WARPLINK™ REFERENCE DESIGN PLATFORM

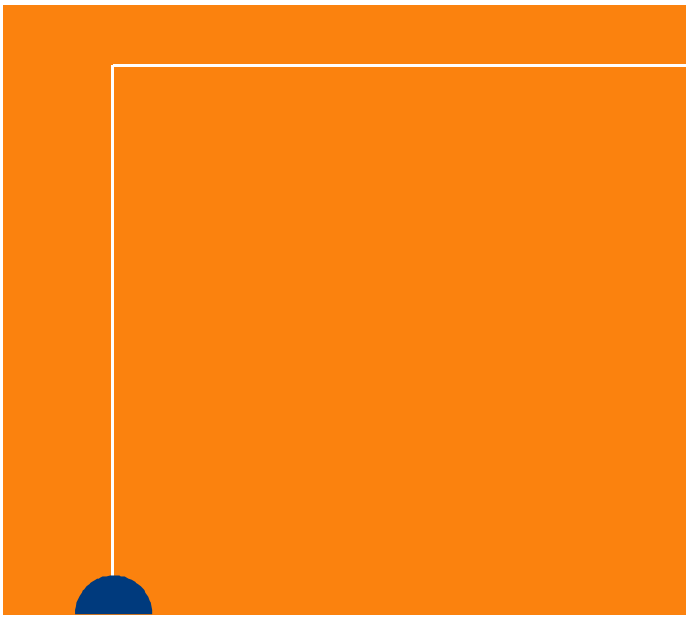

WHITE PAPER

BR1570/D, Rev 0

TABLE OF CONTENTS

ABSTRACT

This paper describes the technical design process used in the development of a WarpLink™ Reference Design Platform showcasing Motorola's WarpLink 2.5 Gbps Quad, a Serializer-Deserializer (SERDES) data interface device. The Reference Design Platform utilizes a combination of good design practices to make it feasible to use an FR-4 backplane with standard components and fabrication processes to realize 3.125 gigabaud (Gbd) XAUI-compliant channels.

A brief description of critical component selection is followed by a discussion of design considerations related to high-speed signal integrity performance. The technical "Right-by-Design" process utilized by North East Systems Associates (NESA) included simulations that were verified through measurements of the fabricated hardware. The simulation and experimental results highlight the special features of the WarpLink 2.5 device making it possible to have 3.125 Gbd wire-speed transmissions through FR-4 PWB material.

INTRODUCTION

WARPLINK 2.5 QUAD DEVICE

The WarpLink 2.5 Quad device is a SERDES interface that transfers data between chips across a board, a backplane, or cables. It handles four full-duplex redundant data links. Serial transceivers transmit and receive 8B/10B coded data at a nominal rate of 2.5 gigabits per second (Gbps) through 3.125 Gbd links.

In the transmit direction, the near-end WarpLink device receives data on its four parallel transmit interfaces which accommodate 8-bit uncoded or 10-bit precoded data bytes. When configured for the 8-bit mode, the device performs 8B/10B encoding on the uncoded data. It then serializes the coded data and sends it onto the four corresponding primary serial differential transmitters. Data can be sent out on the four redundant transmitters or on both primary and redundant transmitters simultaneously. Coded serial data comes out of each transmitter at 3.125 Gbd wire-speed carrying 2.5 Gbps of user data through channels across a board, a backplane, or cables to far-end WarpLink devices' serial receivers.

In the receive direction, serial coded data coming from far-end WarpLink devices' serial transmitters are received by the near-end device on one of four primary or redundant serial receivers. The near-end device de-serializes the data and, if configured for the 8-bit mode, performs 8B/10B decoding. The device then sends the data out on the four corresponding parallel receive interfaces.

The WarpLink 2.5 Quad is packaged in a 324-pin PBGA, with a 19 mm x 19 mm (0.75 x 0.75 in) body and a 1 mm ball-to-ball pitch. The device typically uses 1.8 watts. Its core and link power supply inputs require 1.8 V. The HSTL I/O power supply inputs (for the parallel and digital I/Os) use either 1.5 V or 1.8 V.

For further detail on the WarpLink 2.5 Quad's rich feature set, including Selectable Speed Range, Double Data Rate, 8B/10B, Link Synchronization and Recovered Clock Mode, refer to the WarpLink 2.5 Quad User's Manual.

WARPLINK REFERENCE DESIGN PLATFORM GOALS

With the introduction of the WarpLink SERDES family, consisting of the WarpLink Quad (1.0 Gbps), WarpLink Quad Double Data Rate (DDR) (1.0 Gbps), and WarpLink 2.5 Quad (2.5 Gbps) devices, designing boards and backplanes capable of handling up to 3.125 Gbd channels could prove challenging. The WarpLink Reference Design Platform was developed to assist designers in laying out the interconnection system as well as to provide an example of a system designed with a backplane, daughter cards, and connectors that make 3.125 Gbd channels feasible. The development work served as a learning experience for Motorola, with the aim of sharing the final outcome with WarpLink customers.

A critical portion of the development effort centered on the simulation. In addition to simulations performed at 3.125 Gbd, simulations at 5Gbps were also carried out to give a better understanding of the current capabilities and limitations of the technology.

The platform was designed to accommodate WarpLink 2.5, as it is the fastest device in the family. FR-4 was chosen for the backplane and boards for three reasons: 1) material availability, 2) relative lower cost when compared to other fabrication materials, and 3) most PWB fabrication houses can process the material. Several channels were designed to be long enough to demonstrate WarpLink 2.5's XAUI compliance (50cm or 19.68in).

