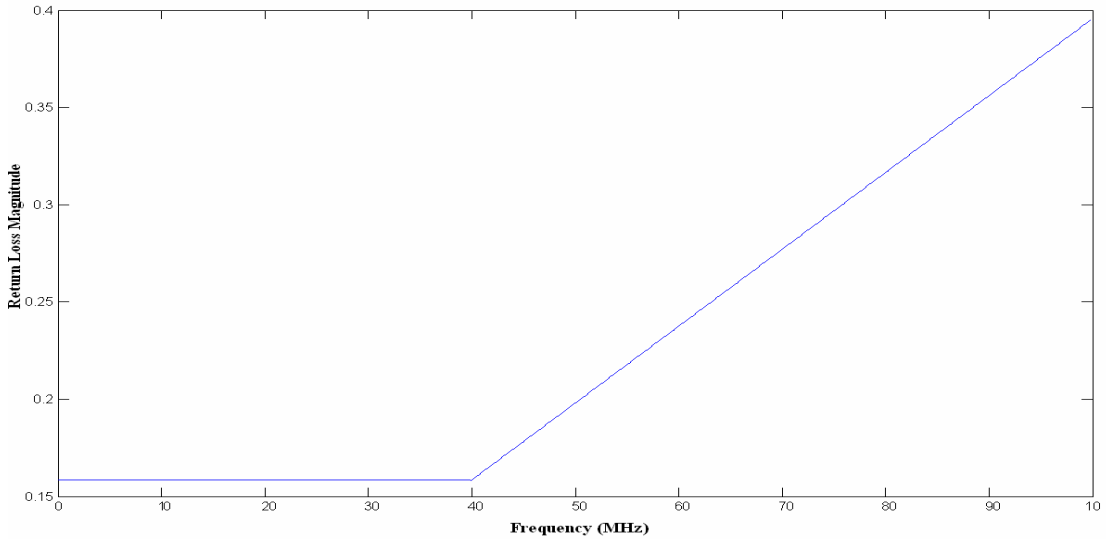

Figure 3: IEEE 802.3.40.1.3 MDI Return Loss Specification

The most notable characteristics from the graph are that the specification is constant for low frequencies and there is an elbow at 40MHz. It can be seen that the specification is frequency dependent and relative to the impedance of the balanced cable, which can be up to ±15% deviation from the nominal value of 100 Ohms differential. The following figure 4 shows the schematic of the equivalent circuit.

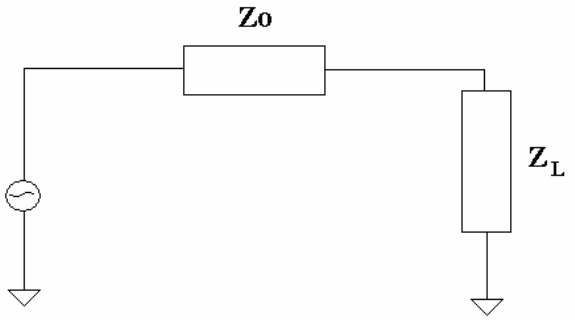

Figure 4: Series Impedance Circuit Block Diagram

In the above figure, Z_o represents the CAT5 cable and Z_L represents the equivalent impedance due to the magnetic jack, the PHY and the termination load. From microwave theory, it is known that the return loss is minimized when these two impedances are matched. This can be shown from the following definition of the reflection coefficient Γ [4], also known as the magnitude return loss S11. In (2), Z_L is the impedance under test, or the load impedance, and Z_o is the line impedance, or the optimal impedance.

$$\Gamma = S_{11} = \frac{Z_L - Z_o^*}{Z_L + Z_o^*} \tag{2}$$

The S21 transmission, which is also the forward power gain, is also defined here.

$$S_{21} = \frac{Z_L}{Z_L + Z_o} \tag{3}$$

By inserting a reference load Z_L of 100 Ohms differential and applying a test voltage, a measurement of the power at the end of the PHY cannot guarantee system return loss compliance due to the unknown transmission line CAT5 impedance. Thus, the only way the return loss can be measured is by somehow including the transmission line impedance as well as the PHY impedance in the calculation. This can be accomplished to a very high degree of accuracy by measuring the return loss of many sample systems and finding the corner cases. By using baluns to balance the transmission line to a differential 100 Ohms, the return loss of the MDI Ethernet port can be measured. This will be shown in more detail in the following paragraphs. However, the Gigabit Ethernet port consists of the magnetics, used to transform the signal, the PHY which transports the signal to the receiver, and the termination load which receives the signal. Thus, all of these components must be taken into account when trying to find total impedance or return loss

seen by the CAT5 cable. The parasitics associated with the vias and other elements must be taken into account as well. This is illustrated in the following figure.

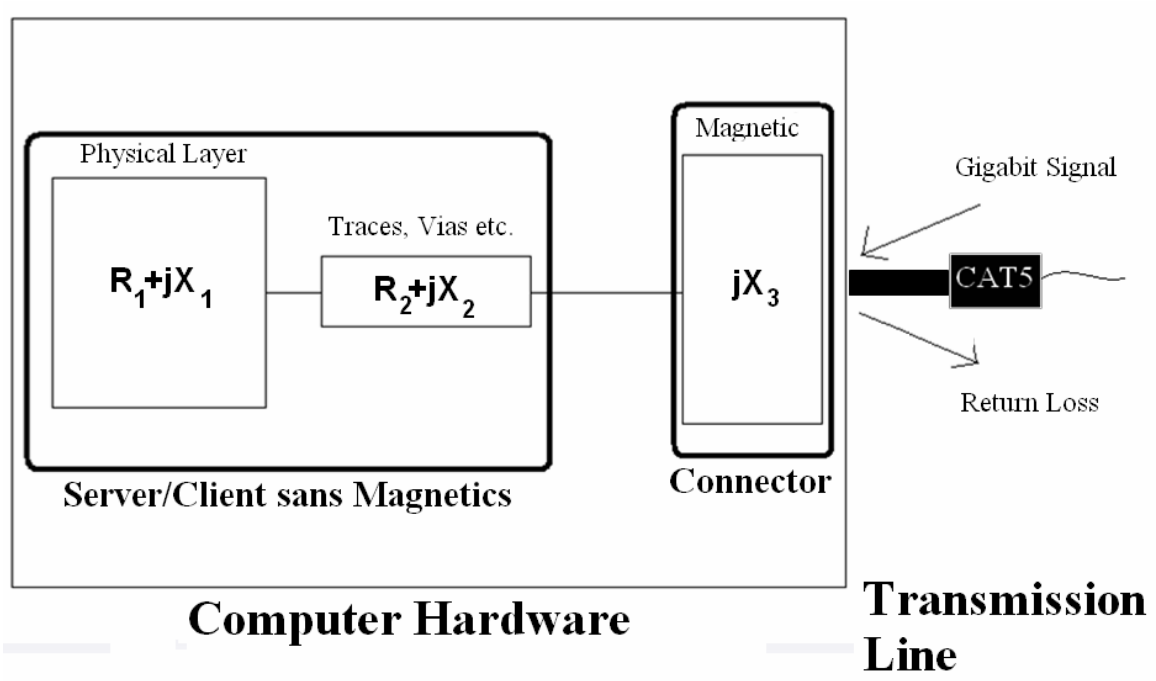

Figure 5: Detailed Ethernet Port Structure

As can be seen above, the magnetic, also known as the RJ45 connector, contributes to the impedance seen looking into the port of the Ethernet terminal of the computer as well. The following figure shows all of the impedances associated with the Ethernet return loss problem.

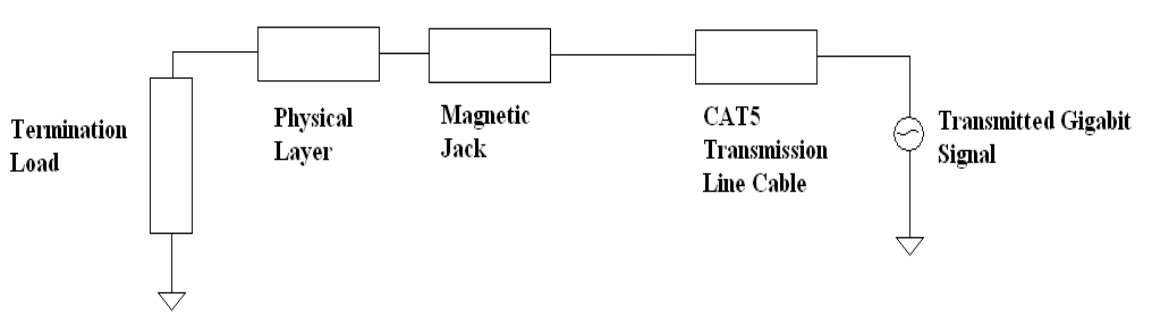

Figure 6: Ethernet Port Block Diagram

It will be apparent from measurements that the real impedance is minimal and the imaginary contribution from the RJ45 to the total impedance is significant. This could not be seen from the magnitude return loss only, and thus this is one of the advantages of using real and imaginary impedance to characterize the overall system. The magnetic is the very front end, and it is what transforms the signal from the CAT5 transmission line to the PHY. Thus, the impedance must be measured after the RJ45 to find the true PHY impedance. However, because the magnetics are actually twisted pairs of wires with a transformer, it cannot be totally removed from the system and thus something has to take its place as the interface between the hardware and the CAT5. The following figure shows the schematic of a normal Ethernet magnetic, where the signal would propagate from the CAT5 cable on the right to the PHY on the left.

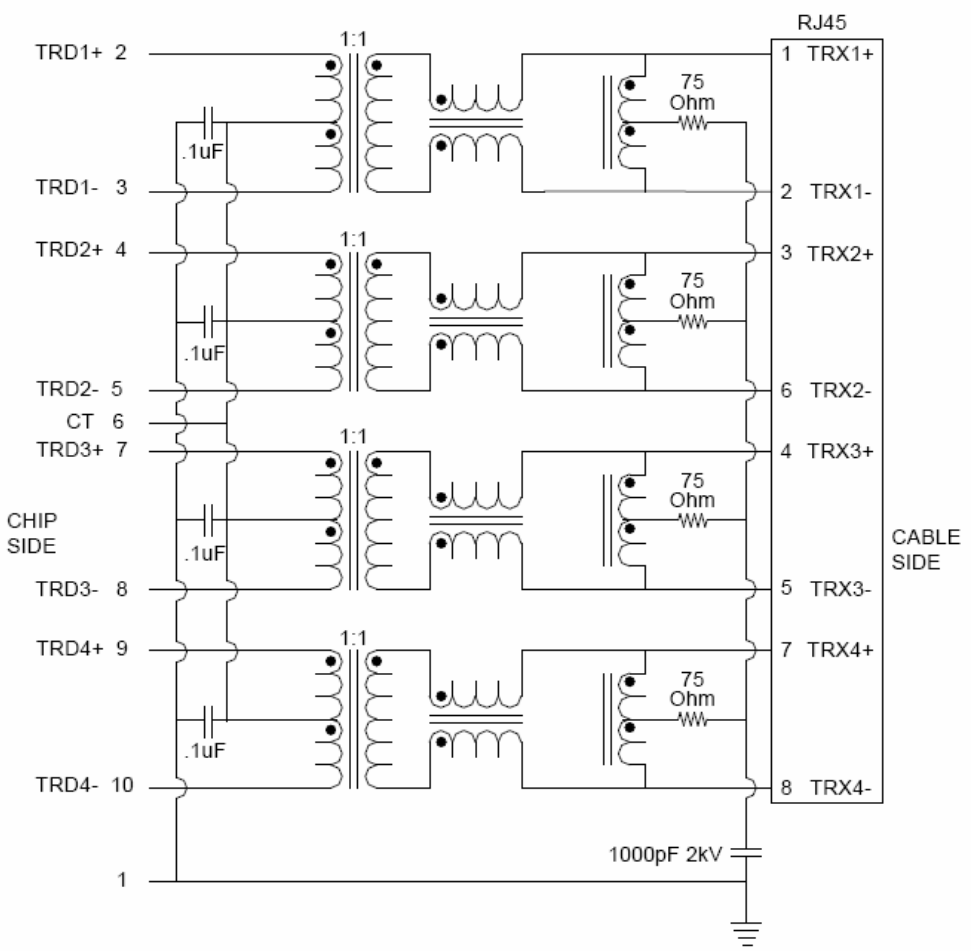

Figure 7: Falco Electronics 1M093_LS2L18 Magnetic

In figure 7, the right side shows the connections to the CAT5 transmission line and the left side, or the chip side, shows the connections to the PHY. In the middle are the 1:1 transformers which are the major reactance contributors. To minimize this added impedance, a replacement without the transformers must be used. Dummy RJ45 connectors were created by running twisted pair wire straight from the CAT5 interface to the output of a magnetic shell according to the design schematic from Falco Electronics in figure 7 [5]. The schematic (figure 8), along with the finished products, can be seen below.

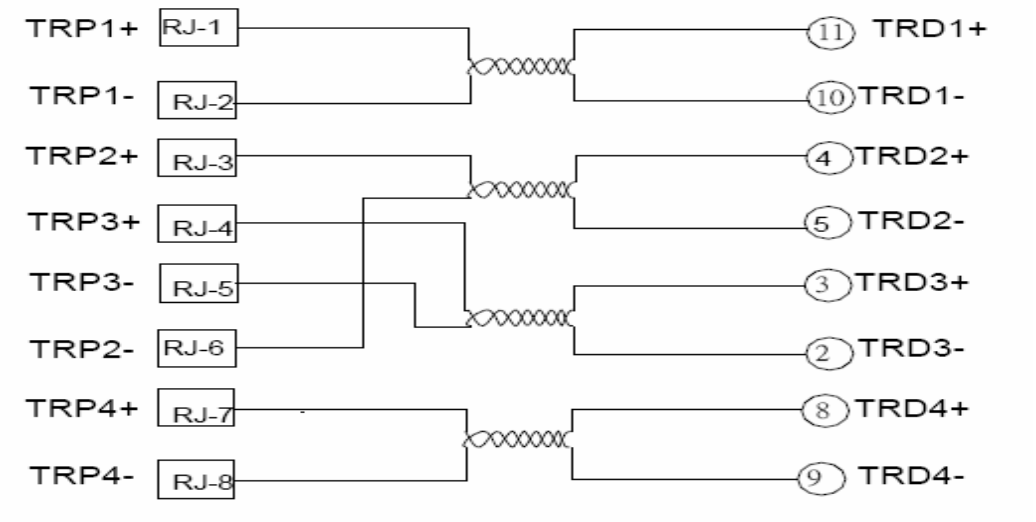

Figure 8: Dummy Magnetic Schematic

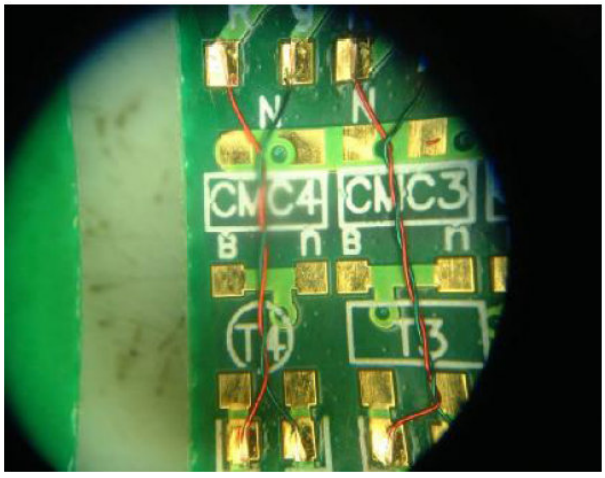

Figure 9: Dummy RJ45 Connector sans Magnetic- Internal View 1

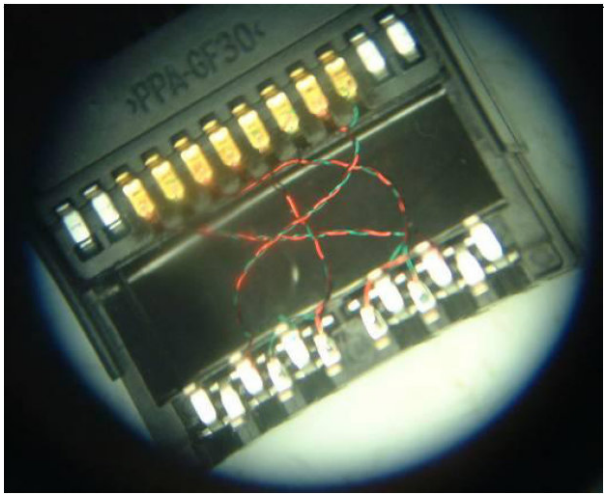

Figure 10: Dummy RJ45 Connector sans Magnetic- Internal View 2

Now that the new twisted pair RJ45 connectors were finished, they had to be characterized so that a tare return loss measurement could be taken when integrated with the PHY. For a true PHY only return loss measurement, the small impedance added by the dummy connector must be subtracted from the overall impedance seen at the CAT5 cable. To this end, a test board was fabricated using the optimum 100Ohm ±1% termination. This board was used to measure the return loss of all dummy and real magnetics and the following figure shows the schematic for this setup.

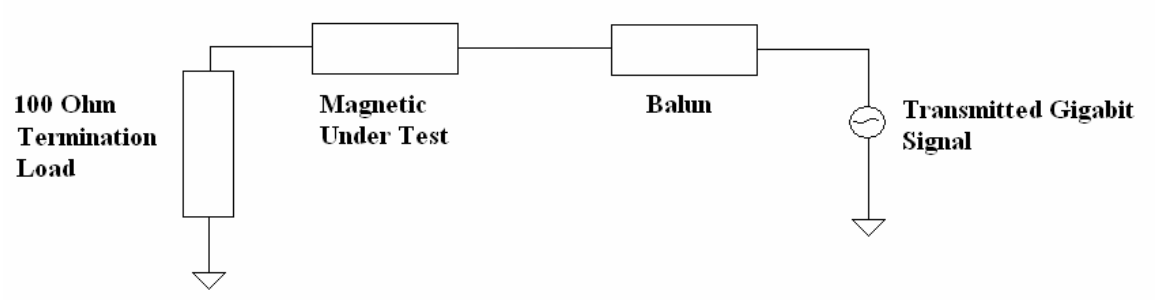

Figure 11: Magnetic Test Setup

Once this board was created and the magnetics were characterized, the magnetics were taken off each individual motherboard or network interface card (NIC), and replaced with a dummy RJ45. The actual return loss of each individual PHY was then measured. The setup for this test is shown in the following figure.

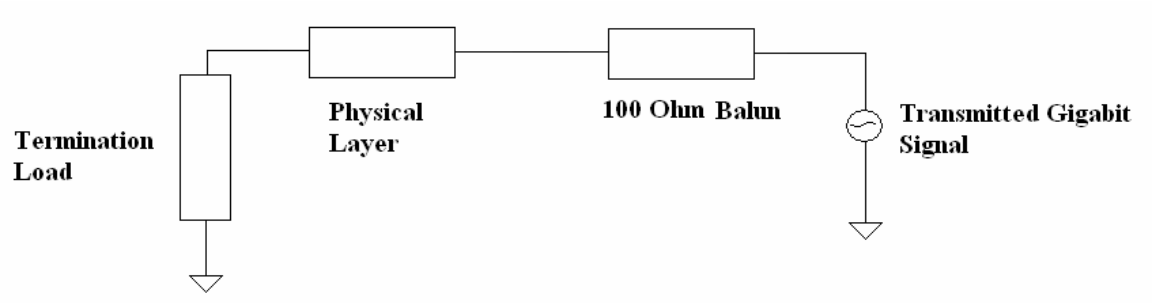

Figure 12: PHY Test Setup

This was done for as many product lines as possible, as well as all magnetics. The following table 2 shows the characterization progress. The systems are described in detail and the quantity characterized is shown. The total number of each system made by Dell is also shown, which gives a percentage of systems characterized.

Table 2: System Characterization Progress Chart

System	Characterized	Total
Server	10	17
NIC	12	12
Desktop	13	16
Notebook	5	19
Total	40	64
Percent Characterized		62.5
Magnetic	400	NA

The return loss information was collected with two Agilent E5062A ENA-L RF Network Analyzers. Using a frequency sweep from 300kHz to 100MHz, 1000 data points were taken per channel, and the real and imaginary return loss data was stored to the hard drive. To bring this data to the computer for processing, a PERL program was written to control the entire test network. It then downloaded the data points into a Microsoft Excel file. A MATLAB program was written to read the excel files and download the data points into an array. This array then could be displayed against the IEEE return loss specification for system compliance. An example of these results is shown below in figure 13.