Standard available components were selected, specifically the ERNI HMZD connector, a high-density connector intended for high-speed signaling. Motorola's Timing Solutions devices, the MC100ES6111, the MC100EP222, the MPC9456, and the MC100EP8111, were examples of currently available clock network distribution devices that are intended for use in such a system.

WARPLINK REFERENCE DESIGN PLATFORM OVERVIEW

Architectural Overview

Mesh and Fabric

The WarpLink Reference Design Platform consists of a 9U 19" Compact PCI chassis with eight double-height slots. The slots accept three types of cards: the Line Card, the Switch Card, and the Test Card. WarpLink devices on the cards are connected to one another through the backplane interconnect, which supports both mesh and fabric switching schemes.

For the mesh-switching scheme, each slot has eight mesh ports. The mesh channels in the backplane connect seven mesh ports per slot to a mesh port on each of the other seven slots. The one remaining unconnected port is looped back to itself. The connection is done such that slot-*m*, port-*n* is connected to slot-*n*, port-*m*. Figure 1 and Figure 2 illustrate the mesh interconnects of the backplane for slot 1 and 2, respectively. In this scheme, 2.5 gigabits of link data can go directly from any slot to any other slot.

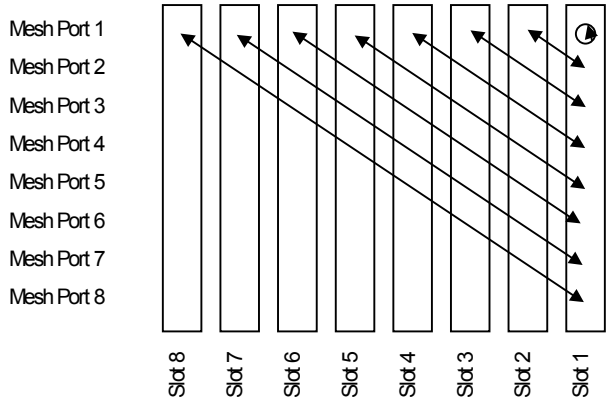

Figure 1. Mesh Interconnects from Slot 1

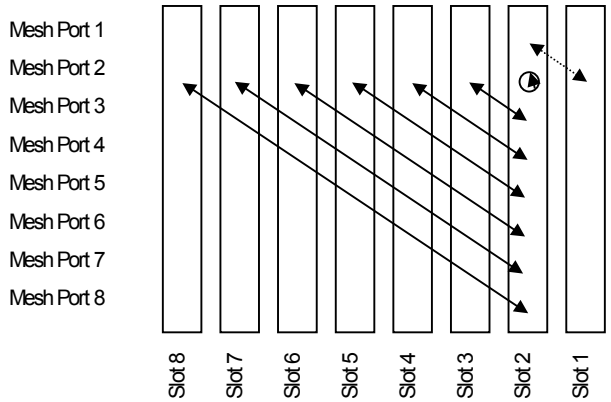

Figure 2. Mesh Interconnects from Slot 2

In the fabric-switching scheme, each slot has eight line ports: four primary, four redundant. Slots 7 and 8 each have 32 fabric ports in addition to the line ports. The fabric channels of the backplane connect the four primary line ports of every slot to four of the 32 fabric ports on slot 8. The four redundant line ports of every slot are connected to four of the 32 fabric ports on slot 7. Figure 3 shows the fabric interconnects of the backplane. In this scheme, 10 gigabits of link data can go from any slot to any other slot through slots 7 or 8. Data links going through slot 8 are primary and those through slot 7 are redundant.

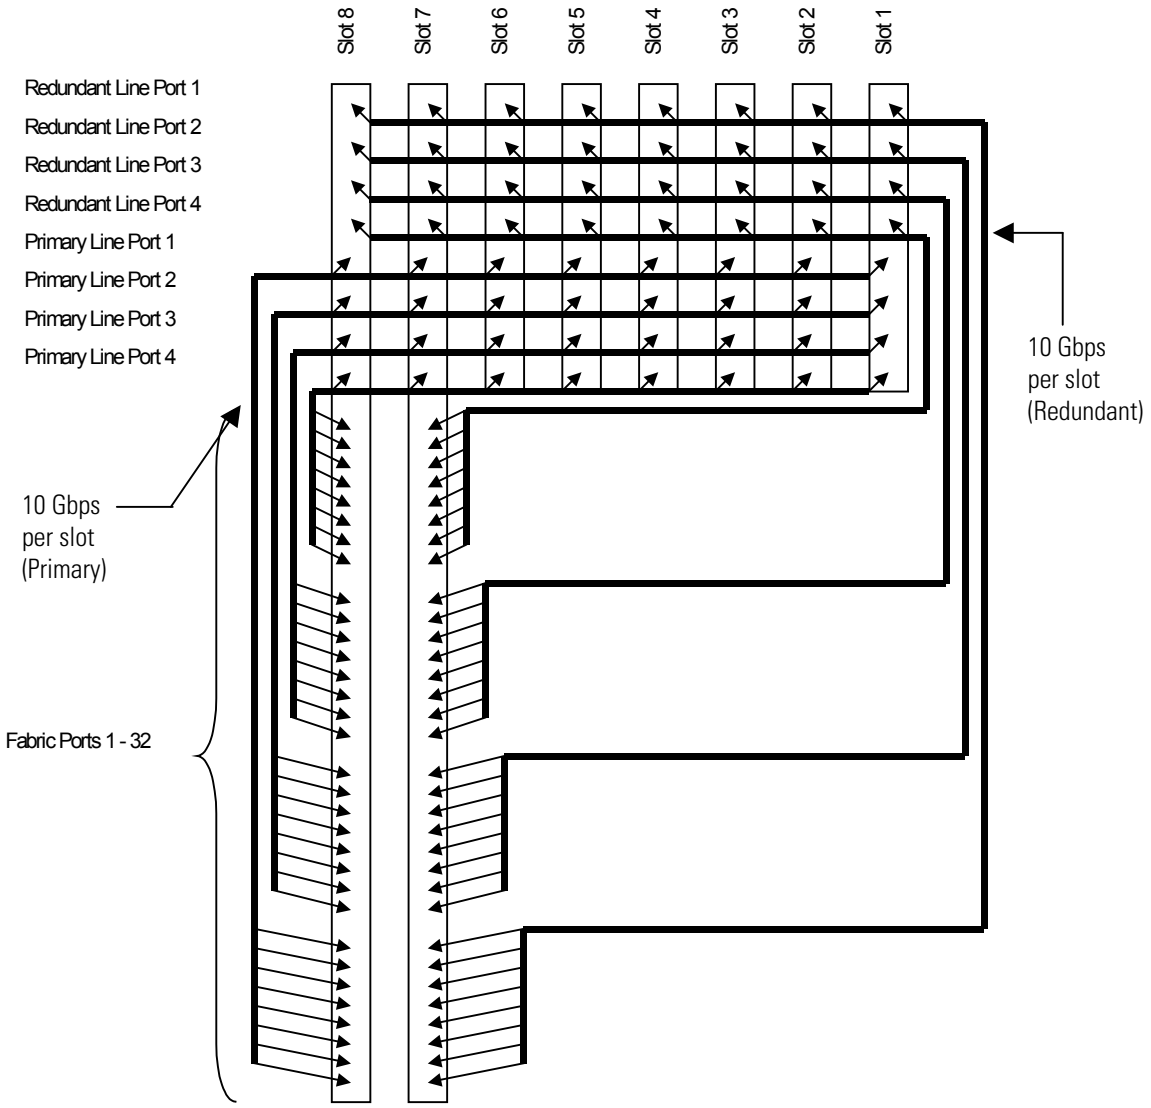

Figure 3. Fabric Interconnects

Clock Distribution

The system has two point-to-point distributed clock networks. The primary network originates from slot 8 and is distributed to all slots. The redundant one is from slot 7, which is also distributed to all slots. The card in each slot selects its own on-board clock, or the clock distributed by either the primary network or the secondary network. Motorola's MC100ES6111 device (on the card in slot 7 or 8) is used to drive the clock distribution network. An MC100EP222 clock receiver on each Line Card is used to receive the distributed clock, and the MC100EP8111 HSTL clock buffer on the Line Card brings the clock signals to HSTL level to be used by WarpLink.

Processor Bus

The system has a multi-drop 32-bit bus going to all slots that can be used as a processor bus. A host card in slot 7 or 8 with a processor can send and receive control and status messages to all slots. Slot 8 houses the primary control card, and slot 7 the redundant control card. The multi-drop 32-bit bus is designed to terminate Gunning Transceiver Logic Plus (GTLP) drivers.

Backplane

The key subassembly of the WarpLink Reference Design Platform is the backplane. It is a 14.96 in x 15.57 in FR-4, 28-layer Printed Circuit Board (PCB). Eight columns of connectors form the eight slots into which the cards are inserted. The majority of the backplane traces form the 3.125 Gbd channels that interconnect the WarpLink 2.5 devices on the daughter cards through the mesh and fabric topologies described previously. Other traces form the clock distribution network and the processor bus.

An important aspect of the backplane design was the selection of the high-speed differential connector from ERNI. The ERmet® ZD 8x10 connector is used as a differential board-to-backplane connector rated to support electrical signaling up to 5 Gbps. The ERmet® ZD is compatible with the mechanical characteristics of the ERmet® 2 mm-HM specifications as defined in IEC 61076-101.

Line Card

One of the cards that reside in the WarpLink Reference Design Platform is the Line Card. It can go into any of the eight slots. In the mesh topology, all eight slots can be populated with Line Cards. In the fabric topology, slots 1 - 6 can be populated with Line Cards, while slots 7 and 8 must be populated with the Switch Cards. Each Line Card has two WarpLink 2.5 Quad devices on it.

Each WarpLink 2.5 Quad device has four SERDES devices each capable of supporting one full-duplex data link. The parallel side of each SERDES device is not redundant. Therefore, each SERDES device has only one parallel port. This parallel port of the full-duplex data link supported by a SERDES device consists of 8 or 10 pins for the parallel transmit interface and 8 or 10 pins for the parallel receive interface. The serial side, however, is redundant. Therefore, each SERDES device has two serial ports: a primary one and a redundant one. Each of the serial ports consists of 4 pins: positive-transmit, negative-transmit, positive-receive, and negative-receive.

A Line Card with two WarpLink 2.5 Quad devices, therefore, can support eight full-duplex data links, to be used in either the mesh or fabric configuration.

The eight serial primary ports provided by the two devices are connected to the eight mesh ports of the slot when the card is inserted. The mesh ports are then connected to other mesh ports as previously discussed. In the mesh configuration, the primary serial ports are selected, and data always passes through the two WarpLink devices via their parallel ports and primary serial ports.