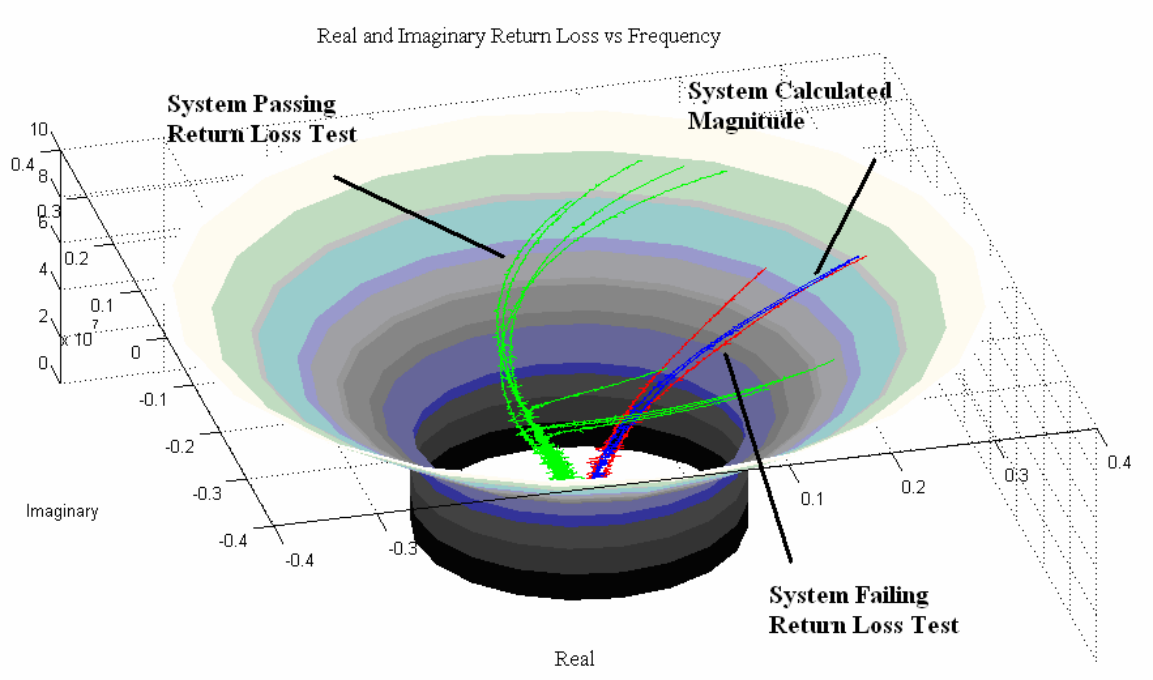

Figure 13: Server Real and Imaginary Return Loss Data versus Frequency with Specification Boundary

In figure 13, the real and imaginary axes are in units of normalized return loss and the Z axis is frequency in 10s of MHz. The cylindrical shape is the IEEE 802.3 magnitude return loss specification applied to all phase angles. The lines inside this cylinder are individual channels of an Ethernet system. The lines on the left represent a system that passed the return loss measurement and the lines in the middle represent a system that failed, as noted on the figure. The lines on the right are the calculated magnitude of the two systems. All server, desktop and NIC return loss data is plotted versus frequency in the following figures. The first plot, figure 14, shows the impedance tree created from all of the system return loss measurements. The lighter lines on the left are the individual channels of each system measured, and the darker lines are the magnitudes of these channels, as noted in the figure.

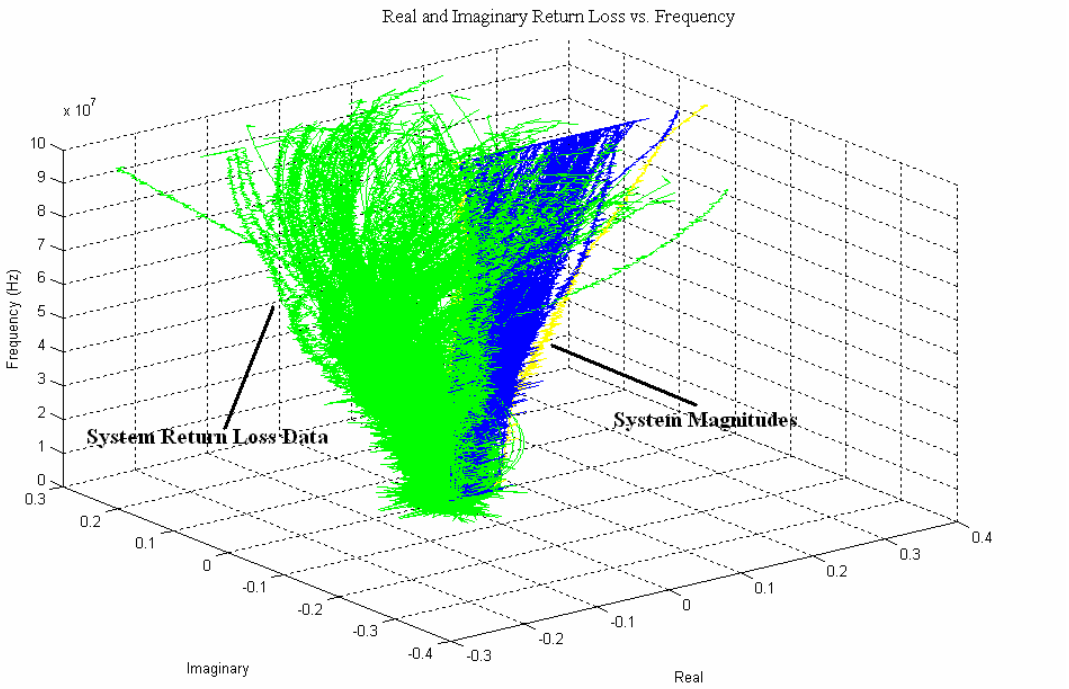

Figure 14: System Real and Imaginary Return Loss versus Frequency

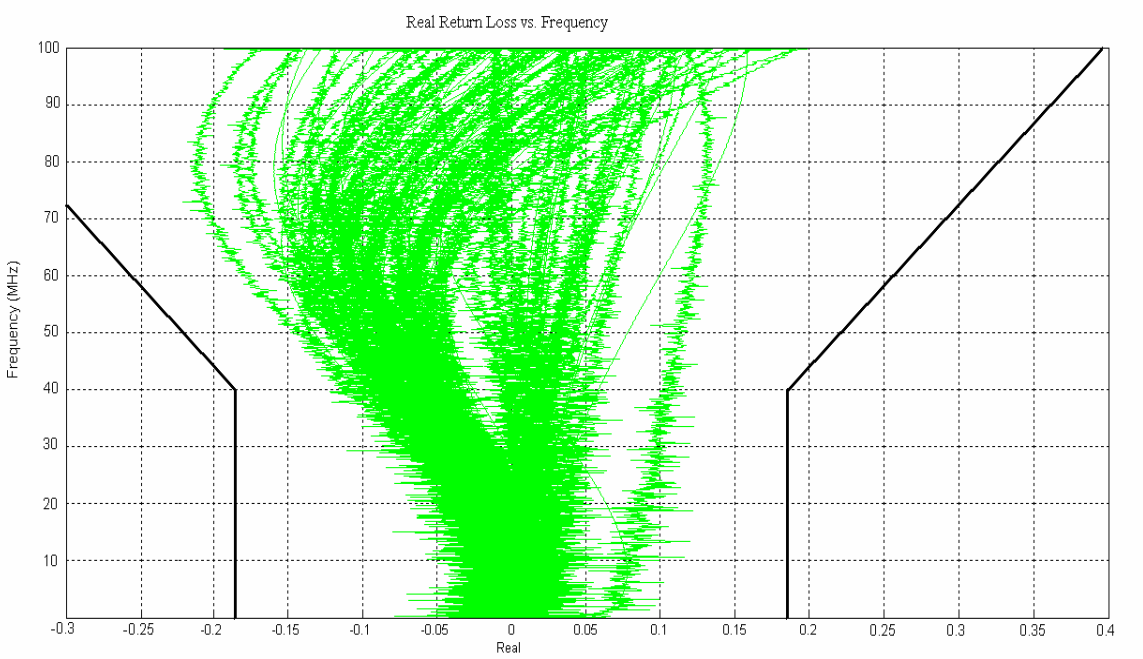

Figure 15: System Real Return Loss versus Frequency

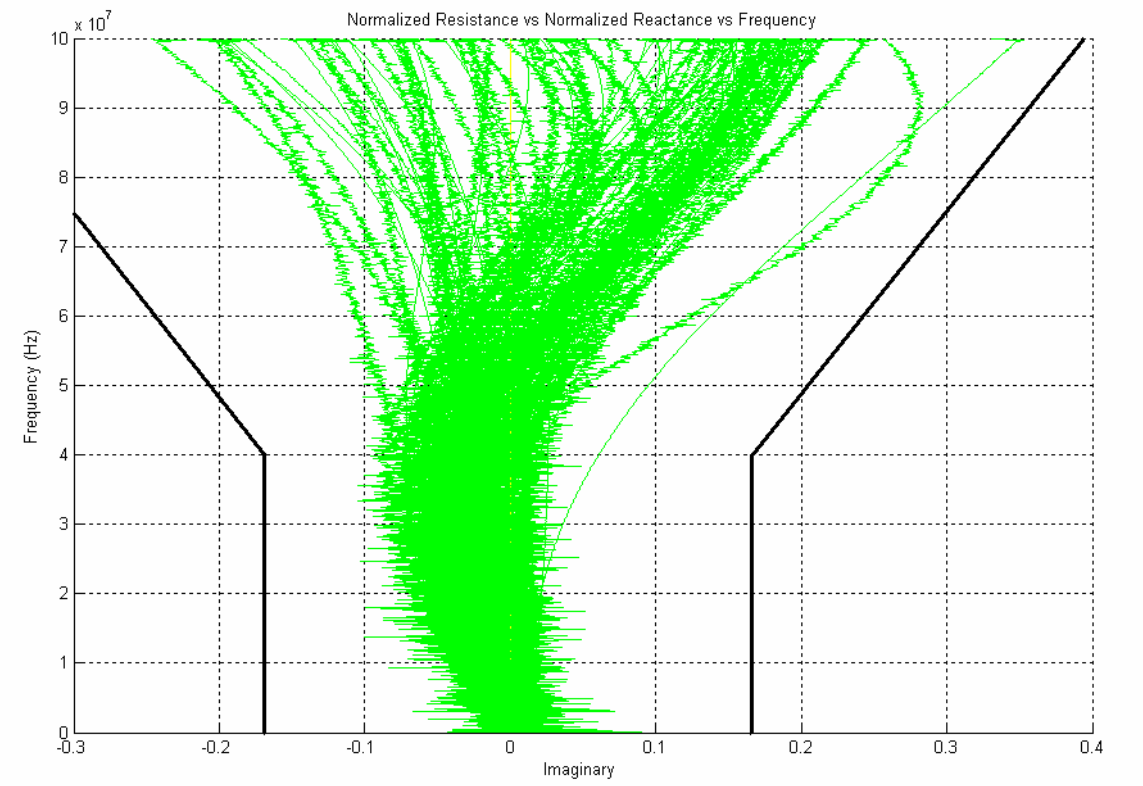

Figure 16: System Imaginary Return Loss versus Frequency

Figures 15 and 16 show only two dimensions of figure 14. Figure 15 shows the real impedances of the measured systems, and figure 16 shows the imaginary impedances of the systems, both of which are absolute values of the return loss. In each of these figures, the magnitude return loss specification is also shown. While the real and imaginary parts of the return loss meet the specification separately, their magnitude must be found and compared to the specification. In addition, the worst case magnetic must be taken into consideration with each system as well, and this is the value that must meet the return loss specification. The need to measure both the real and imaginary return loss variables separately is apparent from the range of values within each graph. Without this information, a system failing magnitude return loss testing could not have its impedance adjusted to any degree of certainty because the engineer would not know how to adjust the system. The Ethernet system cannot be totally characterized by the magnitude return loss only, and these plots prove this.

The magnetics return loss is plotted in the next set of figures. Figure 17 shows the impedance trees of the magnetics with 100 Ohm differential termination, as in figure 11. The light pink lines are the impedances of the individual channels of the magnetics, as noted in the graph. The dark blue lines are their respective magnitudes. Both of the return loss axes are measured in terms of the absolute value of the return loss.

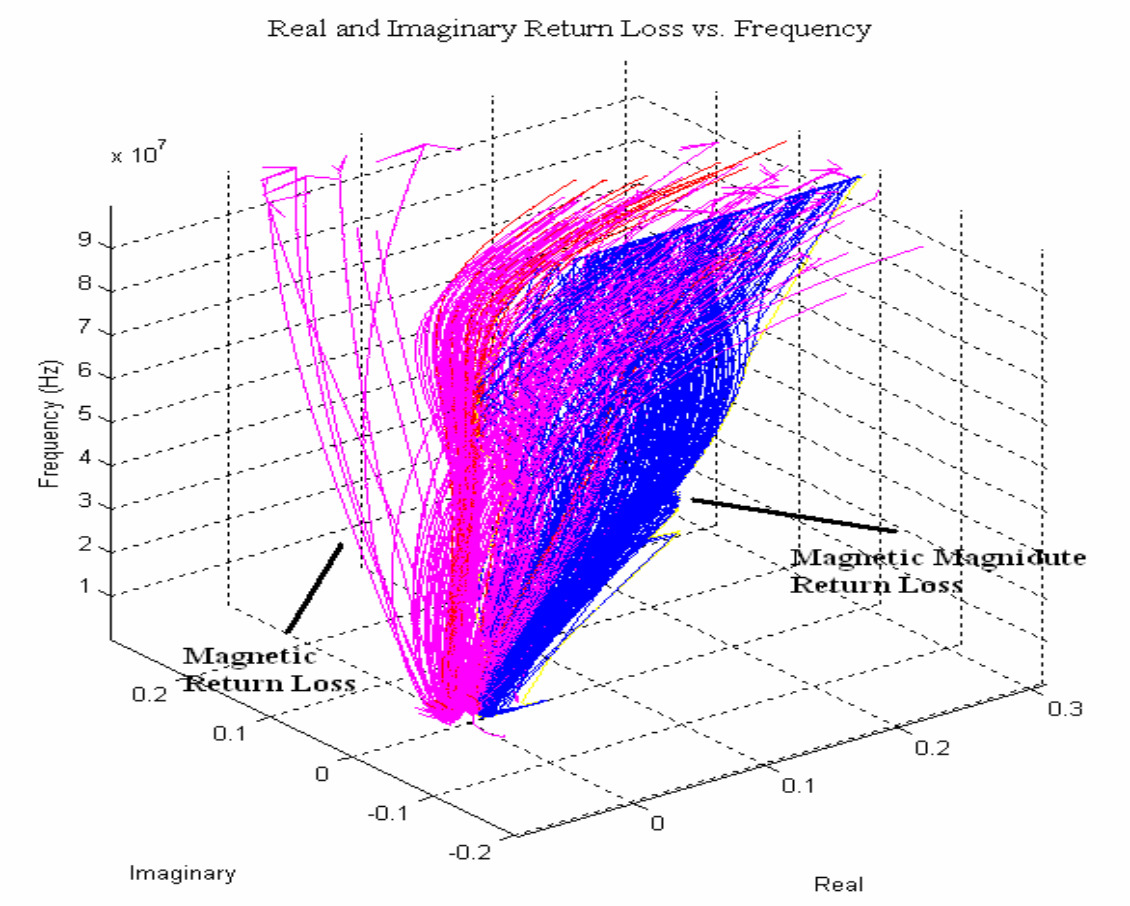

Figure 17: Magnetic Real and Imaginary Return Loss vs. Frequency

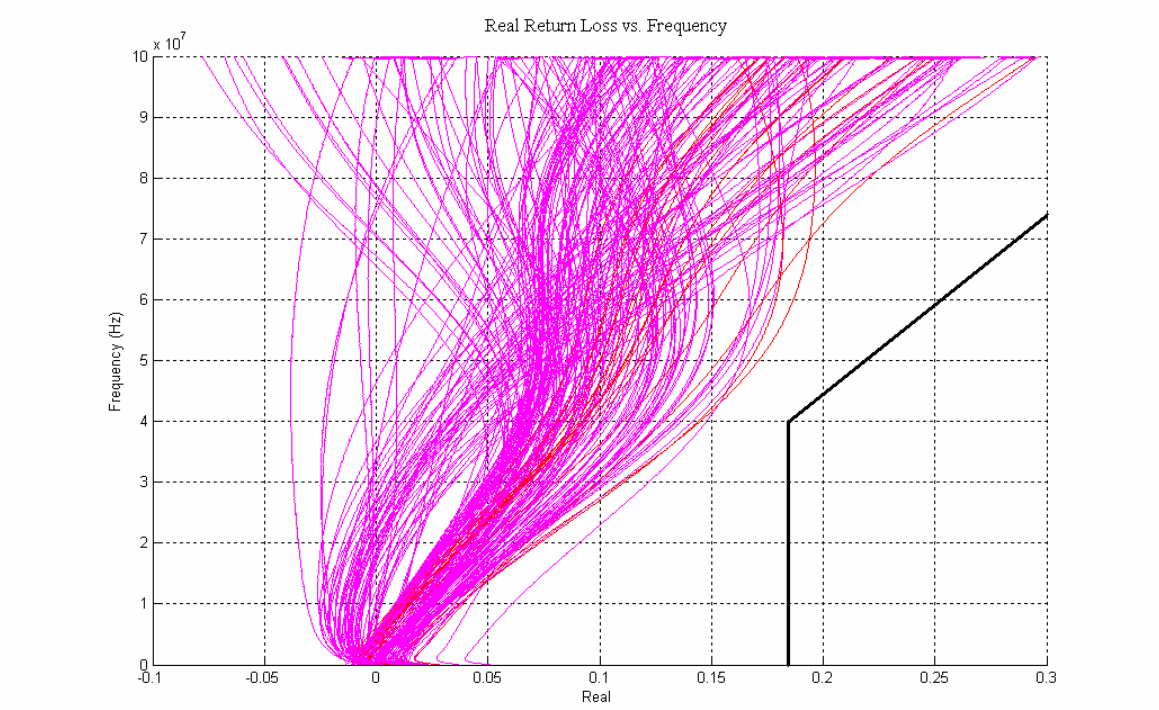

Figure 18: Magnetic Real Return Loss vs. Frequency

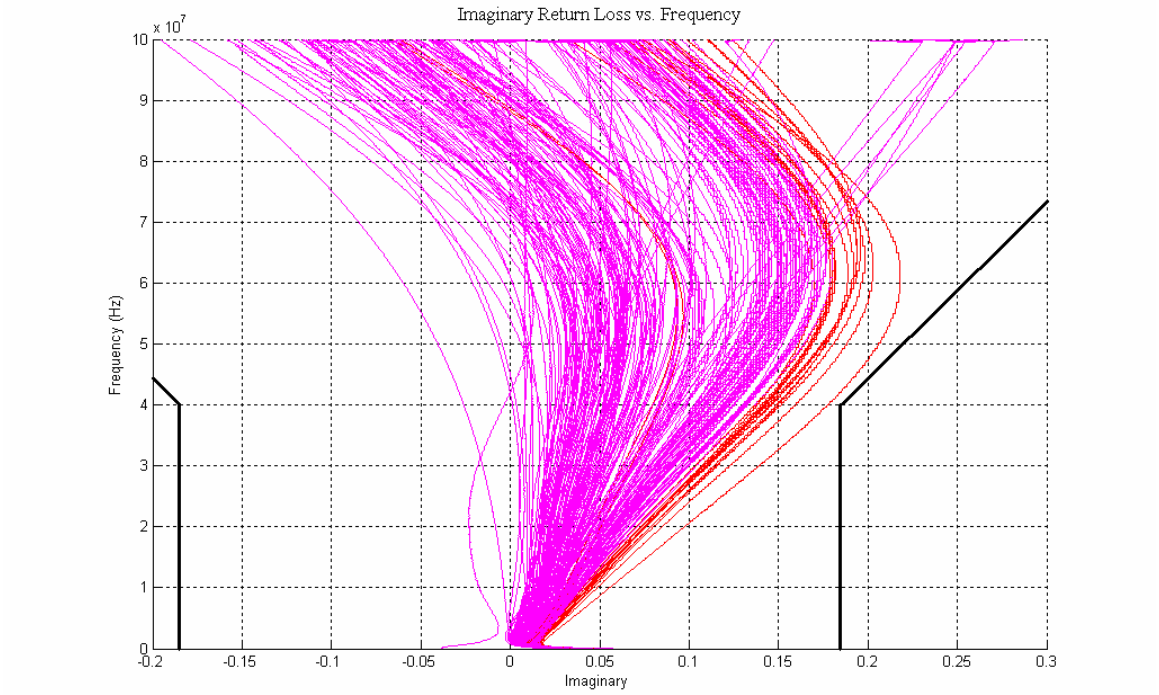

Figure 19: Magnetic Imaginary Return Loss vs. Frequency

Figures 18 and 19 once again illustrate why the magnitude return loss measurement is not sufficient to characterize the magnetic. Without the two variables measured separately, it becomes unclear as to why a certain magnetic would fail testing. With both real and imaginary return loss measured, the magnetic can be seen to be inductive or capacitive, and how large the resistance is. This would more clearly define the magnetic if it should fail the return loss test. Figure 20 shows the worst case system magnitude return loss without magnetics below the specification. By itself, the PHY return loss is within the IEEE specification for magnitude return loss. However, figure 21 shows the worst case system return loss coupled with the worst case magnetic return loss which crosses the specification. This overall return loss clearly does not meet the IEEE return loss specification, and thus the impedance must be compensated such that the return loss is lowered.

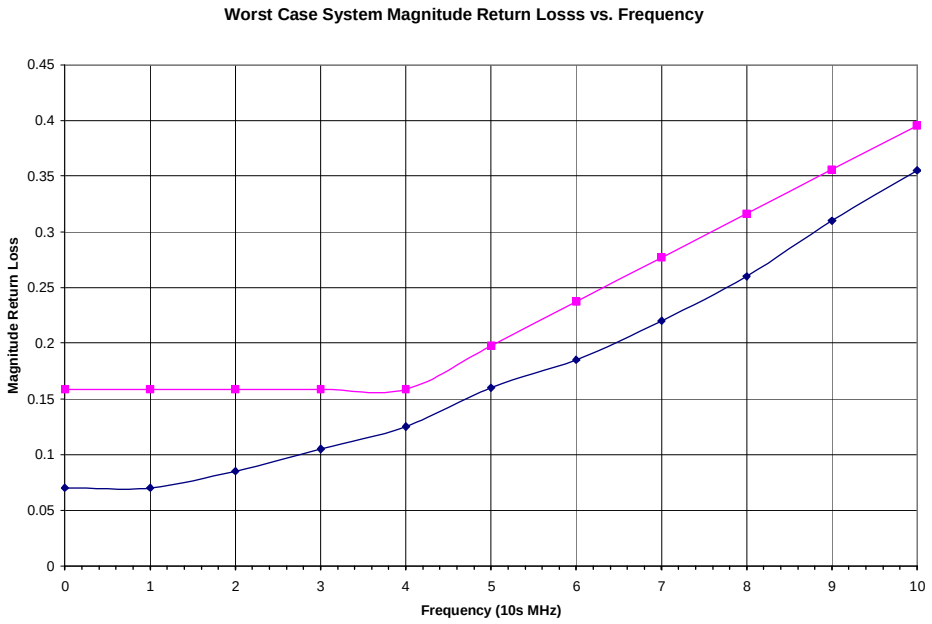

Figure 20: Worst Case Magnitude PHY Return Loss

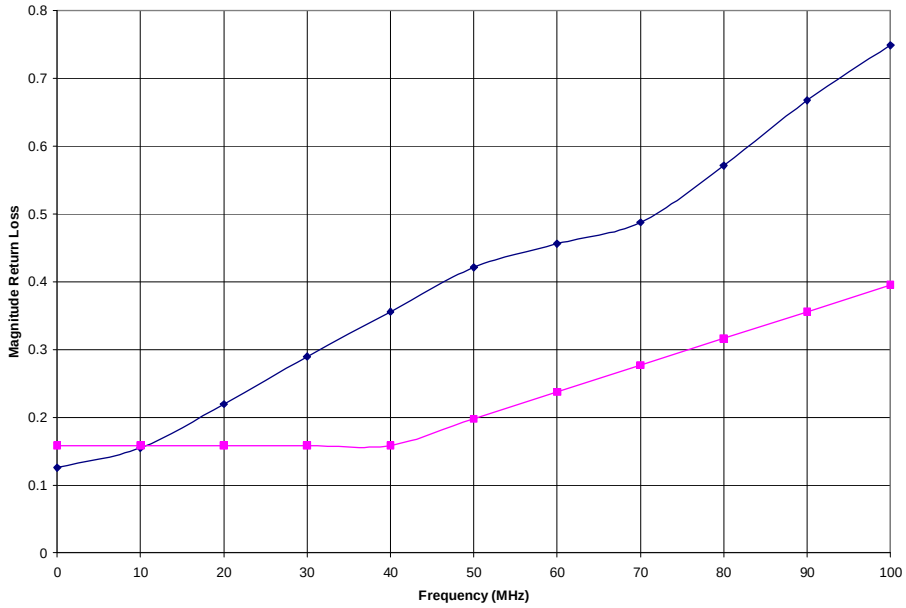

Figure 21: Worst Case Magnitude PHY with Magnetics Return Loss

The variance in real and imaginary return loss measurements shows that testing magnitude return loss only is not enough to totally characterize the systems. The MATLAB program also collected the data points from all systems tested and found the corner case real and imaginary return loss arrays. Figure 22 shows the worst case dark red funnel created that defines the specification for platform return loss, as noted in the figure. The light green and turquoise lines show the individual system return loss. The dark blue shows the magnitude of the respective return loss.