The eight serial redundant ports provided by the two devices are connected to the eight fabric-topology line ports of the slot when the card is inserted. The line ports are then connected to the fabric ports on slot 7 and 8 as previously discussed. In the fabric configuration, the redundant serial ports are selected, and data always passes through the two WarpLink devices via their parallel ports and redundant serial ports.

On the parallel side of the WarpLink devices, the HSTL parallel interfaces can be connected to ASICs, FPGAs, or headers to be controlled and monitored by pattern generators and logic analyzers.

Switch Card

The Switch Card is a combination of a Line Card and a Fabric Card. This card can only go into slots 7 and 8. There are 10 WarpLink devices on the Switch Card. Two of the 10 devices are used to form the Line Card portion: the devices' serial primary ports are connected to the mesh ports and the redundant ports to the line ports. The 32 serial redundant ports of the remaining eight devices are connected to 32 fabric ports of the slot into which the card is inserted. The 32 primary ports of these eight devices are unused.

On the parallel side of the WarpLink devices, the interfaces can be connected to ASICs or FPGAs.

Test Card

The Test Card is a passive card with SMA connectors and 2x10 headers. It does not have any WarpLink or active devices on it.

The card is used to characterize channels from slot to slot. For example, Time Domain Reflectometry (TDR) measurements can be performed with this Test Card to show via, connector, and backplane trace impedance.

The card can also be used to independently characterize transmitters and receivers operating in the system. When the card is in a slot, controlled test signals from a signal generator can be fed into the near-end Test Card's transmit port of a link to characterize the corresponding receiver of a far-end WarpLink device. Also, signals from a far-end WarpLink transmitter can be measured off the near-end receive port, through a pair of SMA connectors on the Test Card.

DETAILED DESIGN DESCRIPTIONS

WARPLINK REFERENCE BACKPLANE

The WarpLink Reference Backplane is specifically designed for the Motorola WarpLink 2.5 Quad component. The overall design is representative of a *typical* backplane used in present-day communications-system chassis. Figure 4 details the overall backplane as well as component selection and placement on the WarpLink backplane. Figure 5 details the layer stackup of the backplane.

Backplane/Chassis Design Considerations

Before we discuss the details of the Motorola WarpLink backplane, it would be helpful to give an overview of the process and constraints under which a backplane and chassis are designed. The design of the backplane must consider the design elements of the system as a whole. The major parameters that shape a chassis/backplane system must be established early on in the design, once the intended market is established. These include 1) the total power dissipated by the system, 2) the system logical architecture, 3) the approximate dimensions of the chassis, and 4) the subsequent regulatory requirements for the intended market. One of the outcomes of coupling the system architecture with the total power dissipation is the determination of the number of daughter card slots within the chassis dimension limitations. Knowing the maximum power dissipation of the various types of daughter cards and the number desired in a system helps to establish the constraints that limit the possible system configurations. At times, thermal constraints can directly impact the electrical architecture of a system. This is especially true if the chassis is sized so that it places limits on the total power that can be dissipated.

A first pass formulation of a system takes place once a balance is reached between the requirements set by the system architecture coupled with the intended market and the total power that can be dissipated within the chassis enclosure. The thermal analysis and subsequent physical constraints on the slot pitch are required to determine the total number of daughter cards possible in the system. Thus, the ability of a chassis to cool itself can limit the number and possibly the type of components used, including logic devices.

The design of the backplane begins in parallel to the thermal evaluation of the system, but is not set until the number of daughter card slots and overall system configuration are established. Once the number of slots is determined, the next major step is the identification of the type and number of signals per slot. A definition of the signaling needed in the WarpLink Reference Backplane was based on the architecture defined by the Motorola design team. The general architecture of the WarpLink Reference Design Platform is discussed in the *Introduction* of this document.

WarpLink Backplane Physical Description

The physical configuration of the Motorola WarpLink backplane used in the reference system contains eight slots on a 1.6-inch card-to-card pitch. The Motorola WarpLink Reference Design Platform was designed to suit both switch fabric and mesh architectures. The logical configuration of the backplane defines slots 1 through 6 as pin-compatible Line Card slots, and slots 7 and 8 as pin-compatible Switch/Line Card slots. When a switch fabric is inserted into slot 7, 8, or both, the system can function in either the fabric or mesh architecture. If a Line Card is inserted into slot 7, 8, both, or neither, the system will operate in the mesh architecture. The ability of the WarpLink to do this is due to the primary and redundant serial I/O characteristics of the device.

The signal allocations for each slot type are detailed below in Table 1.

Table 1. Signal Allocations

Signal Types	Slot 1 - 6 (Line Card)	Slot 7 - 8 (Switch/Line Card)	Signal Speed
Differential Gigabit Serial Links (Mesh/Fabric)	32 pairs	96 pairs	3.125 Gbd
Differential Clocks	2 pairs	8 pairs	156 MHz
Single-ended Address/Data/Control Busses	36 pins	36 pins	<50 MHz

The signal speed identification is critical to the design rules formulated for the layout of a gigabit backplane. For example, the addition of unnecessary vias in the gigabit lines would adversely affect their signal transmission characteristics. Each via presents an impedance discontinuity that will contribute to the distortion of the opening of the receiver input eye pattern. This can reduce available voltage margins as well as time margins through an increase in signal jitter. Therefore, all differential routes of the WarpLink Reference Design Platform backplane had vias placed only where absolutely needed - at the device pads and at the connectors (plated-through holes).

A number of issues were assessed in arriving at a final connector selection for the Motorola WarpLink backplane. The ERmet® ZD high-speed backplane connector was chosen for the gigabit nets for its high-frequency signal transmission characteristics and because it is side stackable with a standard 2 mm hard metric connector. The ERmet® 2 mm hard metric connector was selected for the slower speed single-ended and differential clock nets.

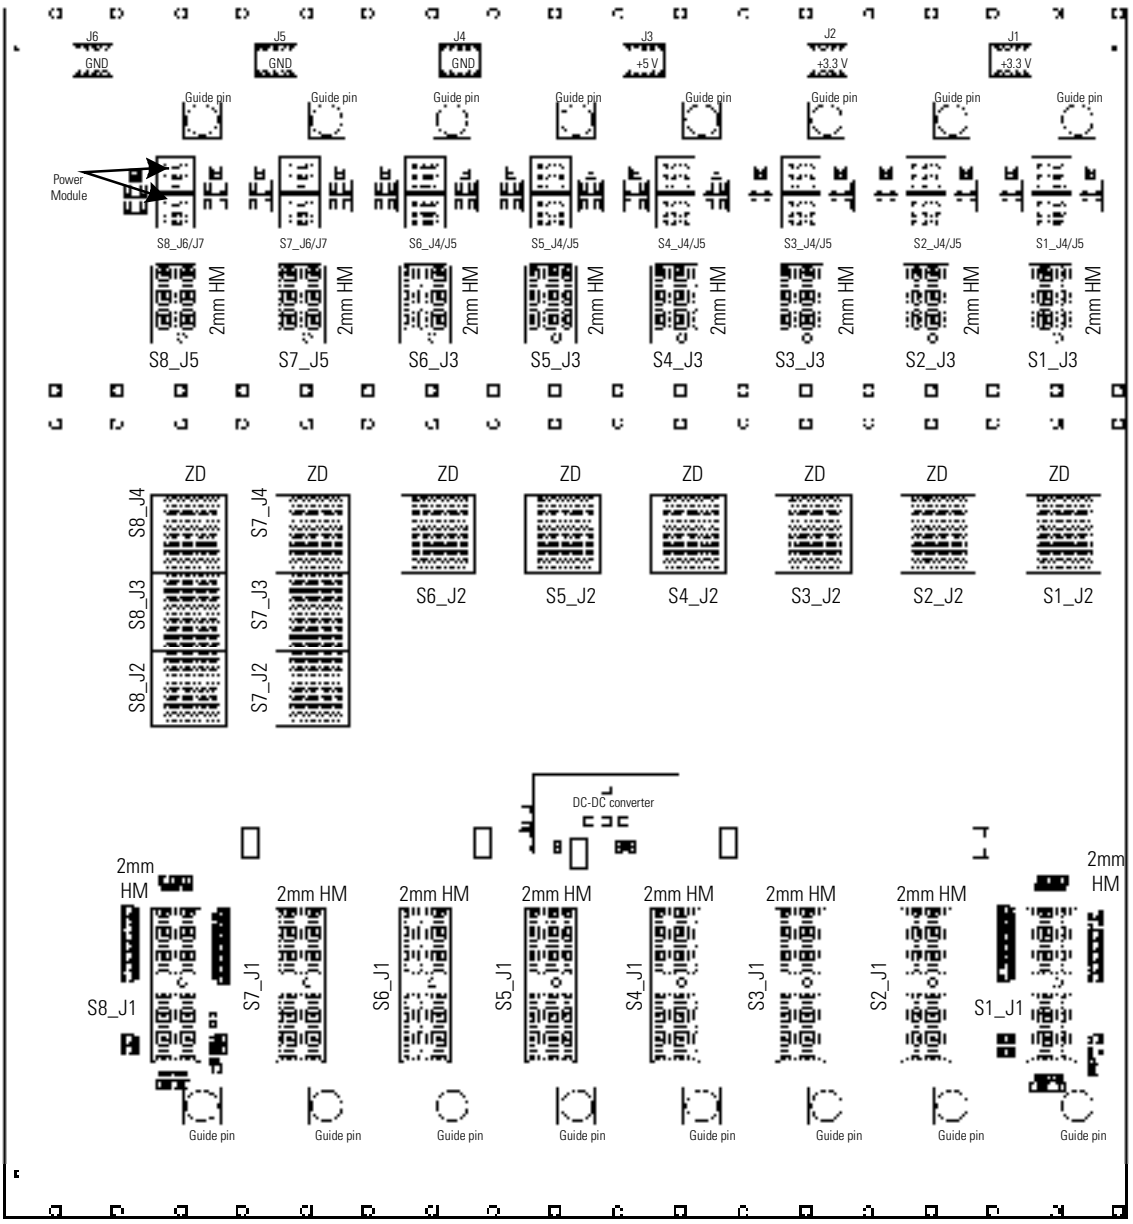

Figure 4. WarpLink Reference Backplane (Front View)

The backplane not only connects the eight slots of the Motorola WarpLink system; it also has been designed with features that help exhibit its gigabit transmission capabilities. The longest gigabit traces required to connect slots 1 to 8 are about 12 to 15 inches. The WarpLink backplane also contains four differential nets that are approximately 0.5 meter long. These nets demonstrate WarpLink's ability to operate over a long lossy channel. They are also used to test WarpLink's compliance to XAUI specifications. The backplane also contains four traces designed to inject crosstalk onto gigabit links. This serves to demonstrate WarpLink's noise tolerance. These noise-generating nets have been placed 5 and 10 mils away from the victim gigabit nets. A recommended design rule spacing of 15 mils (minimum) was implemented for the remaining gigabit nets in order to minimize any crosstalk.