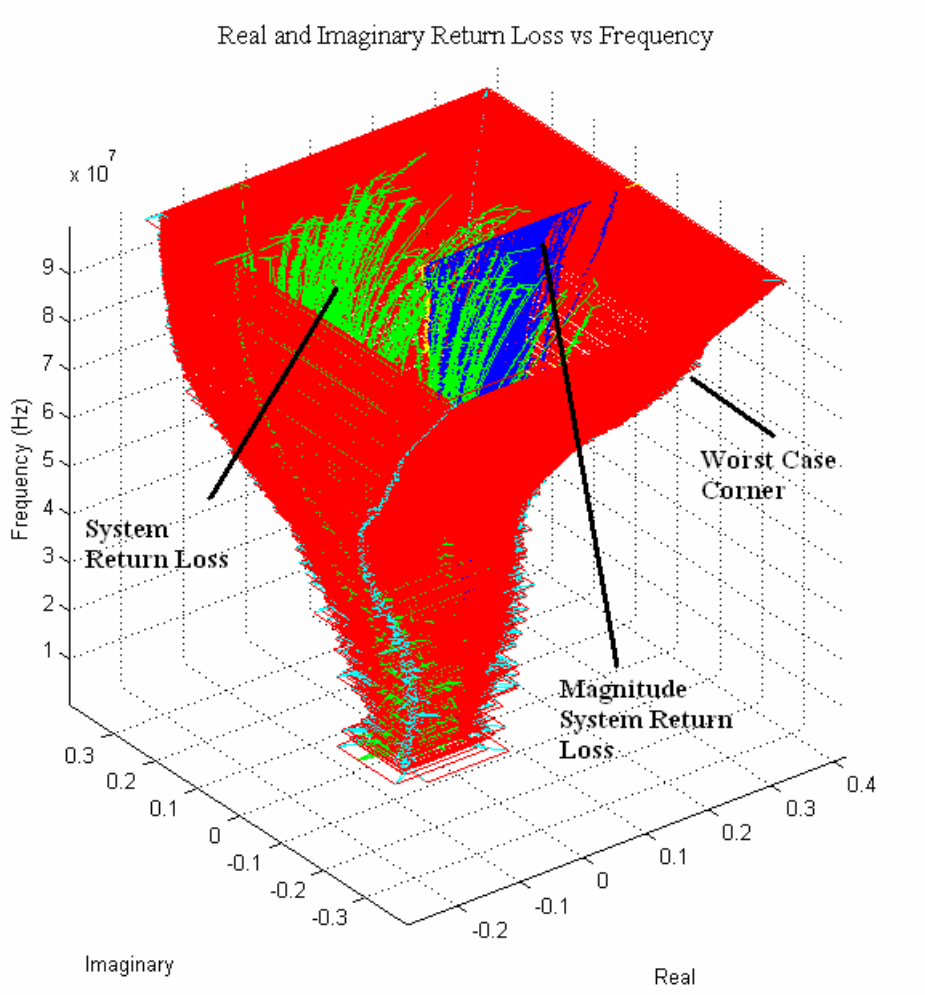

Figure 22: Dell Platform Return Loss Worst Case Funnel

From the above figure, the following table is created. By taking the highest and lowest real and imaginary return loss parameters from the above figure, the absolute worst case return loss situation is defined. Knowing that all systems will fall between the worst case corners, a proprietary specification can be created that defines these real and imaginary corners separately. Instead of just the magnitude, these two variables can more clearly define what part of the impedance is not matched correctly to the CAT5 cable. This proprietary return loss specification can also be used to define the range of possible magnetics return loss allowable when coupled with the worst case system.

Table 3: Dell Computer Proprietary Real and Imaginary Return Loss Characterization

System Frequency (Hz)	Real (High)	Real (Low)	Imaginary (High)	Imaginary (Low)
1 MHz	0.061119343	-0.050362267	0.046368148	-0.021702223
10 MHz	0.077251982	-0.031255372	0.020414868	-0.041692208
20 MHz	0.064084932	-0.048213216	0.021067197	-0.064664492
30 MHz	0.052870154	-0.071210773	0.03432247	-0.093303837
40 MHz	0.053446165	-0.10191251	0.02599266	-0.081756063
50 MHz	0.058374778	-0.143506624	0.04604584	-0.087161839
60 MHz	0.072332121	-0.17573411	0.072388649	-0.107135274
70 MHz	0.083412498	-0.20069395	0.122141942	-0.117430978
80 MHz	0.103824012	-0.210990127	0.1636215	-0.141440392
90 MHz	0.124790199	-0.202992642	0.208927915	-0.167482063
100 MHz	0.205758035	-0.170430664	0.241548091	-0.220784374

The worst possible real and imaginary return loss parameters are now defined for the systems sans magnetics. This data will be useful in creating an onboard active control of the PHY impedance. To translate the new individual real and imaginary return loss specification, the system with the worst case magnitude at each frequency is used. The following table gives this new proprietary specification.

Table 4: Dell Computer Proprietary Magnitude Return Loss Specification

Frequency (Hz)	Magnitude	dB
1 MHz	0.07	-23.098
10 MHz	0.07	-23.098
20 MHz	0.085	-21.4116
30 MHz	0.105	-19.5762
40 MHz	0.125	-18.0618
50 MHz	0.16	-15.9176
60 MHz	0.185	-14.6566
70 MHz	0.22	-13.1515
80 MHz	0.26	-11.7005
90 MHz	0.31	-10.1728
100 MHz	0.355	-8.99543

This worst case system magnitude return loss data defines the new bound that all systems sans magnetics must meet. While it is not as helpful as separate real and imaginary return loss specifications would be in defining the impedance mismatch between the CAT5 and the PHY, these are not covered in the IEEE specification and thus are not widely accepted. Using this and the IEEE 802.3 specification, a proprietary magnetics return loss specification can be created as well for distribution among vendors. By subtracting the system impedance magnitude from the total impedance, the remainder becomes the largest magnitude impedance allowed by the magnetics. This is shown in the following magnitude formula and schematic.

Total Impedance= [Magnetic Impedance] + [PHY Impedance] (4)

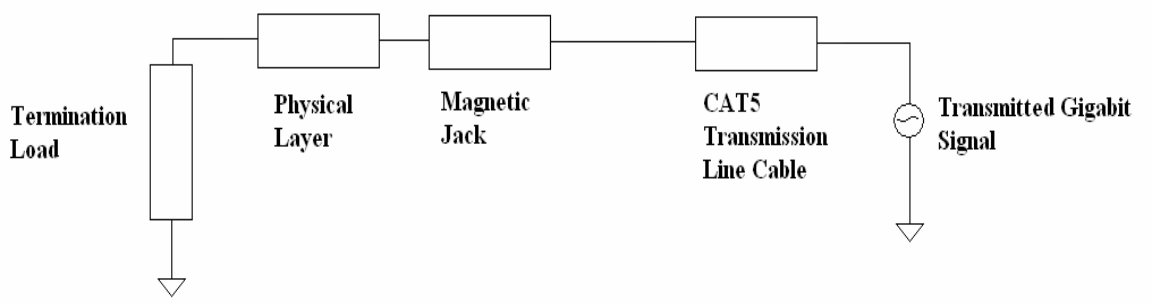

Figure 23: Ethernet Port Block Diagram

By using equation (2), the impedances can be translated to return loss, and the largest magnetic magnitude return loss can be found. By having systems that fall under the new proprietary return loss specification, and magnetics that fall under this new magnetic return loss specification, all servers, desktops, notebooks and NICs are guaranteed to fit

under the IEEE return loss specification. The following is the new proprietary magnetics return loss specification based on the characterized data.

Table 5: Proprietary Magnetics Magnitude Return Loss Specification

Frequency (Hz)	Magnitude	dB
1 MHz	0.088489	-21.0622
10 MHz	0.088489	-21.0622
20 MHz	0.073489	-22.6755
30 MHz	0.053489	-25.4347
40 MHz	0.033489	-29.5019
50 MHz	0.037642	-28.4865
60 MHz	0.052171	-25.6514
70 MHz	0.056699	-24.9284
80 MHz	0.056228	-25.001
90 MHz	0.045756	-26.791
100 MHz	0.040285	-27.8972

After the creation of the proprietary system and magnetic specifications, a fiscal analysis was conducted to see the benefits possible by minimizing the return loss of the magnetic and the PHY. Currently, computer companies use separate RJ45 connectors on each platform. For example, Dell Computer has about 80 different systems being produced at a total rate of over 30,000 computers a day. A solution that could normalize the impedance seen at the input Ethernet port, and thus the return loss, could save hundreds of thousands of dollars per system lifetime and millions of dollars in total. Currently, each RJ45 costs around $3.69. By standardizing the impedance, one magnetic could be used on all system lines. Because the parts are bought in such bulk, cost of scale could reduce the cost of each RJ45 to around $2.70. This would be a savings of $0.99 per part, which over a 600,000 part production lifetime would save around $600,000 dollars per system. The following schematic illustrates this point as well.

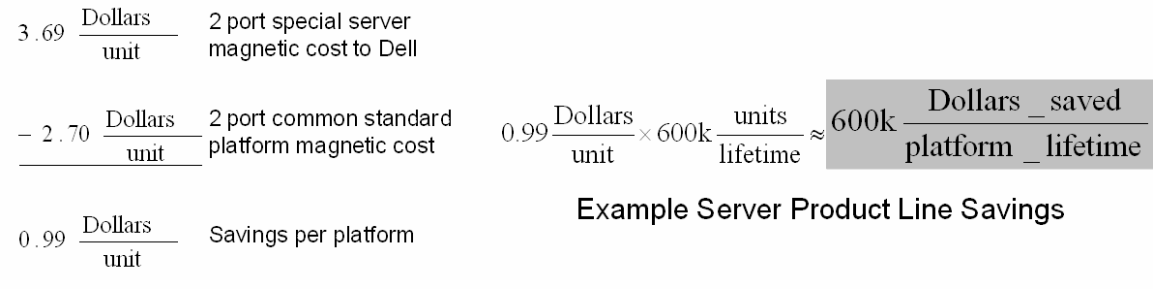

Figure 24: Fiscal Analysis

It can be seen that normalizing the input impedance can impact the bottom line in a significant way. This is the impetus for the creation of an impedance matching circuit for the physical layer and a loop to control it.

CHAPTER III

PHY MODELING AND LADDER NETWORK THEORY

3.1 Problem Overview and PHY Model Creation

While these new proprietary system and magnetic magnitude specifications do comply with the IEEE 802.3 MDI return loss specification, they do not solve the impedance variance problem associated with the magnetic and the PHY. To ensure that the impedance of all systems fall within the bounds of the IEEE return loss specification, the creation of impedance matching circuitry and a control loop becomes necessary. With this in mind, the real and imaginary return loss data can be used to create a model of the impedance seen at the Ethernet port. This solution, as stated previously, should be bounded by cost, area, power and the feasibility of an on board solution. These factors directly rule out using a programmable FIR filter, due to its large area and relatively large cost of adding another chip. Therefore, a solution will be investigated that tries to create poles and zeros while manipulating those inherent in the PHY impedance. While the proposed solution will not have the same accuracy compared to a multiple tap FIR filter, the tradeoffs would be decreased complexity and lowered cost.

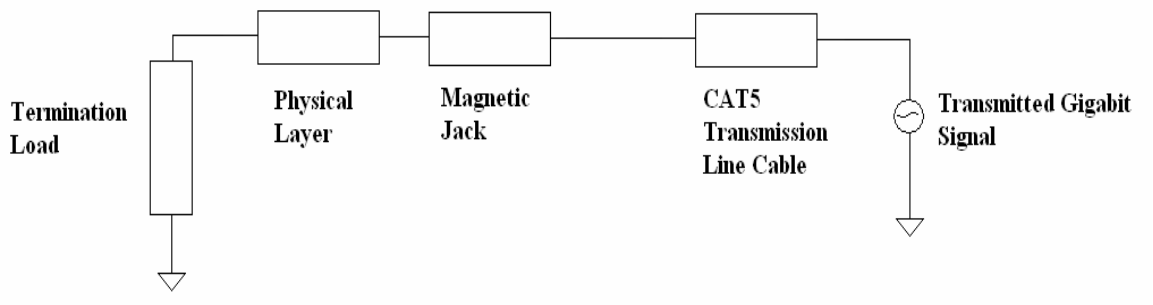

Figure 25: Ethernet Port Block Diagram

To begin the process of creating an impedance matching circuit, the PHY itself will be examined and a circuit model will be created. From the overall Ethernet port schematic, reprinted above, the impedance from the magnetic RJ45 connector will be ignored. This is done because an active impedance matching circuit could not be implemented prior to the magnetic due to its physical positioning on the motherboard. The RJ45 is a magnetic coupling transformer that electrically isolates the CAT5 line from the PHY. A system that matches the overall port impedance including the RJ45 would have to involve multiple outside manufacturers and this is an unreasonable request at this point of the project. If the magnetic is removed from the equation and the coupling transformer put within a variable impedance chip solution, then the proposed solution could reduce the S11 return loss seen at the input even more. In future work, the magnetic impedance could be taken into account and addressed with a more accurate solution. The difference between the two solutions is illustrated in the following figure.

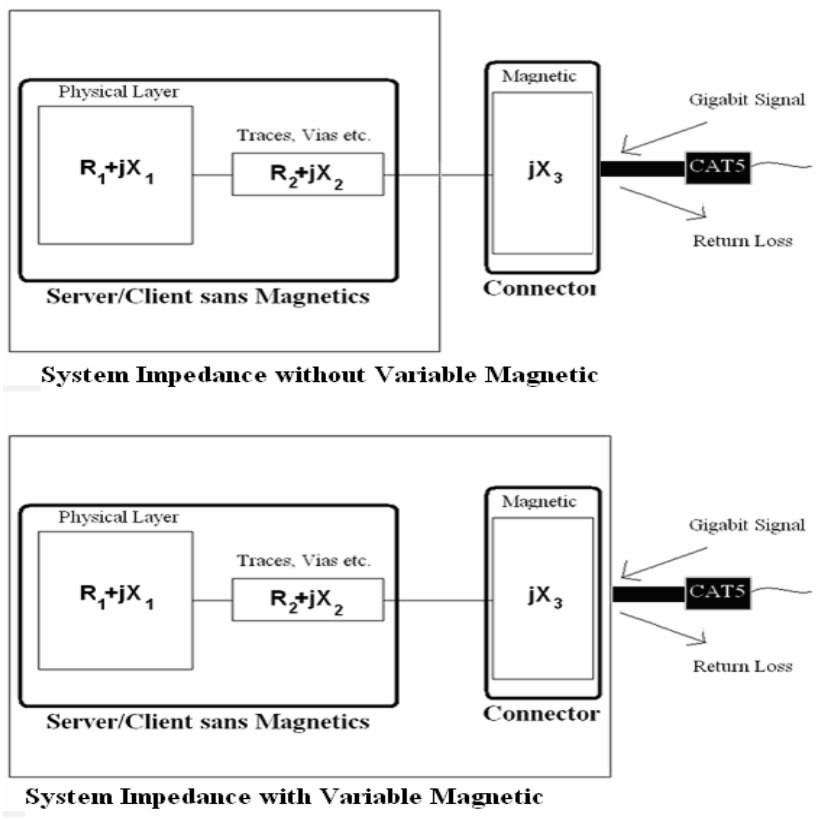

Figure 26: Ethernet Port Structure

The first step in the process of modeling the PHY impedance is to note that with real and imaginary return loss measurements instead of just magnitude, a greater amount of information about the nature of the impedance can be derived. This return loss data can be translated to impedance from the reflection coefficient which is also known as the S11 return loss. The following schematic shows the optimal and load impedance schematic.

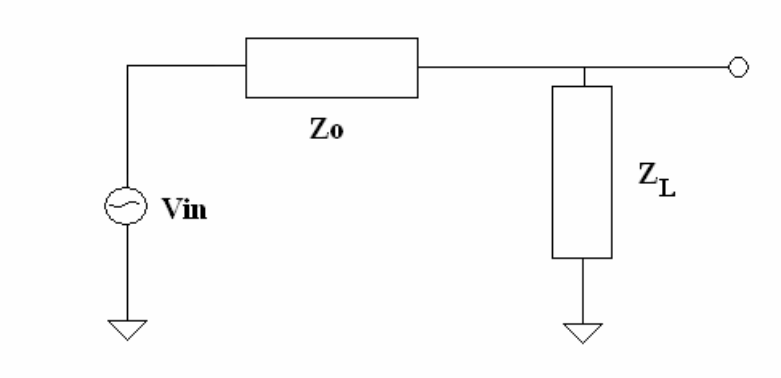

Figure 27: Reflection Coefficient Schematic

The reflection coefficient equation is now repeated for simplicity, where Z_L is the load impedance and Z_o is the optimal or line impedance.

$$\Gamma = S_{11} = \frac{Z_L - Z_o^*}{Z_L + Z_o^*} \tag{5}$$

From the figure, the return loss is seen as a reflection at the input of Z_L. From the IEEE 802.3 Gigabit Ethernet Specification, it is known that ideally Z_L is a termination of 100 Ohms differential. Therefore, for each line, the optimal impedance is 50 Ohms. By substituting this optimal impedance into the equation and solving for the load impedance, the following equation can be used to find the resistance and reactance of the PHY from its return loss data.

$$Z_L = \left[\frac{1+\Gamma}{1-\Gamma}\right] Z_o^* = \left[\frac{1+\Gamma}{1-\Gamma}\right](50 + j0) \tag{6}$$

It can be seen that the conditions of minimum return loss happen when the Z_L is equal to the complex conjugate of Z_o. If the S11 real and imaginary return loss parameters are known, the real and imaginary impedances can be found as well. From (5) it can be seen that by setting the Z_L load equal to the complex conjugate Z_O^* source

impedance, the return loss approaches zero. Because the CAT5 cable defines the source impedance as 50 Ohms, the load resistance ideally is set to 50 Ohms. The reactance of the PHY must be minimized to create the lowest possible S11 return loss. However, the range of values in this impedance is +-7.5% due to the variance allowable in the CAT5 cable, which means that an absolute value for the reference cannot be used. The reference must be based on the CAT5 transmission line impedance, and the PHY impedance should be matched to this. Using Matlab and the characterized return loss data from the previous chapter in table 3, the following table is created to show the highest and lowest real and imaginary impedances found from the PHY characterization.

Table 6: Range of PHY Impedances

Frequency (MHz)	ZL High (Ohms)	ZL Low (Ohms)
1	56.241 - 5.2046i	45.196 + 1.9946i
10	58.292 - 2.3465i	46.833 + 3.9447i
20	56.784 - 2.3958i	45.054 + 5.8955i
30	55.461 - 3.7863i	42.672 + 8.0471i
40	55.517 - 2.8970i	40.244 + 6.7151i
50	55.905 - 5.1716i	36.910 + 6.6094i
60	57.114 - 8.3106i	34.336 + 7.6734i
70	57.155 - 14.256i	32.549 + 8.0485i
80	57.989 - 19.766i	31.472 + 9.4860i
90	58.557 - 24.937i	31.554 + 11.321i
100	65.119 - 35.043i	32.264 + 15.486i

The following figure is a reprint of the circuit in question.

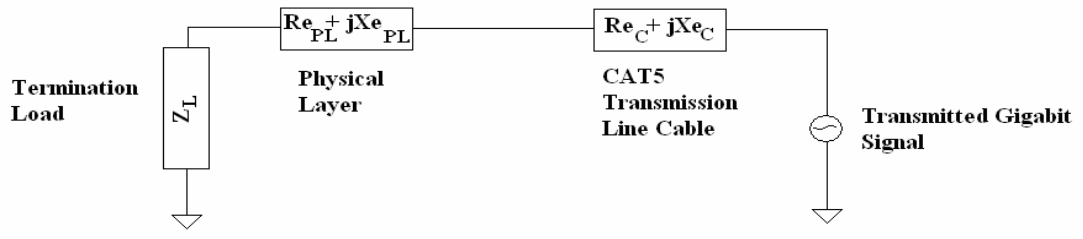

Figure 28: Ethernet Test Structure

There are a few pieces of information that can be found from the above table and figure. First, a simple matching of the resistance of the PHY to that of the CAT5 cable resistance will not suffice. The PHY also has a reactance associated with it which plays a major role in the overall magnitude impedance and phase, especially at higher frequencies. By changing the reactance of the PHY, the shape, or the poles and zeros of the frequency response, can be manipulated and ideally the phase linearized. By linearizing the phase versus frequency, this gives a constant group delay. As a first order approximation, the reactance should contribute little in the frequency band of the majority of transmissions such that only the resistance is seen. Conversely, a purely resistive termination has no frequency dependant component; changing its overall resistance will produce a shift of the entire return loss plot, but it will not, by definition, change the shape. If the source and load resistances are matched and the reactance is linearized, the phase change would approach zero due to the limited effect of the imaginary components. Because the reactance plays a smaller role at lower frequencies, this is where the resistance should be matched and vice versa. Thus, it can be seen that two separate matches are required: the resistance and the reactance which would in turn lend itself to two separate control loops.