The gigabit links are confined to the middle of the backplane and are connected to the daughter cards through the ERNI Components ERmet® ZD High-Speed Backplane Connector. The slow-speed (single-ended) buses are positioned in the lower connector section of the backplane and are connected to the daughter cards via ERNI Components ERmet® 2 mm Hard Metric (HM) connectors. The clock nets run across the upper connector section of the backplane and are connected to the daughter cards via with the ERNI 2 mm HM connectors as well. The backplane provides 3.3V and 5V power to the daughter cards by way of the ERmet® power modules at the top of each slot.

Backplane Design Rules and Layer Stackup

The stackup cross-section of the WarpLink Reference backplane board is shown in Figure 5. Both broadside-coupled differential and single-ended routes are utilized in the layout. The characteristic impedance of the single-ended nets was designed to be 50 ohms and the broadside-coupled differential nets 100 ohms (both ±10%). All differential pairs were routed direct without vias. Vias for single-ended, multi-drop buses were used only within the connector regions and at the trace ends for termination access. The WarpLink backplane was designed as a 28-layer, 0.250in thick PCB. The backplane contains six broadside-coupled routing channels (12 layers) and 14 single-ended routing layers (12 of which are shared with the differential pairs). Broadside-coupled differential routing was chosen over edge-coupled or uncoupled differential interconnects for its ability to maintain the routing channel within the gigabit connector footprints - board thickness was not a critical parameter. This is a typical routing scheme when using high-speed differential connectors on backplanes of substantial board thickness.