From table 6 a model of the PHY can now be derived. The S11 return loss plots in the previous section show vastly fluctuating PHYs, which in turn give rise to a wide range of possible impedances. It can be seen from the imaginary return loss plots that both positive and negative return loss values are possible in the same PHY. This means that both poles and zeros are present in the complex impedance, and thus inductors and capacitors will be required in the model design. To create an impedance zero, a series inductor should be used and the impedance pole can be modeled with a shunt capacitor.

According to microwave theory in [6] and [7], a terminated transmission line can be modeled as a lumped RLC ladder network. The following schematic shows this segmented model.

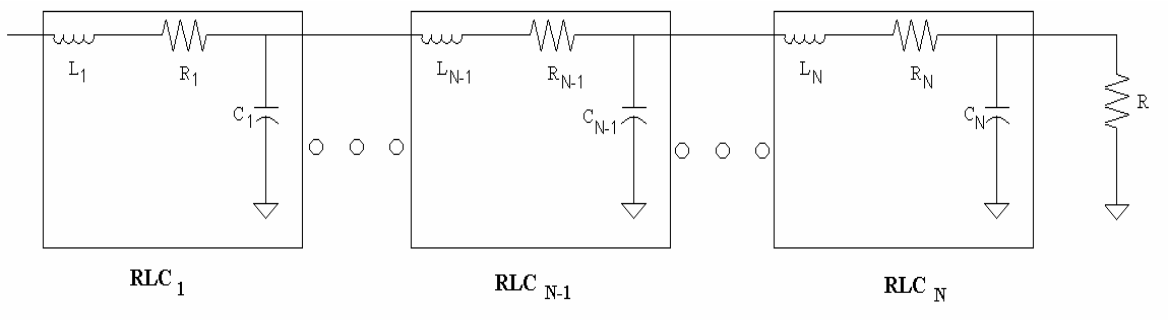

Figure 29: Transmission Line Model

An infinitely long transmission line can be modeled with an infinite number of RLC blocks with ideal resistive termination. One segment of this lumped transmission line will be used as the basis of the simplified PHY model. The impedance circuit model derived so far using ideal lossless elements is shown in the following figure.

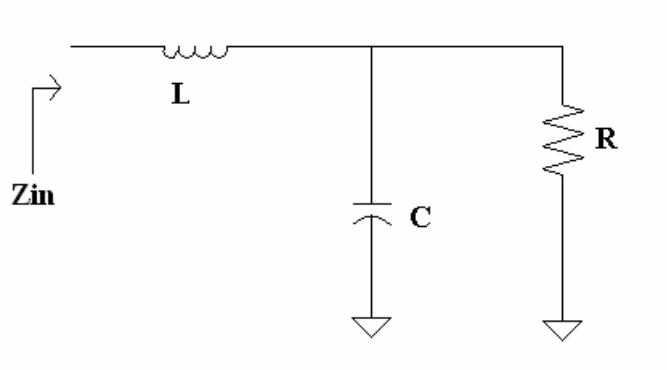

Figure 30: First Approximation Pole/Zero Model

As a first approximation, it is assumed that R is the load termination resistance, and that any parasitic resistance associated with the inductor L can be lumped into this inclusive R. The model in figure 30 gives the following Laplace impedance, where s is the complex variable $j\omega$.

$$Zin = Ls + \frac{R}{1 + RCs} \tag{7}$$

While this circuit can model the general pattern of the PHY impedances, it is not versatile enough to properly model the imaginary return loss family shown in the data reported in the previous section. It can be seen directly from those figures that a higher order circuit is required to more accurately replicate this return loss data. Therefore, another inductor is added to the circuit in series to create more poles and zeros. The following schematic shows the new circuit topology where the parasitic capacitors are ignored. They will be incorporated in the following subsection.

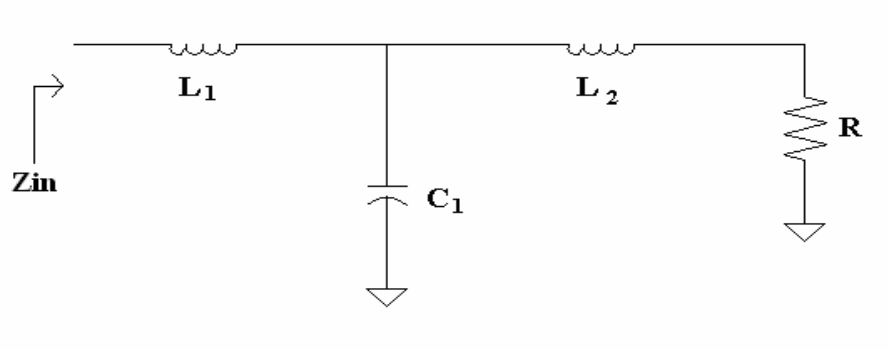

Figure 31: Second Approximation Pole/Zero Model

While the addition of this component does not seem drastic, this new circuit model gives the following input impedance.

$$Zin = \frac{L_1 L_2 C s^3 + RCL_1 s^2 + \left[L_1 + L_2\right]s + R}{L_2 C s^2 + RCs + 1} = \frac{R\left[1 - L_1 C\omega^2\right] + j\omega\left[L_1 + L_2 - L_1 L_2 C\omega^2\right]}{1 - L_2 C\omega^2 + j\omega RC} \tag{8}$$

It can be seen that there are now three zeros and two poles, albeit non-independent, in the new circuit. This higher order model now can properly represent all of the range of PHY line impedances. The problem now becomes a search for the component values needed for the model. This can be determined from the frequency bandwidth of interest and the return loss data plot ranges which in turn gives us midpoint component values and their ranges, respectively. Applying this approach and employing reasonable capacitance values, the following table shows the component ranges used during simulation.

Table 7: PHY Component Ranges

Component	Minimum	Median	Maximum
L1 (nH)	50	100	100
C (pF)	25	45	50
L2 (nH)	50	75	100
R (Ohms)	44	47	NA

The L1 inductor is always kept at 100nH while the L2 inductor is changed; this is done because the zeros and poles are related and thus changing one inductor can adequately change the frequency response while keeping the 3 dB bandwidth around 70 MHz. From the return loss plots, it is seen that the resistance associated with the PHY is on average around 6 Ohms and this is lumped with the load resistance. Therefore, with the matching resistance circuit described later in this chapter, there is no maximum load resistance over the 50 Ohm ideal value. There are a few approximations used in the formation of this model. As stated previously, it is assumed that the series inductor resistances can be lumped together. Also, the pi model would suggest that there are other parasitics associated with the inductors, and these are ignored in the first order approximation. The following figure shows the range of PHY impedance values characterized and the bounds

for the model created. The high and low impedance bounds are using the worst case ranges shown in Table 7.

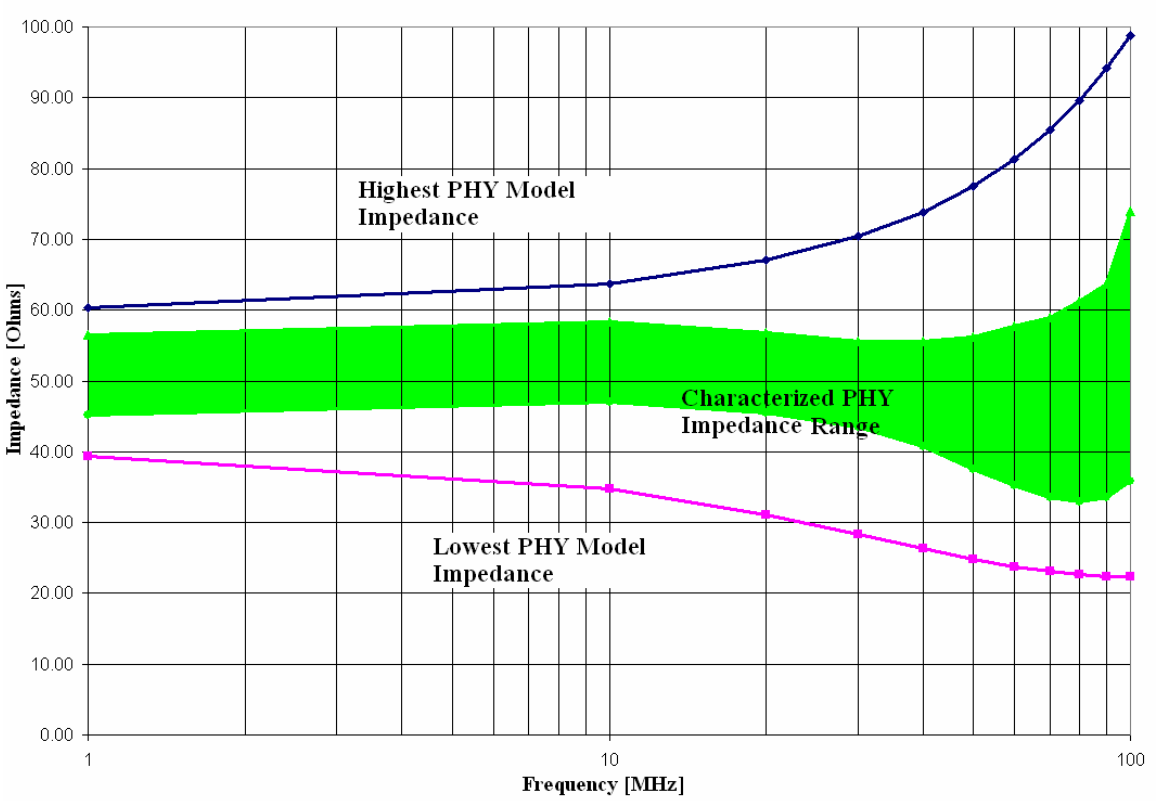

Figure 32: PHY Model and Characterized PHY Impedance

The next chapter goes into greater depth with Cadence simulations of not only the PHY return loss, but the voltages measured and solution presented as well.

The matching circuitry presented in the following sections should be run once when the computer is turned on. This means that the match is not a real time match but a calibration at turn on.

3.2 Impedance Matching

3.2.1 Real Impedance Matching

As stated earlier, the real impedance of the circuit should be matched at low frequency at computer turn on so that the imaginary impedance does not play a large role. In the case of Gigabit Ethernet, there is an elbow in the magnitude return loss specification at 40 MHz. Thus, it is determined that low frequency should constitute 1 MHz to 10 MHz. The following circuit schematic shows the resistances that play a role in this low frequency real impedance match and the introduced current source.

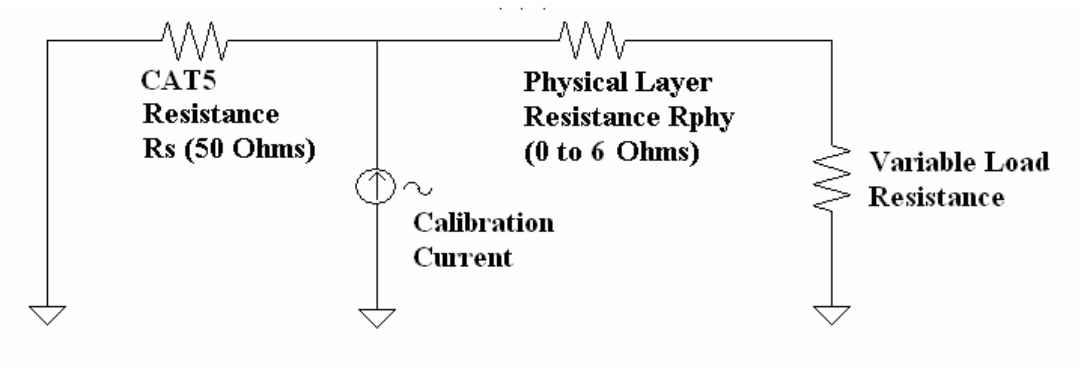

Figure 33: Low Frequency Circuit Schematic

In the above figure, the total resistance seen at the positive port of the input current source is the parallel combination of the CAT5 cable and the PHY. Because the CAT5 and PHY resistances are unknown and the goal is to match these two to each other, the first step should be to create a reference that is not associated with their combination. To take the PHY impedance out of this equation as much as possible, all resistances are removed from the termination of the PHY at computer turn on. The voltage seen at the input port is approximately the input current multiplied by the CAT5 resistance in this

case. This voltage is now compared digitally to the voltage across the reference resistance, and this reference resistance is adjusted until the two voltages match each other. The resistance this voltage corresponds to is then mirrored to the PHY termination resistance, stored and used for calibration. The following schematic shows the setup during the matching of the reference to the CAT5 cable. The iterative resistance matching sequence followed is described in detail in the following section.

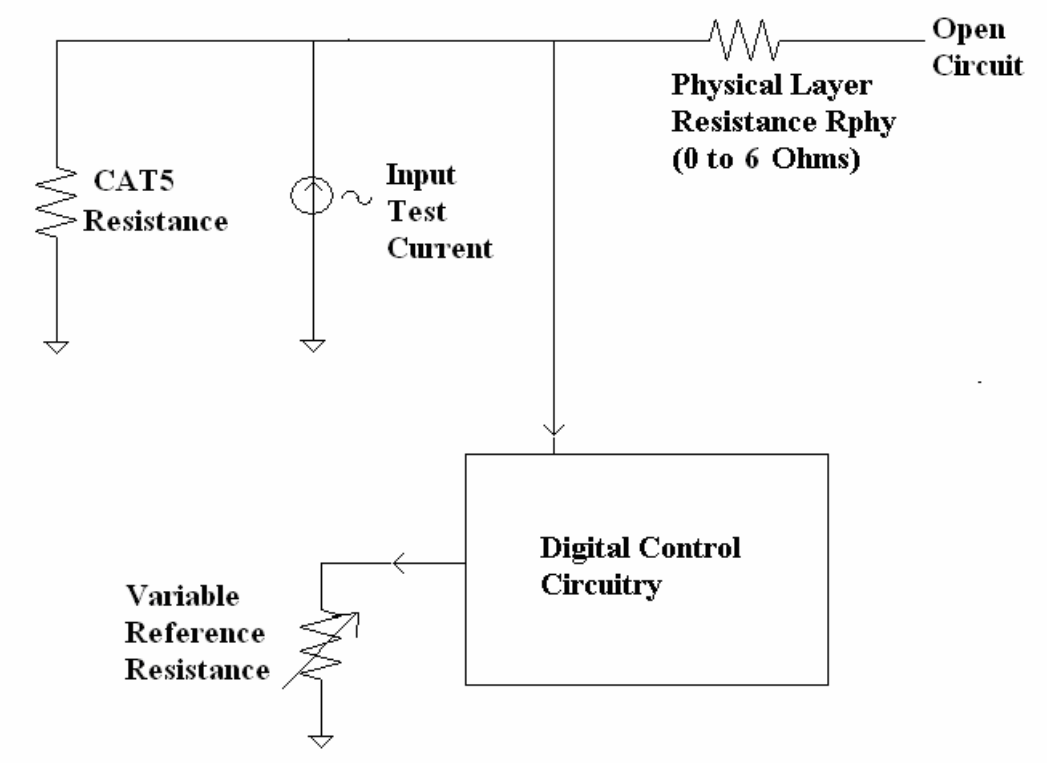

Figure 34: CAT5 Reference Impedance Generation

According to [6], if a test current is applied from a source in parallel with an impedance as shown in figure 34, the voltage seen across that impedance can be used as a measure of the S11 return loss. This technique is employed in the impedance matching of the CAT5 and the reference. The newly created reference voltage could be mapped in the digital

domain to the required resistance needed in series with the PHY. While this would require some processing, it would be the easier and more efficient way of matching the PHY to the CAT5 cable.

From the characterized return loss data, it is known that the maximum resistance due to the PHY itself is 6 Ohms. Thus, the PHY resistance bank needs to add a minimum of 44 Ohms to match an ideal CAT5 cable. If the newly created reference voltage and V_{TEST} are not mapped directly and instead are compared in the analog domain, the impedance of the CAT5 cable cannot be disconnected from the circuit. When they are compared the CAT5 resistance will be in parallel as shown in figure 33. This introduces some mismatch in the PHY voltage measurement, and the worst case of the 6 Ohm PHY in parallel with the CAT5 would be when the CAT5 impedance is the closest to the PHY, or 42.5 Ohms according to the specification. The resistance seen at V_{TEST} in this case would be the parallel combination of 42.5 and 6, which is 5.3 Ohms. In this extreme case there is a maximum mismatch of 12% in the impedance. In the majority of cases however, the average CAT5 impedance is 50 Ohms and the PHY resistance is less than or equal to 1 Ohm, and thus the average mismatch from the parallel combination of the PHY and the CAT5 impedance (50||1 Ohms) is less than 2%. If the comparison is made in the digital instead of in the analog domain and the reference voltage is directly mapped to the resistance needed in series with the PHY, this problem could be averted. The following figure shows the total proposed PHY resistance matching circuit.

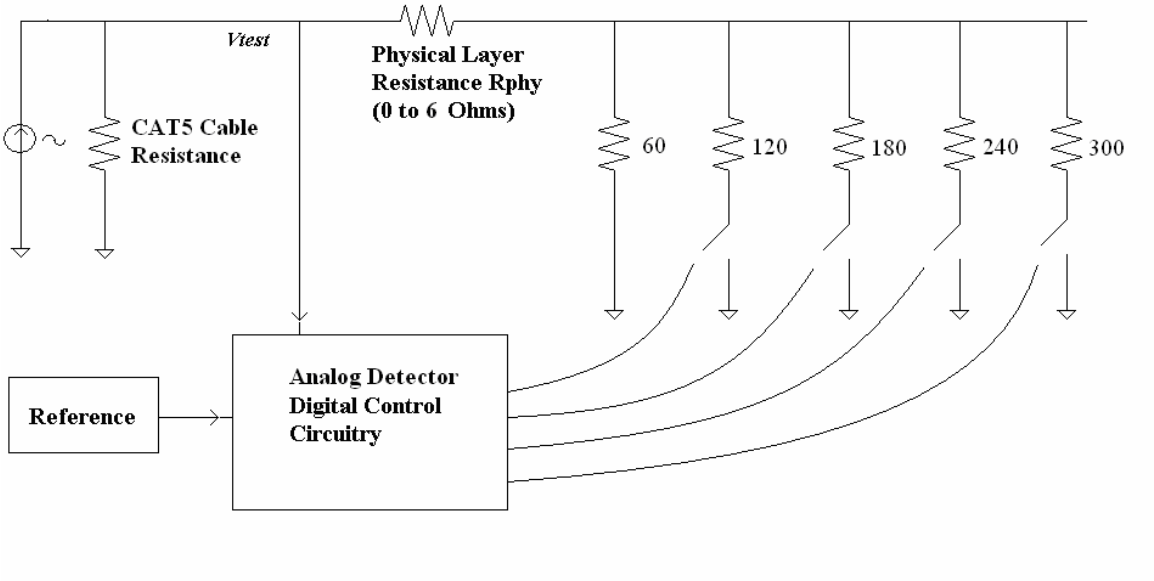

Figure 35: Low Frequency Resistance Matching Network

In this figure, the PHY resistance bank is initially set up with all of the switches open except for the 120 Ohm switch which is in parallel with the 60 Ohm resistor. This gives a total resistance of 40 Ohms for the filter. Because the scattering parameters cannot be easily measured in real time on board, the voltage across the PHY is used as a measure of the return loss. With the first iteration of termination resistors ready, the PHY voltage is then compared to the reference. If the voltage at V_{TEST} is greater than that across the reference load, then the 120 Ohm switch is opened, the 180 Ohm switch is closed and the process is repeated. In this second iteration, the 60 Ohm resistor is in parallel with the 180 Ohm resistor, giving a total PHY termination resistance of 45 Ohms (plus the nominal R_{PHY} as according to the above figure). Through these digital iterations, the resistance can be matched to the reference, and thus to the CAT5 cable. The resistance values possible by closing one switch at a time are 40, 45, 48 and 50 Ohms, but more resolution is possible by combining the switches.