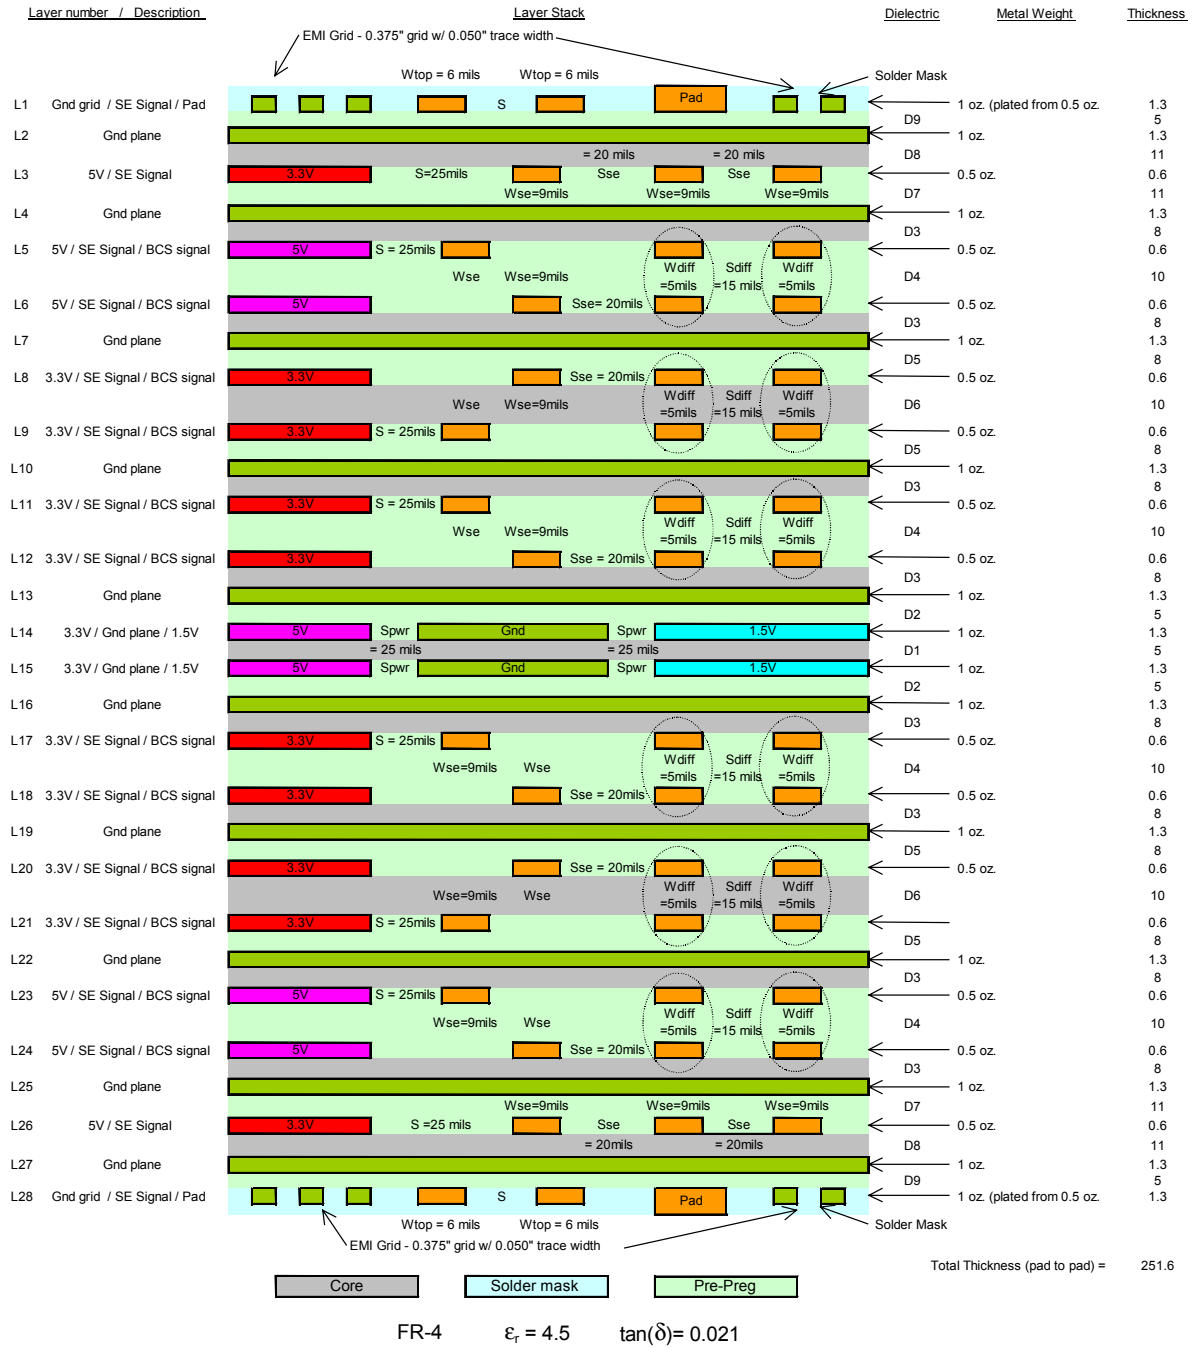

Figure 5. WarpLink Reference Backplane Layer Stackup

The WarpLink Reference Design Platform has been designed for three types of daughter cards. Two of the cards, the Switch Fabric/Mesh Card (Switch Card) and the Line Card contain active components in addition to the WarpLink SERDES. The third card is the passive Test Card. The design intent of the Test Card is to provide access to the Gigabit nets for direct sampling of the WarpLink signals as well as providing access to the backplane for passive signal integrity measurements (via SMA test connectors).

LINE, SWITCH, AND TEST CARDS

Daughter Cards Design Considerations

An important design consideration for the daughter cards should be the pin assignment of the high-speed connector. This pin assignment, in a way, is affected by the number of high-speed routing layers in a card. The number of layers should be minimized so that the card would be thin enough to fit the guiding rails of the chassis. From a signal integrity viewpoint, this should also be done so that the vias' stubs would be as short as possible, especially for the 3.125 Gbd traces that are routed in the upper layers. For the chosen gigabit connector, four was the minimum number of high-speed layers we could have. Once the number of layers is determined, a preliminary hand-routing exercise should be carried out to determine the pin assignment, or to determine whether the existing pin assignment is or is not routable given these four layers. Care should be exercised when assigning pins to make both the Line Card and the Backplane routable. Often, a couple of iterations would be required. This hand-routing exercise should be done early in the design cycle before committing to a pin assignment.

Another trivial design consideration for the daughter cards should be the placement of the WarpLink device. The device should be placed as close to the connector as possible. For the Line Card, one WarpLink was placed close to the connector, and the other one further away to emulate a typical layout in which many WarpLink devices might have been used.

Line Card and Test Card Layer Stackups

The stackups of the Test and Line Cards are presented in this paper to document the standard design practice used in the card layout and fabrication. Figure 6 and Figure 7 detail the Test Card and Line Card stackups, respectively. The gigabit differential nets of the Test Card and Line Card were designed to be the same. The length of the traces in the Test Card emulated the lengths seen on the Line Card.

The major difference between the Line and Test Card stackups is the addition of power planes to the Line Card. The Test Card has thicker dielectric material between layers 6 and 7 to account for the missing plane layers and maintain the same overall card thickness.