3.2.2 Imaginary Impedance Matching

The second part of the impedance matching is the reactance. As stated earlier in this chapter, this is done at higher frequency where the poles and zeros of the PHY are located. Because there is a bandwidth of data transmission with a Gaussian distribution around 70 MHz, the reactance should be matched for not one frequency but the entire band. However, because of the poles and zeros of the PHY network, a filter solution will have minimal effect after the 3dB bandwidth. With these two pieces of information in mind, it will be necessary to sweep not only components but frequency as well. The IEEE 802.3 return loss specification shows an elbow at 40 MHz and it is known from the return loss data that the 3dB bandwidth is around 70 MHz. Directly from this specification and the empirical results, a frequency sweep bandwidth is created that will maximize the component matching in the region that the poles and zeros play the greatest role. This is performed from 20 MHz to 60 MHz in steps of 20 MHz to cover the majority of the high frequency bandwidth without overly complicating the system. The complex Laplace impedance of the PHY model in equation 7 showed that there are 3 zeros and 2 poles in the circuit. To have the greatest chance to balance the reactance, at least 1 pole must be added. This can most easily be accomplished by adding a shunt capacitor at the input to the PHY. In parallel with the entire PHY circuit, this capacitor has great leverage to change the shape of the reactance. At the same time, if a shunt capacitor is added at the termination of the PHY, the circuit becomes more symmetrical for receive or transmit signals. This end capacitor also adds an extra control variable to the circuit. Intuitively, adding additional shunt components at the termination of the PHY will not be as effective as adding them at the input because they will be recursively

in parallel with every other shunt component in the circuit. With this setup there is a component that can change the impedance of the circuit at the input and termination of the PHY. Many other circuit configurations were experimented with, however this arrangement proved to give the greatest leverage with the impedance while keeping the circuit as simple as possible. Adding a series variable inductor would also be a theoretical solution; however this would require dramatic reconfiguration of the board and would be adding direct injection of noise in the signal path. On the other hand, adding a shunt inductor would not help balance the pole/zero ratio. The following equations show the Laplace impedance of the new circuit.

$$\frac{[L_2L_4C_3RC_5]s^4+[L_2L_4C_3]s^3+[L_4RC_5+L_2RC_3+L_2RC_5]s^2+[L_2+L_4]s+R}{[L_2L_4RC_3C_5C_1]s^5+[L_2L_4C_3C_1]s^4+[C_1RC_5+L_4CRC_5+L_2RC_3C_1+L_2RC_5C_1]s^3+[L_4C_3+L_2C_1+L_4C_1]s^2+[RC_3+RC_1+RC_5]s+1} \quad (9)$$

$$\frac{\omega^4[L_2L_4C_3RC_5]-j\omega^3[L_2L_4C_3]-\omega^2[L_4RC_5+L_2RC_3+L_2RC_5]+j\omega[L_2+L_4]+R}{j\omega^5[L_2L_4RC_3C_5C_1]+\omega^4[L_2L_4C_3C_1]-j\omega^3[C_1RC_5+L_4CRC_5+L_2RC_3C_1+L_2RC_5C_1]-\omega^2[L_4C_3+L_2C_1+L_4C_1]+j\omega[RC_3+RC_1+RC_5]+1} \quad (10)$$

In the above equation, the naming convention follows that of figure 36 shown here.

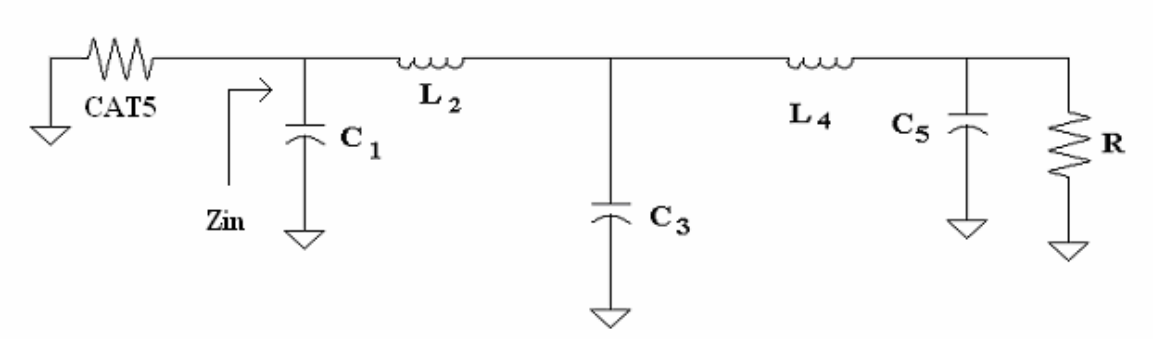

Figure 36: PHY Model with Additional Capacitor Terminations

In equations 9 and 10, C1 and C5 appear separately in most of the coefficients. In general, the RLC network of figure 36 will be unbalanced due to these parasitics and the physical makeup of the circuit components used to model PHY. The objective should be to find the values of C1 and C5 that reduce the S11, which can be accomplished by converting figure 36 into a ladder network.

In [8], a ladder network is shown to be comprised of LC blocks with proper termination. Used as a filter, it is very robust because of its ease and variety of implementation, such as Butterworth and Chebyshev. Both of these filters can accomplish the optimization of PHY impedance; however there is a tradeoff in their operation. While the Chebyshev will have ripple in the passband, the bandwidth of the filter can be increased. Conversely, the Butterworth will have a maximally flat passband, but the bandwidth will be smaller than a comparable Chebyshev. This tradeoff can only be answered by the specification of how much in band ripple can be tolerated, and so both filter implementations are shown in the simulation and experimental results. One of the advantages of the robust ladder filter implementation is that the resulting filter allows for slight deviation in the values of the components. This means that a wide range of PHY component values can be tolerated by allowing for different values of C1 and C5 capacitance. The key to this filter is that it is properly terminated. This means that the CAT5 resistance and the PHY termination resistor must be matched as closely as possible to ensure the best conditions for the LC ladder to work properly. In a properly terminated LC network, the reactive components will resonate at a certain frequency. In this ideal condition, the LC components add no additional impedance to the circuit, thus sending the entire signal to the termination. However, due to parasitics and the resistance in the

inductor, there is some minimal loss even in resonance. In the ladder network, this creates a passband in which the LC resonance minimizes the S11 return loss and maximizes the S21 throughput signal. This is one of the main properties of the ladder network. By compensating the PHY model with the correct C1 and C5 capacitors, the PHY can be optimized so that the passband of the network more accurately fits the IEEE Gigabit Ethernet bandwidth specification. The following Butterworth filter tables (normalized to 1 rad/sec bandwidth) are adapted from [8]. It is assumed that a normalized termination resistance of 1 Ohm is used at the input and output of the network.

Table 8: Butterworth Ladder Filter Component Values

	C1	L2	C3	L4	C5	L6	C7	L8	C9	L10
2	1.4142	1.4142								
3	1.0000	2.0000	1.0000							
4	0.7654	1.8478	1.8478	0.7654						
5	0.6180	1.6180	2.0000	1.6180	0.6180					
6	0.5176	1.4142	1.9319	1.9319	1.4142	0.5176				
7	0.4450	1.2470	1.8019	2.0000	1.8019	1.2470	0.4450			
8	0.3902	1.1111	1.6629	1.9616	1.9616	1.6629	1.1111	0.3902		
9	0.3473	1.0000	1.5321	1.8794	2.0000	1.8794	1.5321	1.0000	0.3473	
10	0.3129	0.9080	1.4142	1.7820	1.9754	1.9754	1.7820	1.4142	0.9080	0.3129

It can be seen from the table that the 5^{th} order filter implementation applies to the PHY model and derived filter in figure 36. In the Butterworth and Chebyshev realizations of the PHY model, it is assumed that the model is symmetrical. To accurately represent the possible worst case PHY impedances, L4 is at least equal to or less than L2. This accounts for all parasitics present that unbalance the PHY. The following table shows the component ratios needed to implement the Chebyshev ladder filter. Once again, it is assumed that a 5^{th} order filter is used.

Table 9: Chebyshev Ladder Filter Component Values

	C1	L2	C3	L4	C5	L6	C7
3	1.5963	1.0967	1.5963				
5	1.7058	1.2296	2.5408	1.2296	1.7058		
7	1.7373	1.2582	2.6383	1.3443	2.6383	1.2582	1.7373

In both cases, the C1 added filter capacitor plays the major role in balancing the network. The main idea seen from these tables is that the ratio of C1 and C5 to that of the C3 capacitor defines the ladder network relationship. By compensating these two capacitors and balancing the network, resonance is achieved over the passaband.

If a more complex PHY model was created, or if a more accurate ladder filter was needed, the above table could be used to add additional RLC blocks. A solution with less error, larger bandwidth and wider range of application could be designed by extending this ladder network line of thought at the expense of complexity, area, settling time and cost. The key to remember in any implementation and model is that the component ratios define the compensation needed and the overall performance of the ladder. The next two sections, the Cadence simulations and the experimental results, show that this filter implementation does indeed improve the Ethernet return loss and bandwidth of the PHY at least 10 dB and 20% respectively.

The next step is to derive the possible range of component values for the PHY model so that simulations and experiments can be performed to verify the ladder network theory. From the return loss data, the range of the PHY capacitor C3 is from 25pF to 55pF. For the added shunt filter capacitors, C1 and C5, the range used is from 0pF to 40pF, in increments of 4pF. These could be implemented with a varactor or a capacitive multiplier. This allows the circuit to address a wide range of possible PHY capacitances

without overcomplicating the decision of what capacitance to add. The ease of implementation allowed with capacitive increments of 4pF comes at the cost of inaccurate capacitance up to LSB/2 (Least Significant Bit), or 2pF, being added to the circuit. This error is tolerable because the impetus of this project is a proof of concept to the ladder filter solution for the PHY impedance matching problem. For a full scale IC solution, smaller capacitor increments could be used for more accurate matching results. While this takes care of the passband, it has little effect after the corner frequency of the filter. To address this problem, the bandwidth would have to be extended. If some ripple can be tolerated, the Chebyshev filter implementation could be used to increase the bandwidth and decrease the return loss even further. Both of these filter types are shown in the following experimental sections. Figure 38 shows the schematic of the entire filter implementation and matching control circuitry. The PHY component and ladder network capacitor ranges are also listed in the following table.

Table 10: PHY Ladder Model Component Ranges

Component	Minimum	Median	Maximum
C1 [pF]	0	20	40
L2 [nH]	50	100	100
C3 [pF]	25	45	50
L4 [nH]	50	75	100
C5 [pF]	0	20	40
R (Ohms)	44	47	NA

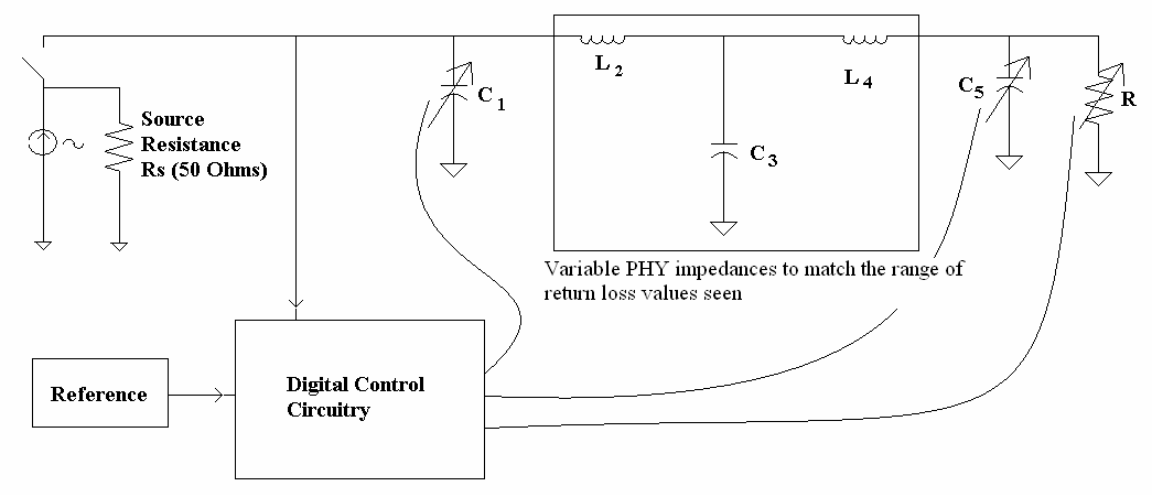

Figure 37: Ethernet Impedance Matching Network with Digital Control Circuitry

3.3 Control Circuitry

This section will show a possible implementation of the control circuitry of the PHY matching. The following figure shows the proposed block schematic of the entire matching circuit.

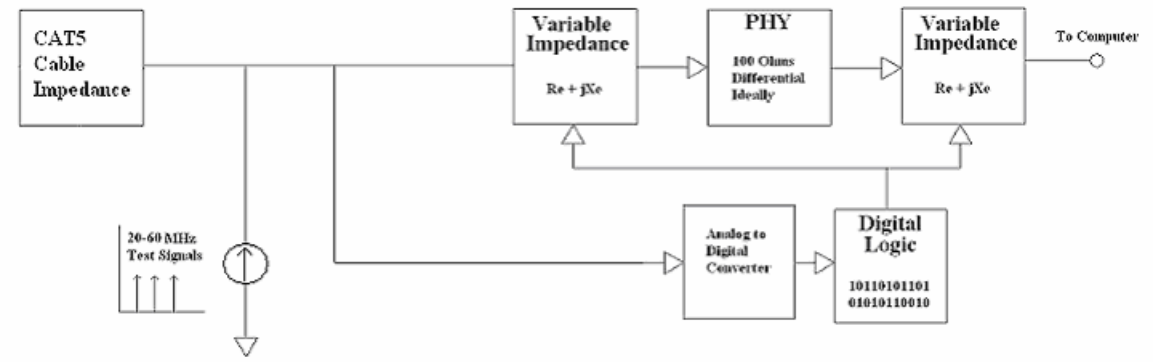

Figure 38: Block Diagram Circuit Schematic

From the diagram, the ADC and the digital logic will be covered in more detail in the following sections. Because the application of this system is for Ethernet communication, the circuitry already on the computer's motherboard includes ADCs and a

microprocessor capable of the conversion and logic needed. The biggest challenge with using already present circuitry would be routing the traces to and from the chips. For the sake of discussion, it is assumed that the control circuitry would not take advantage of any computer hardware and would be implemented by the impedance matching system.

One possible implementation of the ADC and digital comparison circuitry is with an algorithmic ADC as in [9]. In this approach, the voltage signal at the input of the PHY would be buffered and then converted to a DC signal proportional to the root mean square voltage of the signal. This signal is then input to the algorithmic DC-ADC. Such a converter is described in detail in [10]. A block diagram of this approach is presented here.

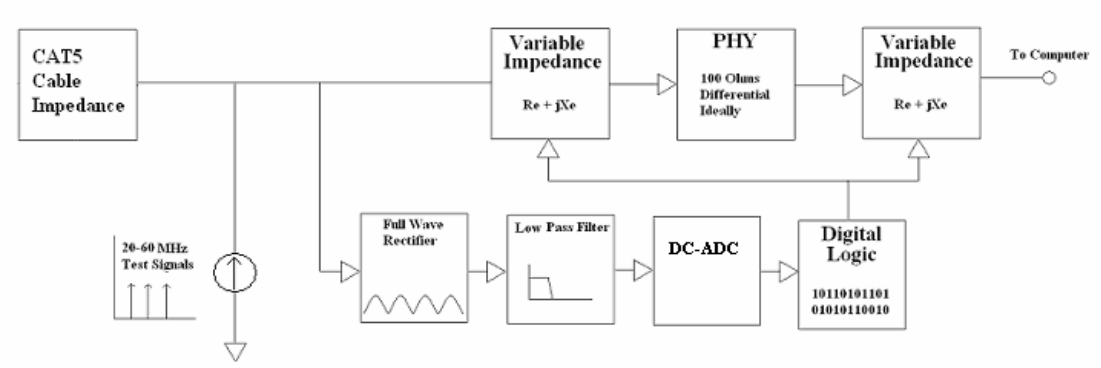

Figure 39: DC-ADC Implementation Block Diagram

In this solution, the ADC would need to convert DC signals not 60 MHz signals, which would greatly reduce the need for high power, high bandwidth performance such as those shown in [11]. While this does ease the required specifications of the ADC, it would introduce more blocks into the signal path. The rectifier, the filter and the ADC would all have an error associated with them and the total error would have to be taken into account which could be large relative to the signal magnitude. In the RF ADC implementation,

the performance specifications are increased but the error is introduced in fewer blocks, which would make it easier to control from a design standpoint, such as in [12]. There are multiple vendors that provide ICs capable of the required RF implementation performance, and a few ICs are listed here [13]. With this in mind, the specifications required of an RF ADC are discussed.

3.3.1 Analog to Digital Converter

The analog signal at the input of the PHY is ready to be converted to digital information for comparison. This can be accomplished by an analog to digital converter, or ADC. The performance specifications depend on the architecture chosen, but a brief discussion of the general considerations is now presented. From [14], the first specification to consider is the resolution, which can be defined by the input range required and the maximum error tolerable. The largest error introduced in the voltage measurement at the input of the PHY is due to the parallel combination of the CAT5, which was found to be around 2% as stated earlier. Thus, the ADC should introduce no more than this amount of error so that it is not a major contributor. The following formula is used to determine the error associated with a certain resolution, N, in bits.

$$Error(\%) = \frac{100}{2^N} \qquad (11)$$

For an error smaller than 2%, this equation shows that the resolution must be at least 7 bits (2^7=128). The range in this case is based on the limitations of the reference generator. The schematic of the circuit under question is shown here.

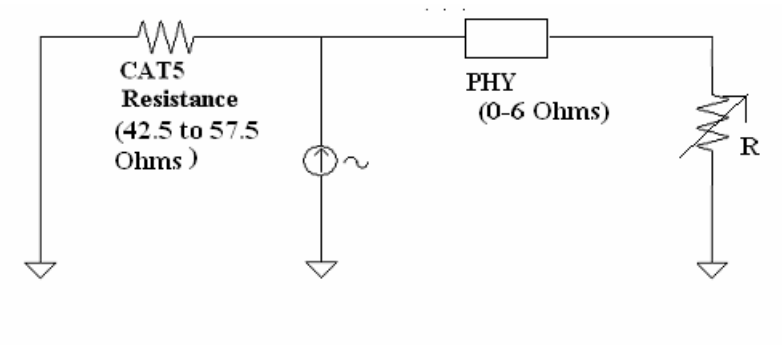

Figure 40: Resistance Seen at the Input Reference Source

This current source has to drive the PHY in parallel with the CAT5 impedance, which averages to about 25 Ohms. In the worst cases, the PHY impedance goes as low as 35 Ohms at 60 MHz, which would give a parallel impedance of 20 Ohms. On the other end of the spectrum, the highest impedance seen by the input source would be when the PHY is very large, in which the worst case CAT5 cable would be seen, or 58 Ohms. It becomes harder to generate large signal amplitudes at 60 MHz, which is the largest input required, so the current input signal should be kept to a reasonable value. If this value is set to no greater than 20 mA, which is a large but bearable value for a CMOS transistor, this gives an input range of 400 mV to 1.16 V with an average around 500 mV. For 7 bit resolution over a 1 V input range the circuit has a LSB of 7.8 mV. With the 20 mA current input, a 1 Ohm difference between PHY and CAT5 impedance would be seen as a 20 mV difference in their voltages which is 2.5 times greater than the LSB. Thus a 7 bit ADC achieves the resolution specification. From [13], there are multiple ICs that have this performance.

3.3.2 Source Generation and Digital Control Algorithm

The control algorithm would only need to be run once at the start up of the computer. This means that there would be a calibration for each CAT5 cable plugged in to the computer every time it is turned on, which makes the system very robust. It also means that the system could take seconds to calibrate without loss of performance seen by the user.

Because the system is an Ethernet impedance matching system, the computer processor generates signals that could be used for the reference. This is possible with the on motherboard oscillator, and a simple MOS gate could be used to transform a voltage signal into current. After conversion of the PHY voltage signal to a digital signal, the signal processing is performed. The first step is to compare the PHY peak voltage to that of the reference voltage with the first resistance configuration as described in the previous sections. The difference between the two is stored as 't0.' The next matching resistor is then connected, as described in the resistance matching section, and the PHY and reference voltages are compared and the difference is stored as 't1'. 't1' and 't0' are then compared against each other and the smaller of the two is stored as 't0'. This process is repeated until all resistor combinations have been compared to the lowest differential value. The control circuit then attaches the resistor associated to 't0' to the circuit and this value is kept for the duration of system operation. This algorithm is then repeated for the capacitors at higher frequencies. However, for the reactance match, the capacitor added that produces the smallest PHY and reference impedance difference at a given test frequency is added to a value called 't_total'. This value is then divided by the number of frequencies tested, 20 MHz to 60 MHz in 20 MHz steps, and the average

capacitance needed over the frequency range is then found. The closest capacitance value implementable is then used. The following table examples illustrate this process.

Table 11: Capacitance Bandwidth Averaging Example

	20 MHz	40 MHz	60 MHz	Capacitor Average over Bandwidth	Capacitor Value Implemented
Case 1 Matching Capacitor (pF)	12	12	16	13.3 pF	12 pF
Case 2 Matching Capacitor (pF)	4	8	12	8 pF	8 pF
Case 3 Matching Capacitor (pF)	8	8	16	10.6 pF	12 pF

A flow chart is presented below to illustrate this logic.

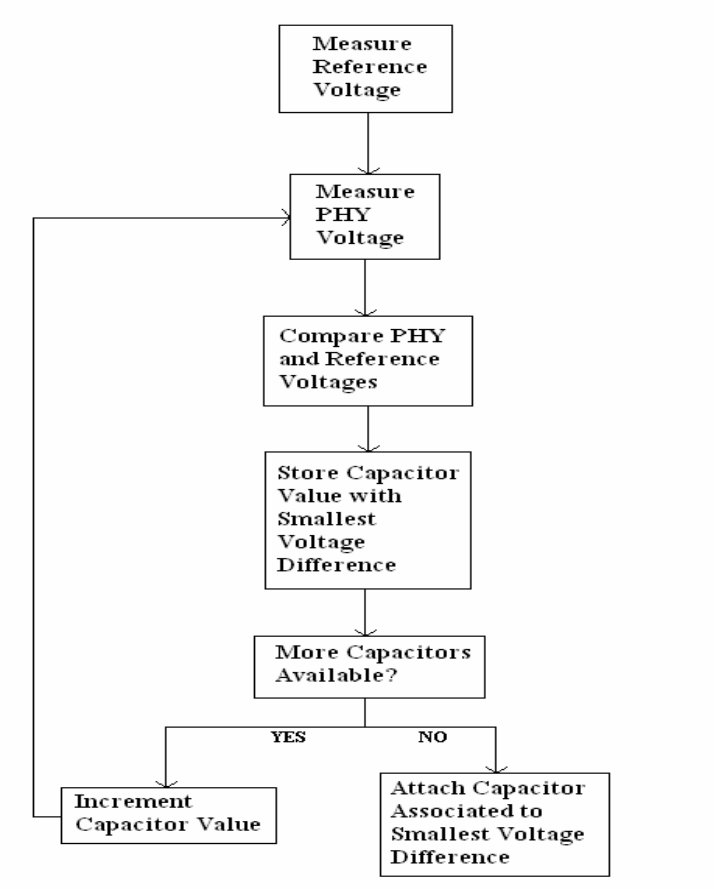

Figure 41: Digital Logic Flow Diagram

CHAPTER IV

CADENCE SIMULATION RESULTS

4.1 Physical Layer Model Creation

Using the return loss data and [15], a model for the PHY was created with inductor, resistor and capacitor value ranges. The derivation of these values was shown in the previous section. The following schematic shows these experimentally derived values. The input is a current source, with 50 Ohm source resistance to model the CAT5 cable impedance. A 5mA amplitude sine wave is used as the input.

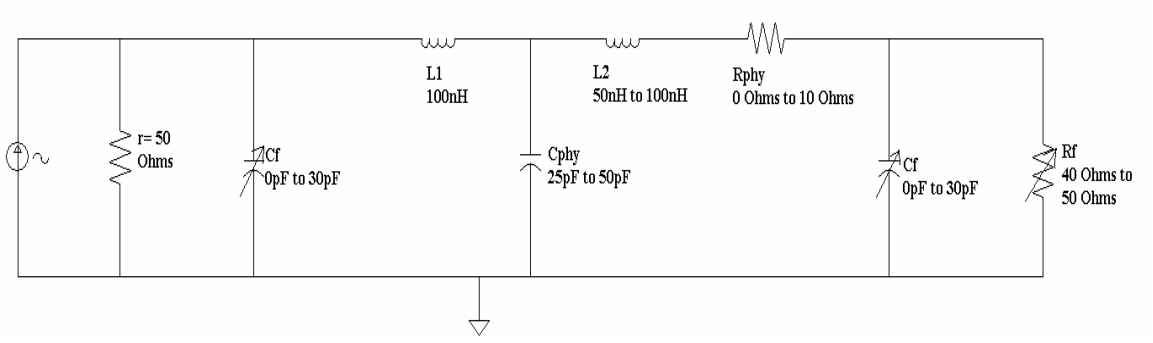

Figure 42: PHY Component Value Ranges

The first step in matching the PHY impedance to the CAT5 cable is to match the resistive component. As pointed out in the previous section, this is done at low frequency so that the reactance does not play a major role. For the return loss simulation plots, a 50 Ohm port with an input of 1 mA amplitude at 1 MHz was used for testing purposes. The filter resistance Rf is swept from 40 to 50 Ohms, and the voltage is measured and compared to a reference with the same mirrored input.

The PHY model also includes the Butterworth impedance matching filter used in the project. The ranges for these values are derived from the PHY values and the tables in [8]. Cadence was then used to simulate this circuit. By measuring the voltage at the input of the PHY, the return loss can be approximated. The following figures show the schematic of the circuit under test and the PHY voltage and the return loss of the circuit with swept filter capacitors Cf for Rf=50 Ohms, L1= 100nH, L2=75nH and Cphy=45pF. These values are in the center of their expected ranges. In this case, both capacitors Cf are swept together.

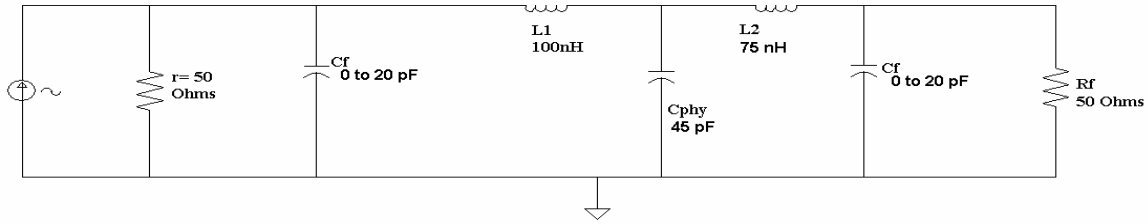

Figure 43: PHY Schematic for Test 1

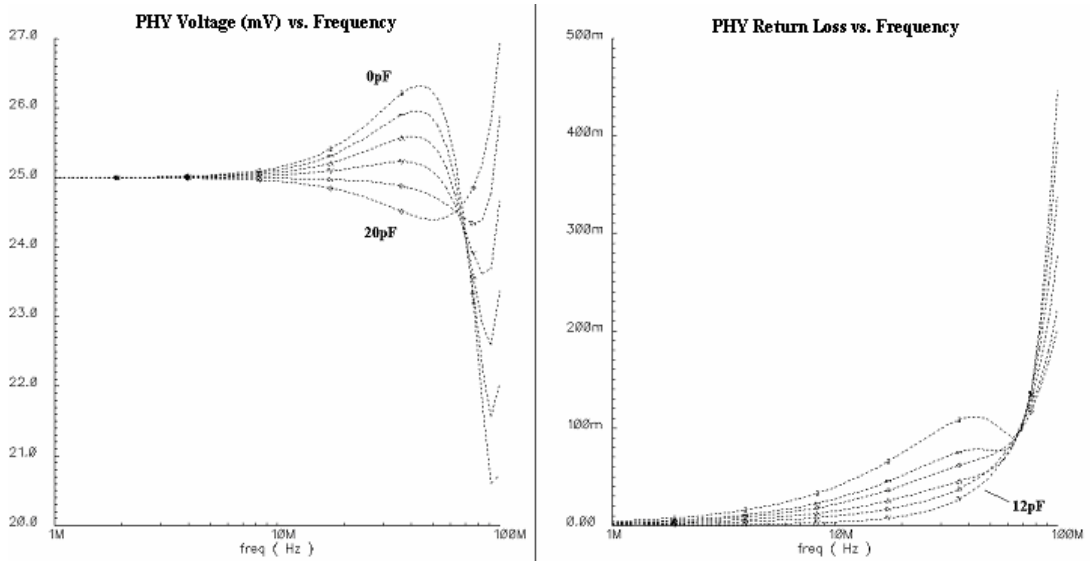

Figure 44: PHY Voltage and Return Loss with Filter Caps Swept from 0 to 20pF

From this figure it can be seen that by adding 12pF capacitors to the front and back of the PHY, the return loss is minimized by as much as .07, or -8.8 dB at 40 MHz. At the same time, the voltage of the PHY is closest to 25mV. Thus, if the voltage of the PHY is measured and compared to the voltage across an ideal 50 Ohm reference with the same source, an ideally matched impedance can be found. This is possible with iterations of the measurement and digital control of the loop. It can also be seen from the above plot that all of the filter implementations have the same passband. Thus, to measure the voltage, and the return loss, frequencies within the passband should be used. To reduce the chance of inconsistent or inaccurate results, a frequency sweep of the input source is used and the best result is taken. A sweep from 20MHz to 60MHz, with 20MHz increments is used. The impedance closest to 25 Ohms is taken, and the results are averaged digitally. In addition, the following plot shows that by unbalancing the capacitor at the front, C1, and back, C5, of the PHY, even better results can be obtained at the expense of ripple in the passband. The values for the PHY are L1= 100nH, L2= 75nH and Cphy= 45pF as corresponding to the circuit schematic in figure 44.

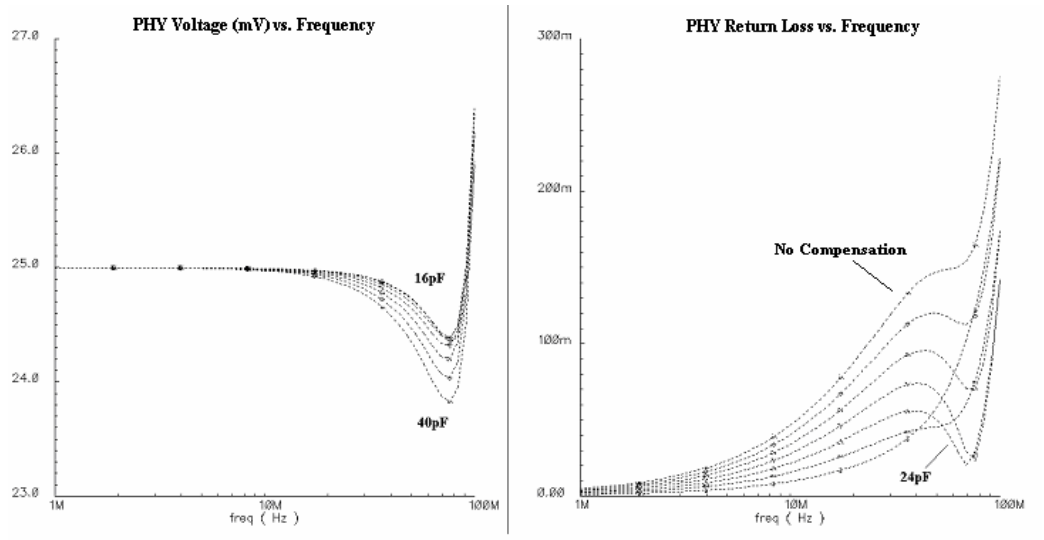

Figure 45: PHY Voltage and Return Loss with End Cap at 12pF and Swept Front Cap from 16 to 40pF

Using different capacitor values at the front and end of the PHY can compensate impedances that could not be compensated with balanced capacitors. It also adds complexity to the design implementation. If the capacitors are swept together, the total number of iterations is equal to the number of capacitor increments, or N. However, with unbalanced capacitors, the total number of iterations increases to N^2, unless some sweep algorithm can be derived to decrease this number. Luckily, the ladder network tables in [8] show that the front capacitor should be greater than or equal to the back capacitor. By first sweeping the capacitors together, the back capacitor C5 is set to the best value. Then, the front capacitor C1 can be swept independently, with the minimum value of the sweep set to the back capacitor C5. This algorithmic approach would also take care of the parasitics found at the PHY terminations as well. The appendix includes more filter implementation sweeps for the range of PHY component values.

The following plot shows that as the input impedance of the PHY is tuned closer to the CAT5 cable impedance of 50 Ohms, the power delivered to the output of the

physical layer gets larger. On each plot, the top line is the voltage across the PHY, which is a measure of the S11 return loss. The line on the bottom is the voltage at the output of the PHY, which is the transmitted signal. Assuming that there is some loss from the PHY due to parasitics and resistance of the line, the PHY output voltage should be as large as possible. The graph on the left is without compensation, the graph in the middle is with balanced 16 pF compensation capacitors and the graph on the right is with unbalanced 24 pF and 12 pF capacitors at the front and back of the PHY respectively. It can be seen that when the S11 return loss is minimized with compensation, the transmission is increased as well especially at high frequencies.

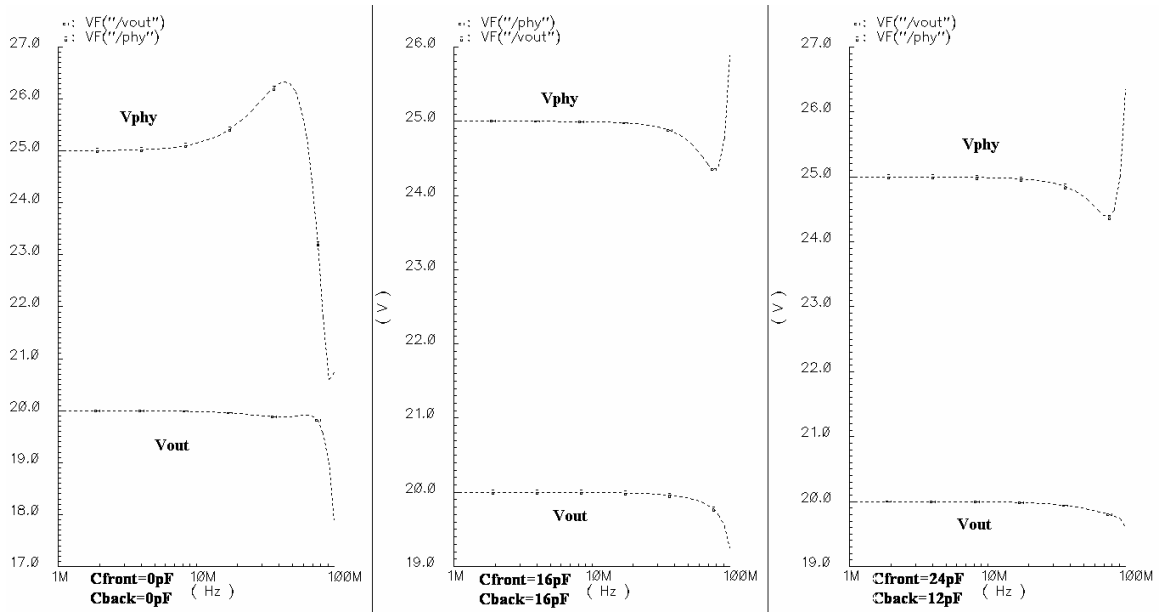

Figure 46: PHY and Output Voltages for Different Capacitor Compensation Values

The following plots show the S21 transmission data and the 3dB Bandwidth of the PHY with different filter implementations. The transmission loss at 100 MHz for the first, second and third capacitor arrangements is -2.9 dB, -2.3 dB and -2.1 dB respectively.

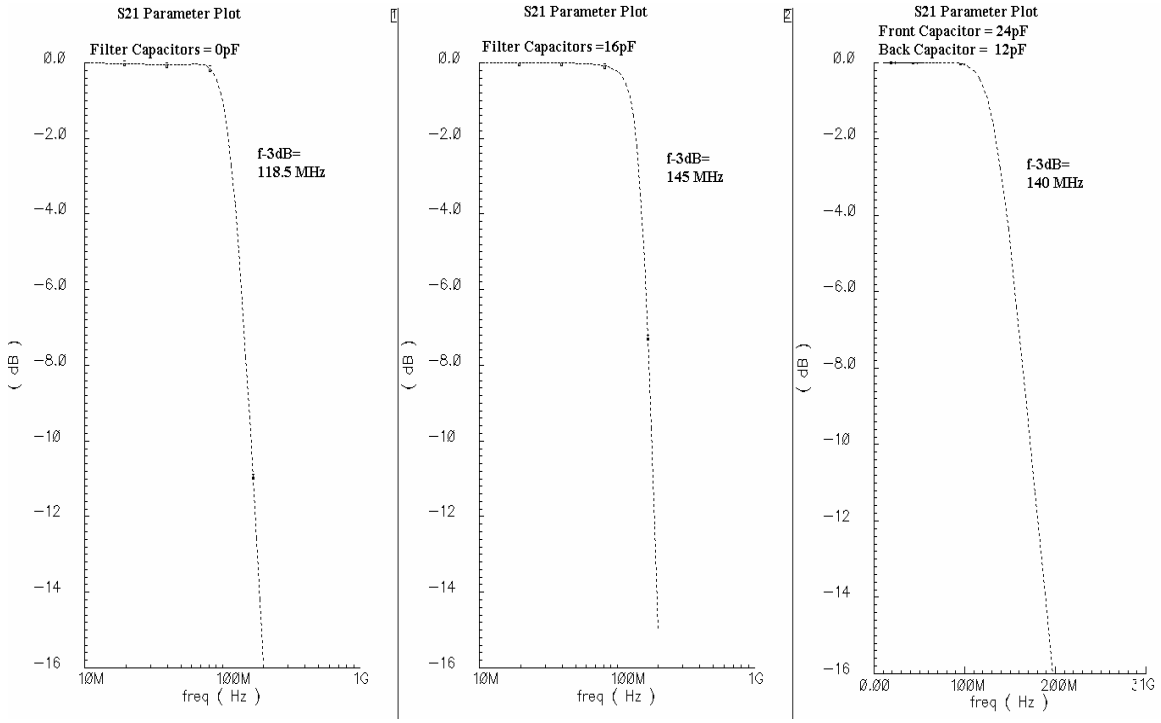

Figure 47: S21 Transmission for Three Filter Types

From the previous two plots it can be seen that with the filter implementation in place to reduce the S11 return loss, there are other benefits that appear as well. The first of these is the increase in the S21 transmission. With greater signal delivered to the output of the PHY, the specifications for the input of the next circuit can be reduced. Also, it is seen that the bandwidth of the circuit increases as well. This would increase the possible data transmission frequency of the system. More plots are included in the Appendix to show the range of PHY impedances that can be normalized with the filter implementation.

CHAPTER V

EXPERIMENTAL RESULTS

5.1 Circuit Model Definition

For the purposes of testing the ladder filter solution to the return loss problem, a discrete component circuit was implemented. Discrete components can accurately replicate the PHY impedance, and thus a discrete model can be used for experimental results. The circuit was scaled down from 100MHz to 1MHz bandwidth to accommodate the tolerances of available discrete inductors and capacitors. Because the capacitance values of the PHY are so small, parasitics from the discrete components become comparable to the parts themselves especially at the higher bandwidths. Other sources of error are those associated with the breadboard and the interconnects used in the experiment. These error sources are what create the small ringing seen in the eye diagrams presented in this chapter. If a PCB was created to test the design, these errors could be further reduced. With the lower frequency range used, the adverse effects of the ground plane and other parasitics associated with the breadboard are reduced. All of the parts used had at least ±3% tolerance. The following schematic shows the PHY values used in the breadboard experimental results.

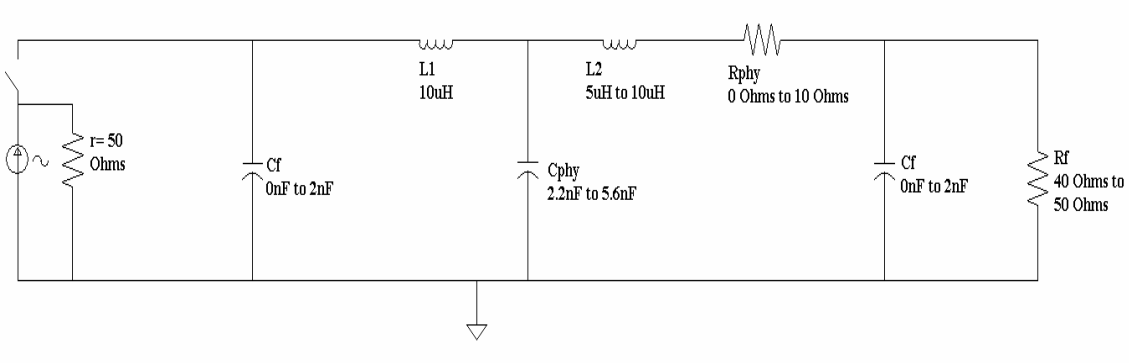

Figure 48: Physical Layer and Ladder Filter Experimental Values

The following figure shows the breadboard test setup of the PHY circuit.

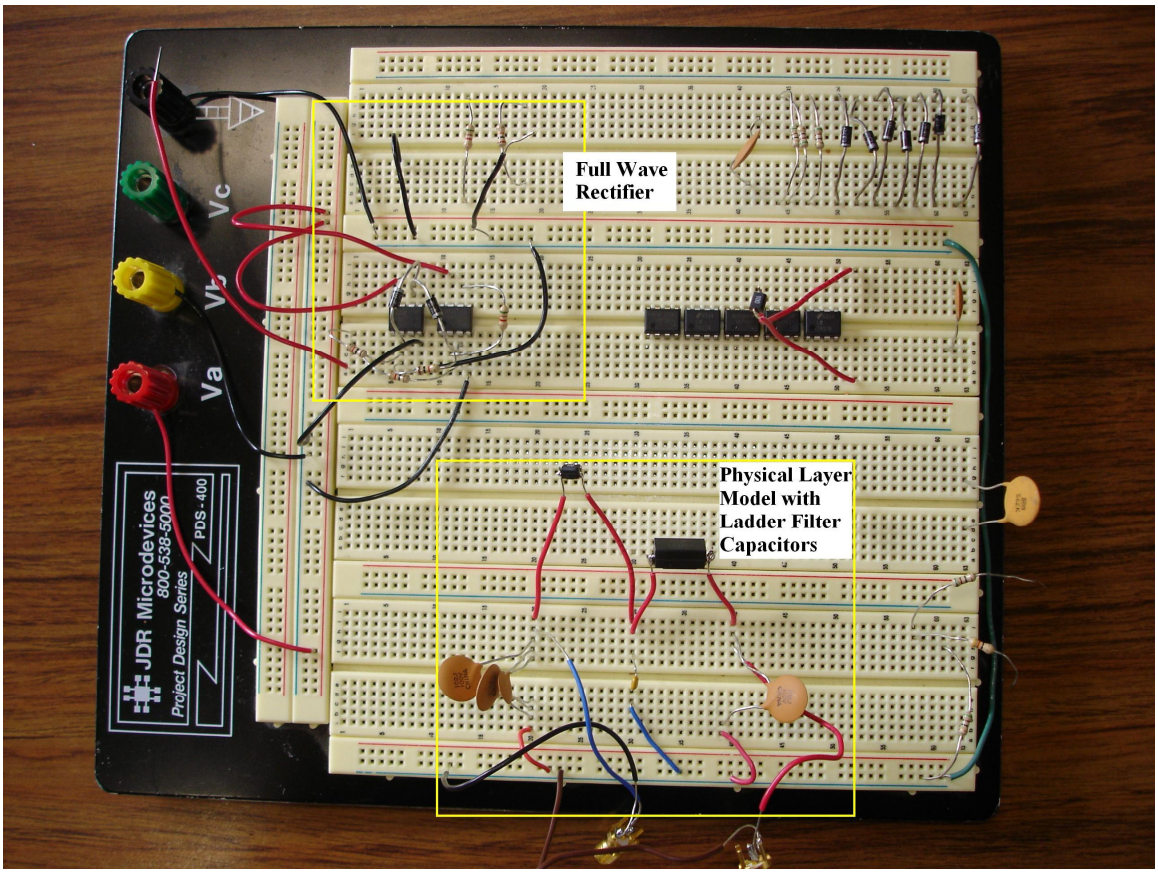

Figure 49: Breadboard Experimental Setup

Multiple tests were performed to see the effect of the matching circuit on the PHY impedance. The first was a scattering parameter test to measure S11 and S21 and the second was an eye diagram of the signal.

5.2 Scattering Parameter Results

The above schematic was created on a bread board with six different PHY configurations. The voltages at the input and output of the PHY were measured, and the S11 return loss and S21 transmission were found with the Agilent 4395A and the Agilent 87511A Network Analyzers. These instruments operate from 100 kHz up to 500 MHz. The first step was to set up and calibrate the analyzers, and this process can be found in [16]. An HP 85052C 3.5mm calibration kit was used for this purpose. The return loss was measured from 100 kHz to 1 MHz and the magnitude in decibels was plotted. The following images show the improvement in return loss found by employing the ladder equalized network to the PHY impedance.

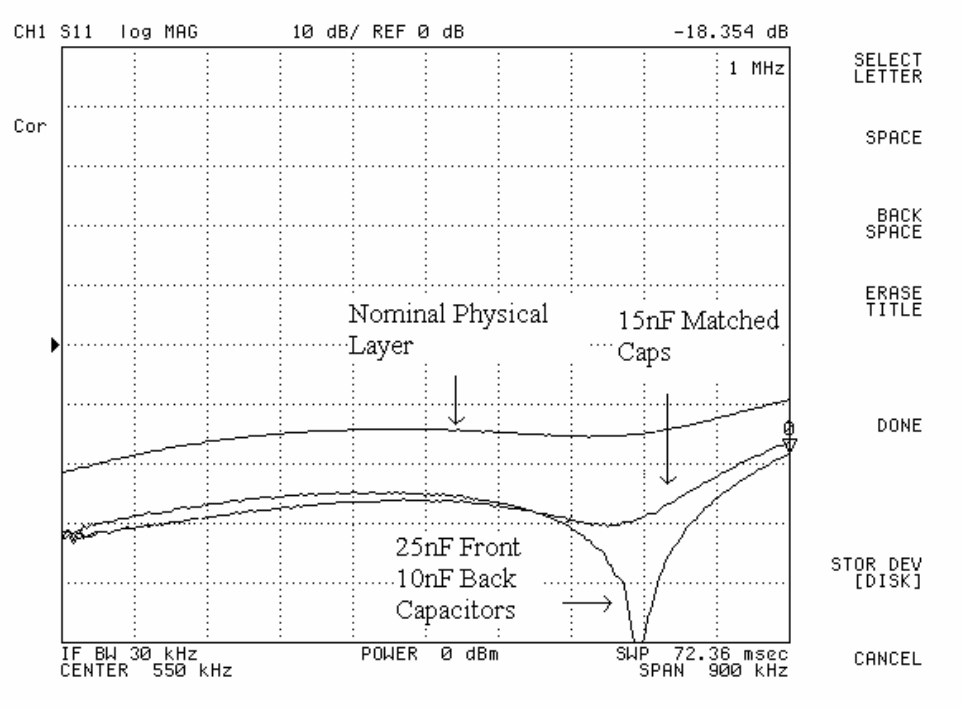

Figure 50: S11 Magnitude Return Loss for PHY, Butterworth Filter and Chebyshev Filter

There are two distinct solutions shown on the plot, the Butterworth filter with 15 nF matched capacitors and the Chebyshev filter with 25 nF and 10 nF capacitors. If some return loss ripple in the bandwidth can be tolerated, then the improvement seen by using the proposed approach is apparent. The unbalanced filter capacitor approach (Chebyshev) adds complexity and increases the time to lock but achieves the best performance. Not only is the return loss across the passband reduced, the bandwidth of the filter increases as well. However, if no ripple is allowable, then the Butterworth filter with matched capacitors still provides greatly reduced return loss in the passband. The other main advantage of this implementation is that the circuit and the digital logic are not complex, and the time to match the impedance is relatively less than the Chebyshev filter.

The next step was to measure the S21 transmission. This was done by connecting the RF output of the Network Analyzer to a power splitter, and measuring the total power

to the power received at the end of the PHY. The following plot shows the S21 curves for the balanced and unbalanced PHY from 100 kHz to 2MHz.

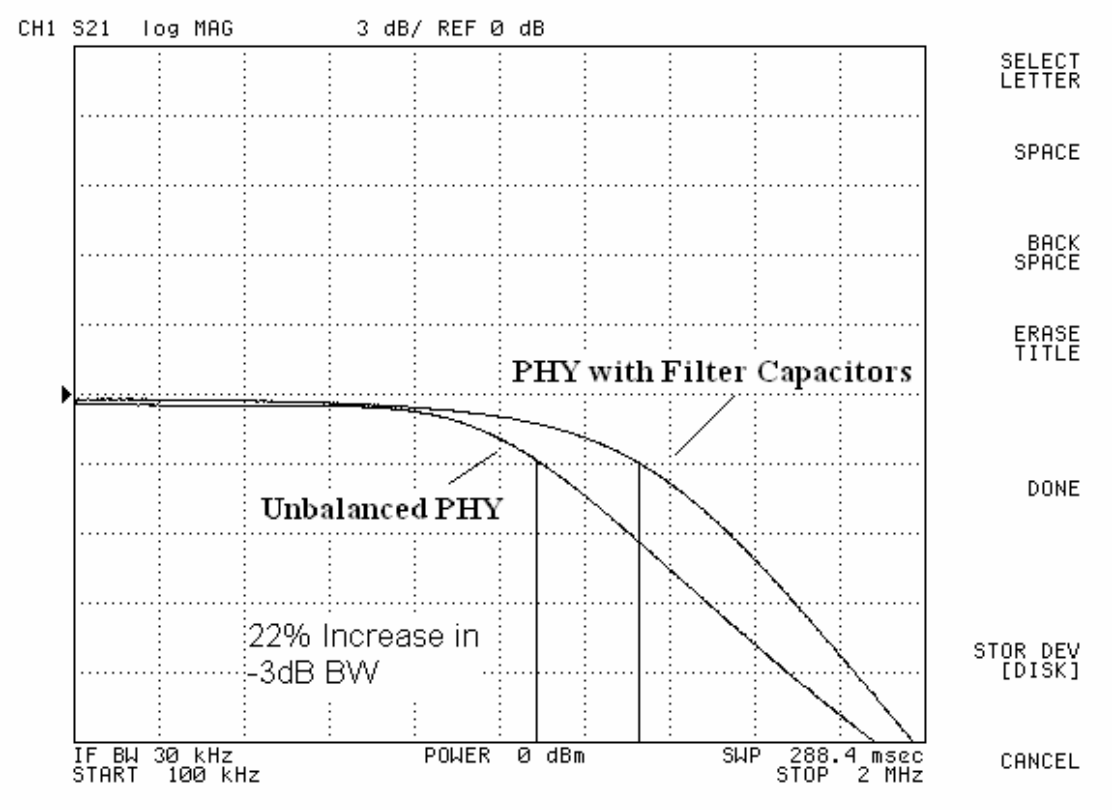

Figure 51: S21 for Balanced and Unbalanced PHY

The above figure shows the improvement in the S21 gained from the Butterworth ladder filter. In addition, the bandwidth of the circuit is increased. Without the filter in place, the S21 corner frequency is 1.125 MHz and when the filter capacitors are added to the circuit, the corner frequency is increased to 1.375 MHz. This is an increase of 250 kHz in the bandwidth, or about 22%.

5.3 Eye Diagrams

The second experiment conducted was the creation of eye diagrams. This was done with an Agilent Infinium 54825A 500MHz 2Gsa/s Oscilloscope and a HP 81130A Pulse/Data Generator. A pseudo-random binary sequence was created, using non-return to zero, NRZ, data with a repetition length polynomial of $2^{15}-1$. The data had a peak to peak voltage of 0.5V up to 1Mbps. Higher data rates were also used to show the effect of the filter at higher frequencies. This setup created a random input for the eye diagrams.

Each of the following eye diagrams showing the effects of the ladder network on the PHY assumes that the load resistance has already been matched. The compensated PHY plots show the optimized capacitances and their effect on the jitter and eye opening. The opening of the eye reflects the increase in bandwidth seen previously. The first eyes shown are at a data rate of 1.4 Mb/sec for the uncompensated PHY and compensated PHY with 12 pF capacitors.

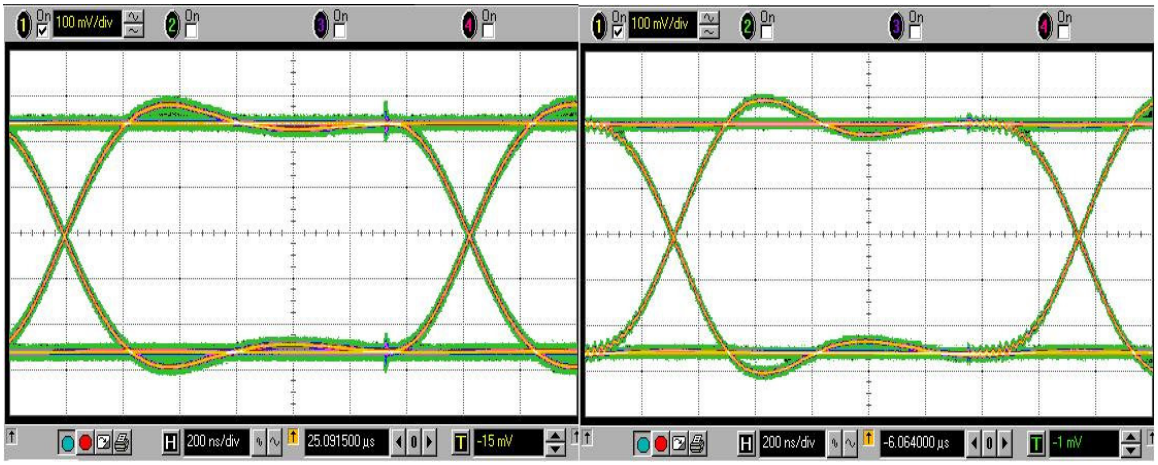

Figure 52: Uncompensated and Compensated PHY Eye at 1.4 Mb/sec

A few observations can be made from the differences in the plots. The first observation is that the zero crossings are more clearly defined with the PHY filter implemented. This decreases the jitter and phase noise seen in the circuit. The second remark that can be made from these plots is that there is some ringing present in both plots. With the compensated PHY, this ringing is increased relative to the uncompensated PHY, however, it resides mostly above the average HIGH or below the average LOW value. This means that for digital signal transmission, it will have no effect on the performance of the system. The third improvement that can be seen especially at the higher transmission frequencies is the opening of the eye when the compensation is performed. The next set of plots is for an input frequency of 2 Mb/sec.

Figure 53: Uncompensated and Compensated PHY Eye at 2 Mb/sec

The same observations can be made at this higher frequency as those at 1.4 Mb/sec. The eye is opened ±25 mV with the addition of the filter capacitors and the jitter is reduced 25%. It also can be seen at 2 Mb/sec that average HIGH and LOW levels are more clearly defined due to the reactance matching, which adds to the opening of the eye. The next figure shows the jitter improvement more clearly.

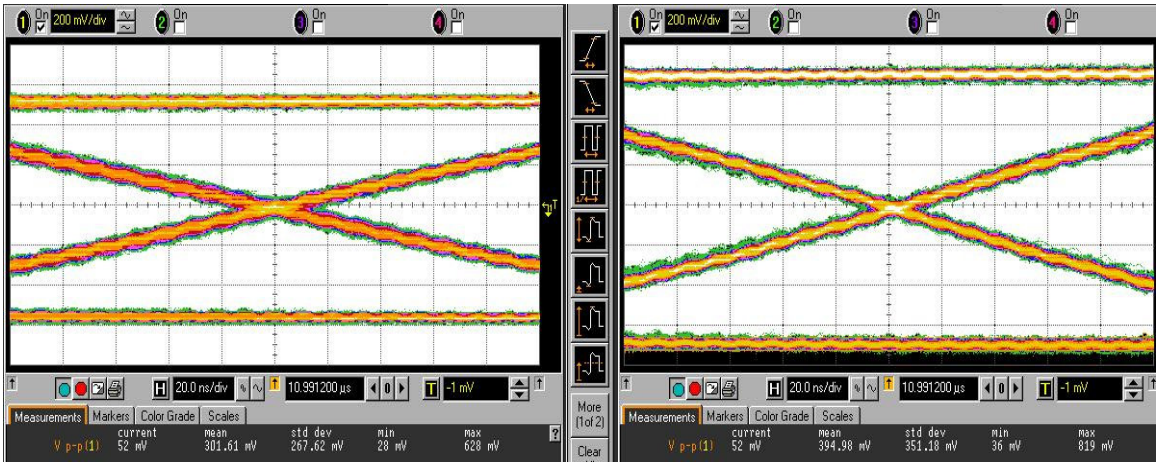

Figure 54: Uncompensated and Compensated Zoomed Eyes at 2 Mb/sec

From the zoomed in eye diagrams above, it can be seen that the voltage amplitude improves to 395mV from 301mV, an improvement of over 31%. Also, because the zero crossing jitter is improved, the 200mV rise time improves to 68ns from 96ns, an improvement of over 29%. The improvements from the filter are seen at higher frequencies as well. Even above the corner frequency of the PHY, the filter helps match the impedance. From the S21 transmission curve shown previously in figure 52, it is known that the 3dB bandwidth of the PHY and filter is around 1.25MHz. The following figure shows the voltage across the PHY with and without the network capacitor terminations at this higher frequency.

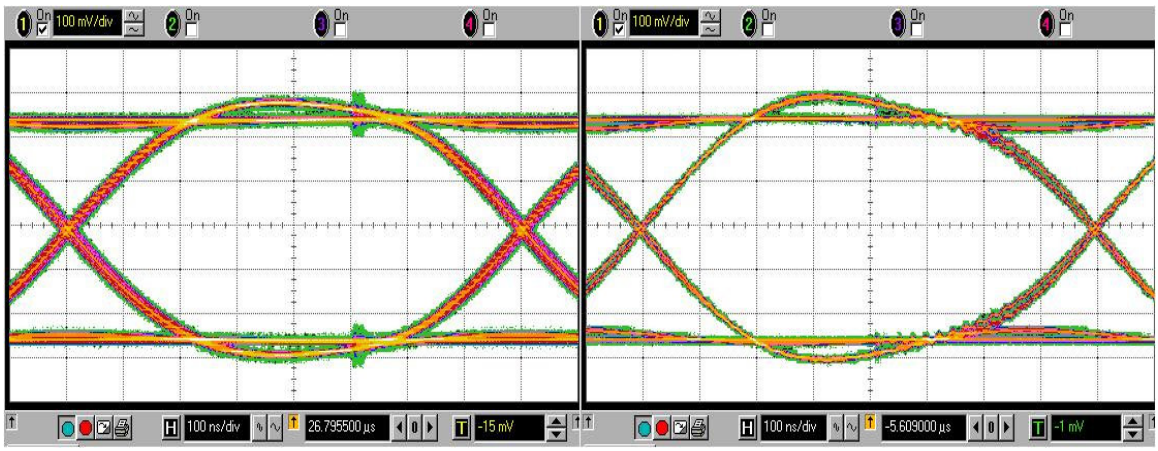

Figure 55: Uncompensated and Compensated PHY Eyes at 2.5 Mb/sec

From the above figure, there is greater than 23% improvement in jitter and the eye is opened about 50 mV more with the compensation in place. All jitter and eye opening improvements are listed in table form at the end of this section. The next figure shows the jitter and eye opening at 2 MHz.

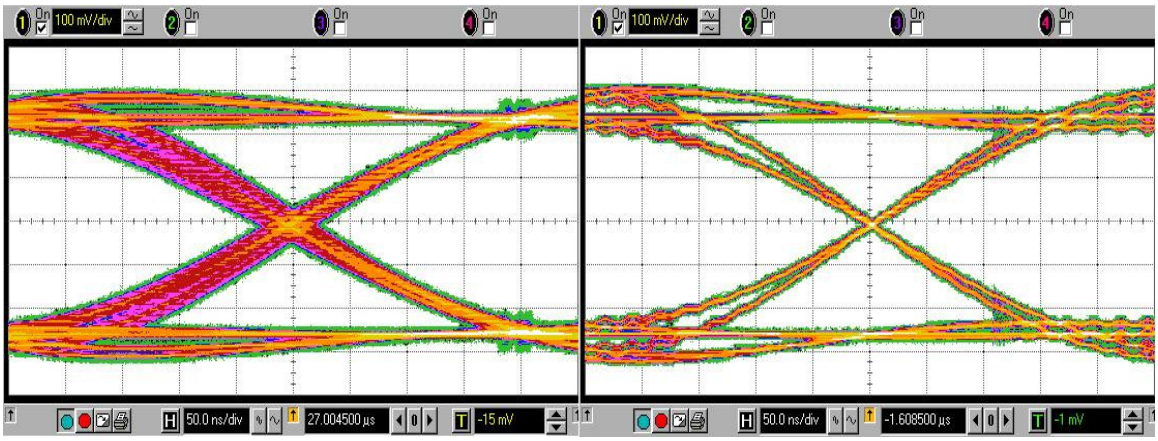

Figure 56: Uncompensated and Compensated Zoomed Eyes at 4 Mb/sec

At this higher frequency the results of the filter are apparent: the opening of the eye is over 10% greater and the jitter in the zero crossings is reduced by 64%. This shows that with the compensation scheme implemented, the bandwidth of operation for digital

signals could be increased dramatically. The following plot shows multiple eyes at 4 Mb/sec, which more clearly shows the opening of the eye.

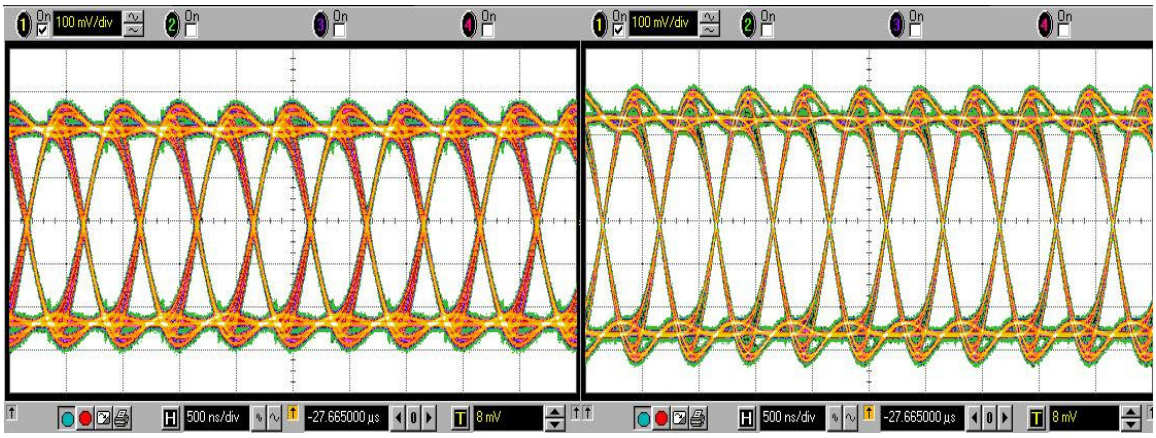

Figure 57: Uncompensated and Compensated Multiple Eyes at 4 Mb/sec

The next figure shows the uncompensated and compensated eye diagrams at 6 Mb/sec. At this higher frequency, the eye is opened by 44% with the filter in place. However, at 3 times the bandwidth of the PHY, the improvement seen in the jitter is no longer there. By extending the ladder network with more compensation elements, the bandwidth could be increased further.

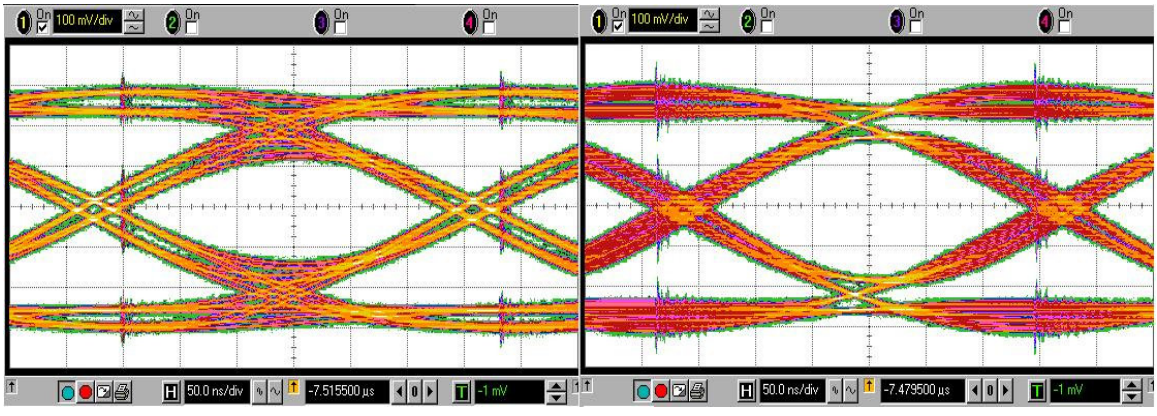

Figure 58: Uncompensated and Compensated PHY Eyes at 6 Mb/sec

The next plot shows multiple eyes at 6 Mb/sec.

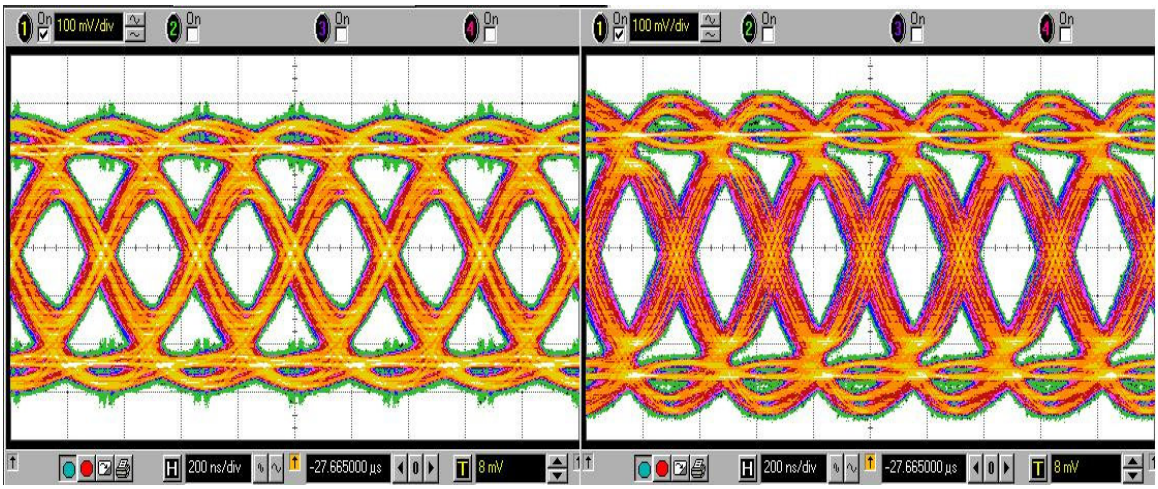

Figure 59: Uncompensated and Compensated Multiple Eyes at 6 Mb/sec

Also seen in the previous two plots of multiple eyes is the effect of the matching on the average HIGH and LOW voltage levels. Because there is a real impedance created from the high order zeros present, the matched PHY has better level definition as well as improved jitter and eye opening. The following table shows the peak to peak jitter reduction at selected frequencies. It also shows the improvement in eye opening by the compensation elements; the front capacitor C1 was 25 nF and the C2 back capacitor was 10 nF. Within the bandwidth of the PHY, the system sees improvements of 23% in the jitter and 12% in the opening of the eye.

Table 12: Jitter Improvement and Eye Opening

Data Rate	Uncompensated PHY Jitter (ns)	Compensated PHY Jitter (ns)	Jitter Reduction [%]	Uncompensated PHY Eye Opening (mV)	Compensated PHY Eye Opening (mV)	Eye Opening [%]
1.4 Mb/sec	25	25	**0**	425	425	**0**
2 Mb/sec	40	30	**25**	400	425	**6.25**
2.5 Mb/sec	65	50	**23.1**	425	475	**11.765**
4 Mb/sec	85	30	**64.1**	375	425	**13.333**
6 Mb/sec	70	70	**0**	225	325	**44.444**

Also from the above table, it is seen that there is an improvement in the jitter and the eye opening not only in the passband of the filter, but up to 2 times the bandwidth as well. Relating this to the S21 transmission bandwidth increase, the implemented ladder filter is shown to increase the operating frequency of the Ethernet circuit.

All of the eye diagrams use two extra capacitors to properly terminate the PHY and convert it to a ladder network. The results obtained show an improvement in the S11, S21 and eye diagrams with minimal additional circuitry, which was the goal of this project. However, greater circuit performance is possible with greater circuit complexity. By following the ladder network theory, additional series inductors and shunt capacitors can be added to correct larger amounts of return loss. An added benefit would be the increase of bandwidth available in such applications as RFIC. If the components of the PHY model could be predicted to within a small range, fixed capacitors could be added to the circuit to improve the return loss. In systems where variable termination impedances would be hard to implement, such as around the bondwire in an IC, this fixed capacitance solution could be implemented to achieve some improvement in performance.

It is seen that the experimental results obtained match with the simulation results in confirming the ladder filter application to the impedance matching problem in Ethernet communication systems. Tangible results may be obtained with discrete components resulting in low implementation cost and complexity. An IC could be fabed in cheaper technology for reduced cost and area as well.

CHAPTER VI

SUMMARY AND CONCLUSIONS

An active on board solution for matching the impedance of an IEEE 802.3 Gigabit Ethernet port to the CAT5 line has been developed. The physical layer of the board can be impedance matched to the transmission line impedance over the bandwidth with the addition of matching circuitry and a control loop, thus minimizing the return loss of the system. A model of the PHY has been created with real and imaginary return loss data employing microwave theory using MATLAB and Cadence; by properly terminating the PHY it is converted into a ladder network. The resulting impedance matched network decreases the S11 return loss over the defined bandwidth, while also increasing the S21 transmission, decreasing the jitter and increasing the eye opening. This solution is realized with minimal cost, area, power and complexity compared to previous design solutions.

REFERENCES

[1]. R. Kamprath, "The Echo Effect". Presented at Dell Computer Server Wired COMMS Engineering, Round Rock Campus, Summer 2005.

[2]. X. Lin, J. Liu, H. Lee and J. Liu, " A 2.5- to 3.5-Gb/s Adaptive FIR Equalizer With Continuous-Time Wide-Bandwidth Delay Line in 0.25-um CMOS", *IEEE Journal of Solid-State Circuits*, Vol. 41, No. 8, pp. 1908-1918, Aug. 2006.

[3]. *Gigabit Ethernet Specification* IEEE STD 802.3 2002

[4]. A. Sedra and K. Smith, *Microelectronic Circuits*, 4th Edition. New York: Oxford University Press, 1998.

[5]. Falco Electronics RJ45 1M093_LS2L18 schematic, available: http://www.falco.com/products/4/pdf/GIGABIT%20RJ45%20SINGLE%20WITH%20LEDS%206X2.pdf

[6]. B. Nauta and M. B. Dijkstra, "Analog Line Driver with Adaptive Impedance Matching", *IEEE Journal of Solid-State Circuits*, Vol. 33, No. 12, pp. 1992-1998, Dec. 1998.

[7]. W. K. Chen (Editor), *The Circuits and Filters Handbook*, 2nd Edition. New York: CRC, 2002.

[8]. L. P. Huelsman and P. E. Allen, *Introduction to the Theory and Design of Active Filters.* New York: McGraw-Hill, 1980.

[9]. A. Valdes-Garcia, F. Abdel-Latif Hussien, J. Silva-Martinez and E. Sanchez-Sinencio, "An Integrated Frequency Response Characterization System with a Digital Interface for Analog Testing", *IEEE Journal of Solid-State Circuits*, Vol. 41, No. 10, pp. 2301-2313, Oct. 2006.

[10]. P. Confalonleri, M. Zarnprogno, F. Girardi, G. Nicollini, A. Nagari," A 2.7 mW 1 MSps 10 b Analog-to-Digital Converter with Built-in Reference Buffer and 1 LSB Accuracy Programmable Input Ranges", *Proc. IEEE 30th Eur. Solid-State Circuits Conf. (ESSCIRC), Leuven, Belgium,* Sep. 2004, pp. 255-258.

[11]. S. Norsworthy, R. Schreier and G. Temes, *Delta-Sigma Data Converters.* New York, John Wiley and Sons, Inc., 1997.

[12]. T. H. Lee, *The Design of CMOS Radio-Frequency Integrated Circuits*. New York: Cambridge University Press, 2004.

[13]. Analog Devices ADC selection table, available: http://www.analog.com/IST/SelectionTable/?selection_table_id=185

[14]. D. A. Johns and K. Martin, *Analog Integrated Circuit Design.* New York: John Wiley and Sons, Inc., 1997.

[15]. Gray, Hurst, Lewis and Meyer, *Analysis and Design of Analog Integrated Circuits.* New York: John Wiley and Sons, Inc., 2001.

[16]. Agilent 4395A Spectrum/Impedance Analyzer Manual, available: http://cp.literature.agilent.com/litweb/pdf/04395-90040.pdf

APPENDIX

Bond Wire Application

The ladder filter impedance matching circuit presented in this project has other very useful applications in addition to PHY return loss minimization. One of these that could have a large impact is on bond wire impedance matching. For systems sending high speed or digital data off chip, the bond wire presents non-ideal frequency dependant impedance. The bond wire can be modeled as in [13], and the ladder filter form [7] can be used to compensate the bond wire impedance to the required nominal value. The following is a schematic of the QFN64 Package Model. This is the model used in simulation to determine parasitic performance results.

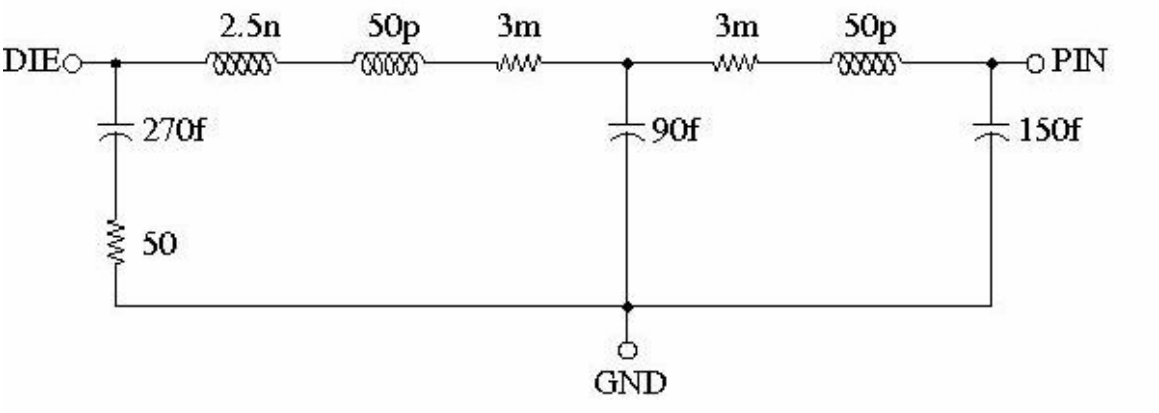

Figure 60: QFN64 Package Model

In the interest of ease of implementation, it is assumed that a first hand approximation of the solution can be accomplished by adding a shunt capacitor at the DIE side and another at the PIN side. The following plot shows the S11 return loss for swept values of the

filter capacitors. The capacitors at the DIE end and the PIN end are swept with equal values.

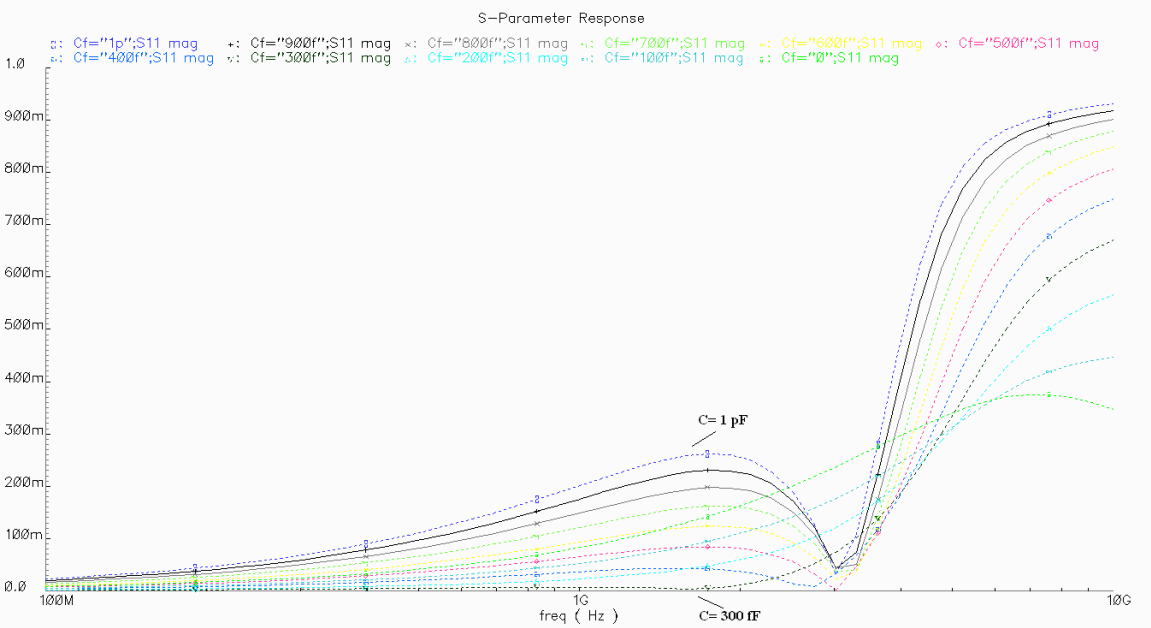

Figure 61: S11 Return Loss for Bond Wire Ladder Filter

It can be seen that by adding a small amount of capacitance before and after the bond wire, the return loss can be greatly diminished. From the plot, it can also be seen that the filter's corner frequency is around 3GHz for this particular filter capacitance sweep. With incremental capacitor values used in the sweep, the ideal values for this bond wire model can be found. The following schematic shows the S21 transmission and S11 return loss plots for the best case filter capacitor value of 300fF.

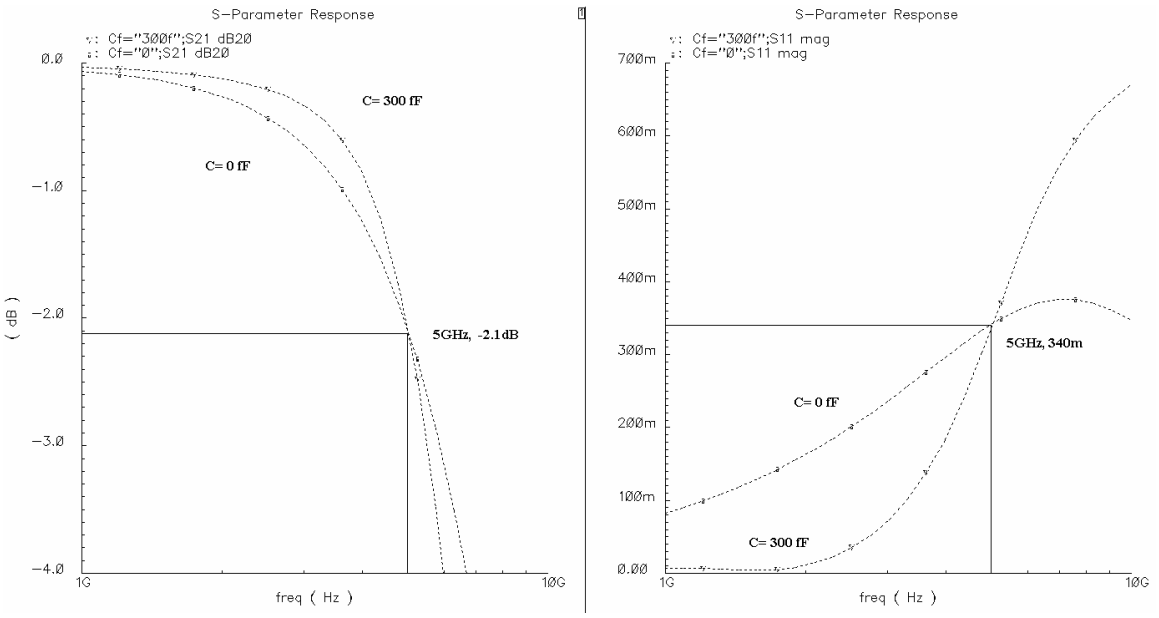

Figure 62: S21 Transmission and S11 Return Loss for Bond Wire with and without Best Case Filter

It can be seen from the S21 and S11 plot that the filter has a useful bandwidth of 5GHz. Above this frequency, the benefit of the filter is lost. However, below this corner frequency, the filter reduces the return loss of the bond wire and increases its transmission. If the capacitors were variable, they could be tuned such that the benefit in transmission would be greatest in a frequency range of choice below the corner frequency. Another condition to consider is that these ladder network implementations only use two component additions. If extra components are added and swept such that the bondwire is matched to 50 Ohms, the bandwidth of operation could be increased even further and the passband characteristics would improve as well.

The largest problem that would have to be overcome with an RFIC bondwire matching circuit is that part of the circuit would be on chip and the other part would have to be off. However, if the control circuitry is put on chip and a digital signal is sent off

chip to turn on and off the impedance banks used for matching, the circuit would create a robust solution to matching bondwire impedance.

Cadence Simulation Plots for PHY Value Ranges

In this section, a few sets of plots will show the PHY voltage for alternate component values. The ranges of components chosen for the filter implementation determine the range of applicability to the PHY impedance. Thus, if 20pF is the largest value that either filter capacitor may be, the ranges of values for the PHY are given in the Cadence Simulation Results chapter. The first figure is the PHY model presented again for the reader.

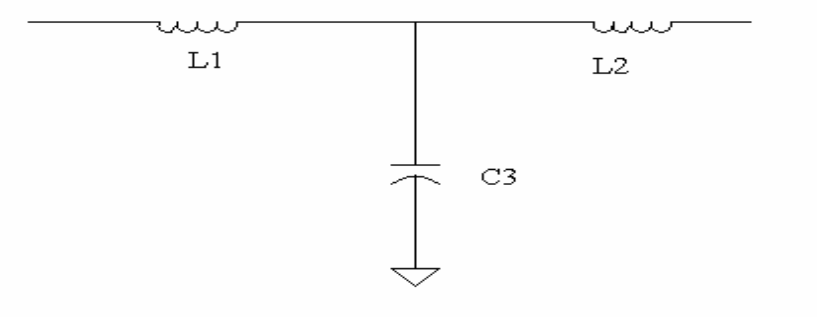

Figure 63: Second Approximation Pole/Zero Model (REPRINTED)

For the following set of plots, L1=100nH, L2=100nH and C3=45pF.

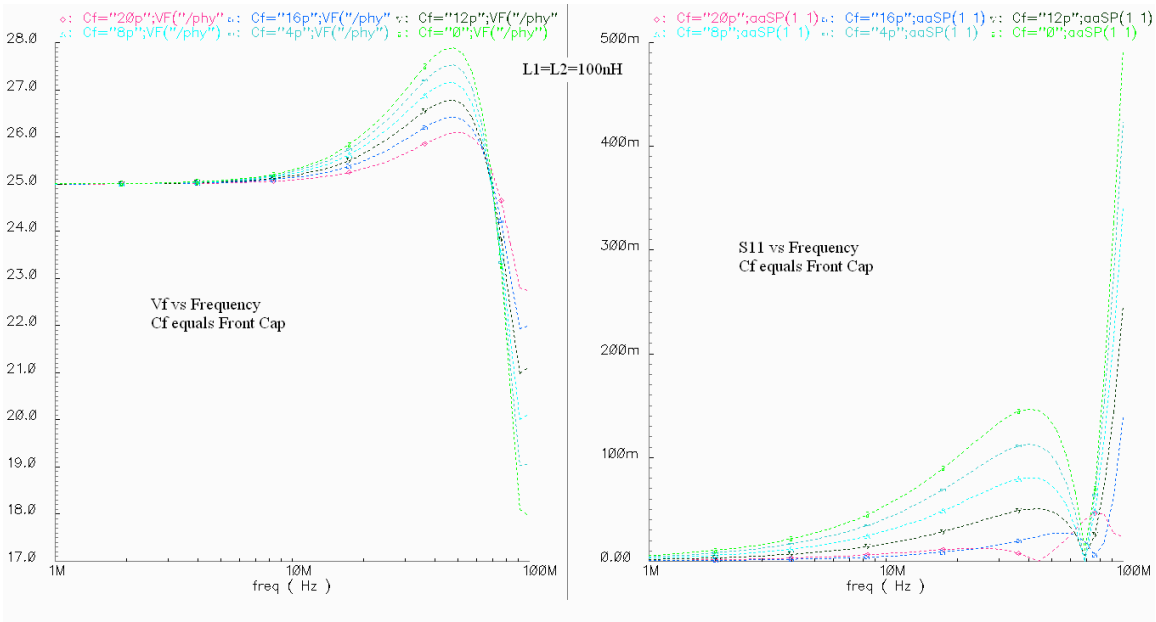

Figure 64: PHY Voltage and Return Loss vs. Frequency for L1=100nH, L2=100nH and C3=45pF

It can be seen above that for these PHY values, the filter capacitors should be set to 20pF. From there, a sweep of the front capacitor could be performed, and a larger value could be added. However, even without optimizing the capacitor values and ratio, it can be seen that there is an improvement in return loss of up to 140m at 40 MHz. At 100 MHz, there is an improvement of 465m in the magnitude return loss. The next plot shows the opposite end of the spectrum for a PHY with component values of L1= 100nH, L2= 75nH and C3= 30pF.

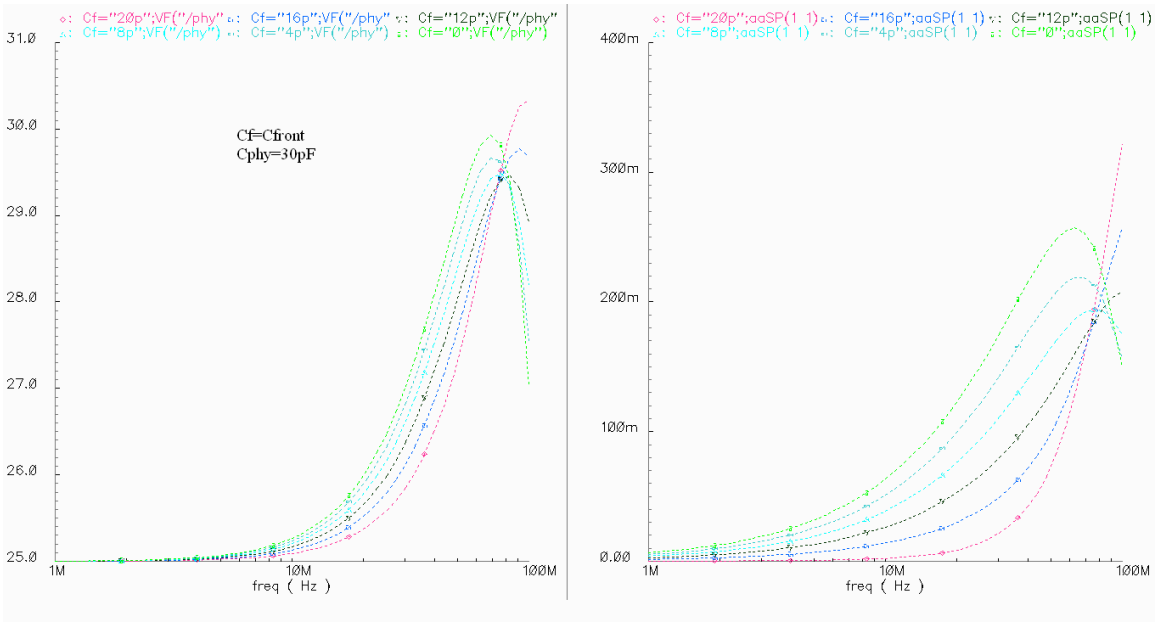

Figure 65: PHY Voltage and Return Loss vs. Frequency for L1= 100nH, L2= 75nH and C3= 30pF

From the above figure, it is seen that at the lowest value of PHY C3 capacitor of 30pF, the addition of equal filter capacitors of 20pF is best. Within the passband of the filter, there is substantial improvement of return loss. Again it is orders of magnitude better than without the filter implemented.

VITA

Richard Alan Kamprath received his Bachelor of Science degree in electrical engineering from Texas A&M University, College Station, Texas in 2004. Since then, he was awarded his Master of Science degree in May of 2007 with the Analog and Mixed Signal Center at Texas A&M University under Dr. Jose Silva-Martinez. His current research interests include Radio Frequency Integrated Circuit Design and their application to future markets. He is currently working for Schlumberger in Sugar Land, Texas in the Nuclear Magnetic Resonance Group. He can be reached through the Department of Electrical Engineering, Texas A&M University, College Station, Texas 77843 or through the Schlumberger Sugar Land Technology Center at 281-285-7548.