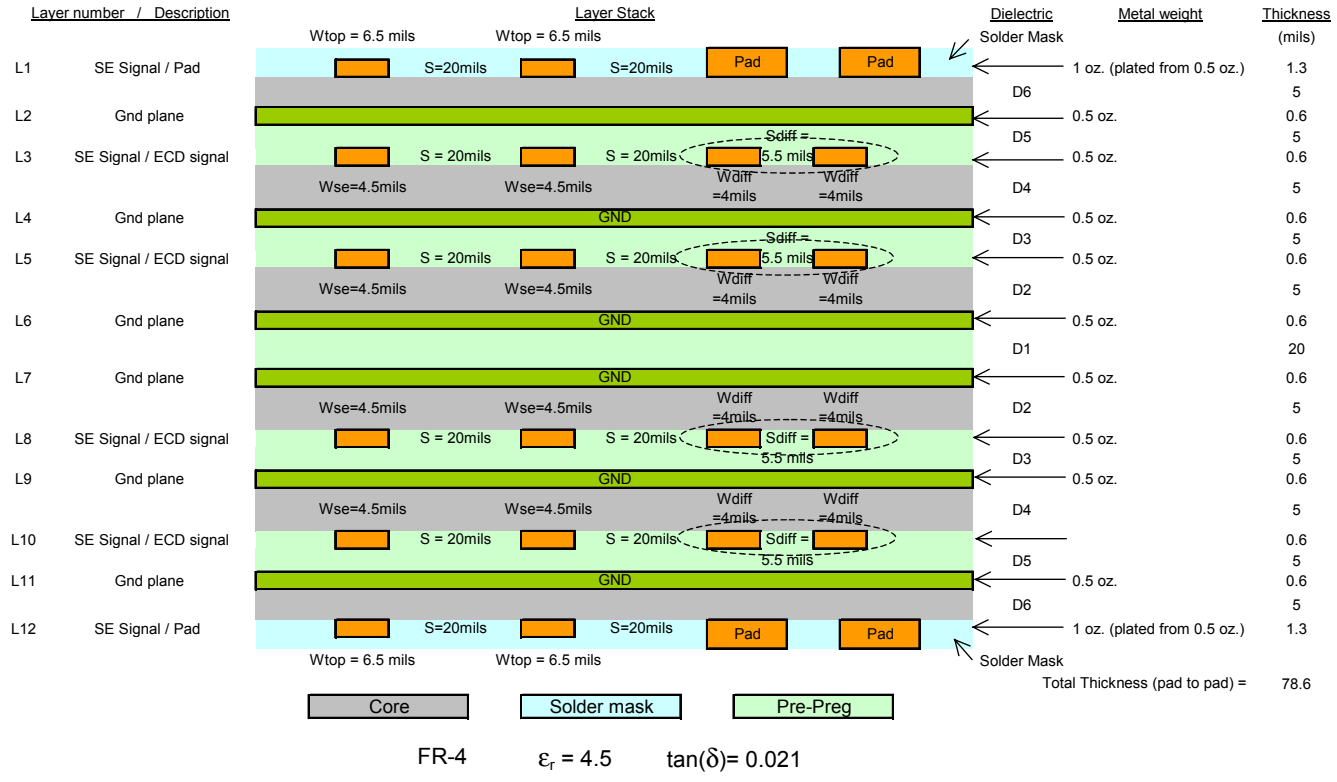

Figure 6. WarpLink Test Card Layer Stackup

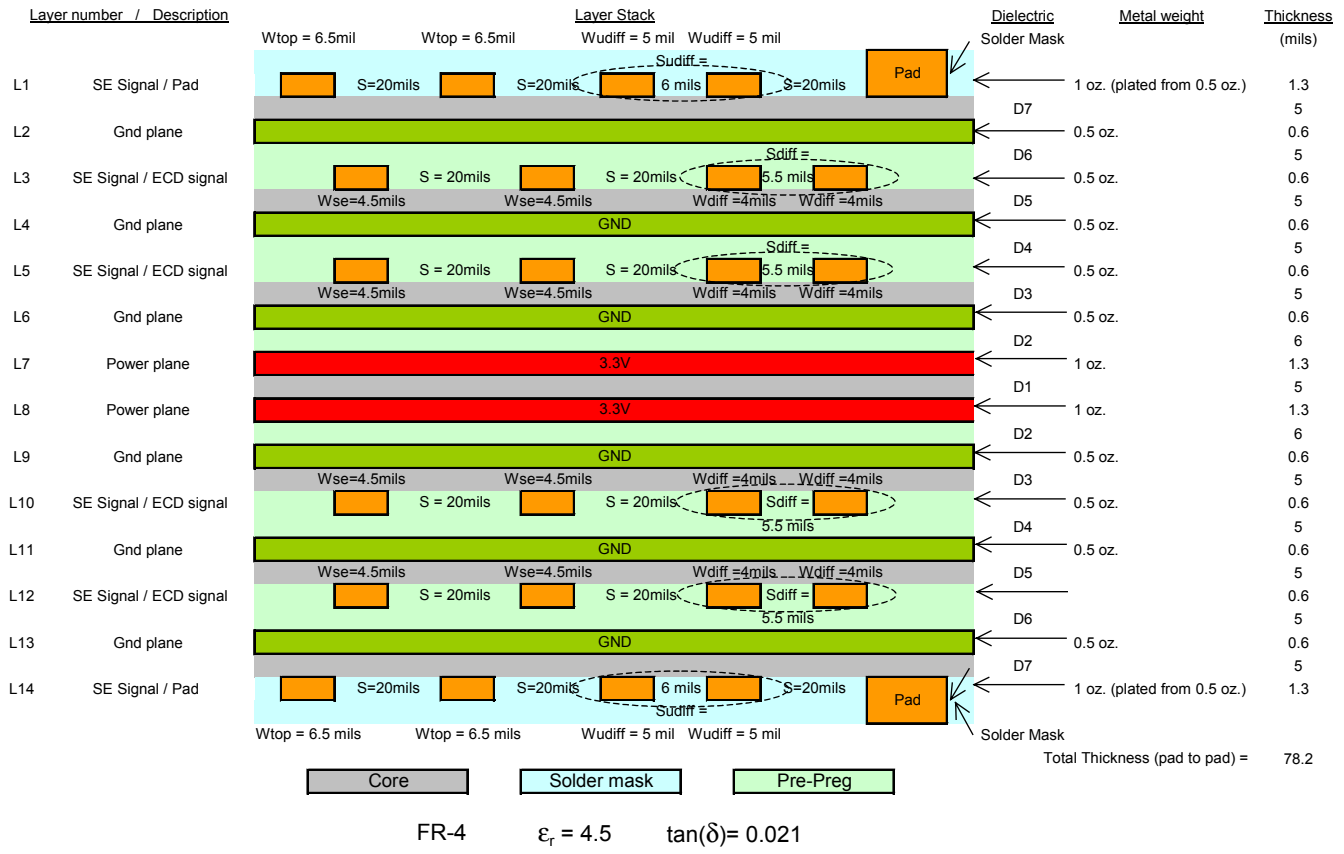

Figure 7. WarpLink Line Card Layer Stackup

WARPLINK SIGNAL INTEGRITY SIMULATION PROGRAM

WARPLINK GIGABIT SIMULATIONS

As part of the Motorola WarpLink Reference Design Platform effort, NESA completed an extensive HSPICE signal integrity study for the gigabit links and differential clocks. Some of the simulation model features include: coupled lossy-line models; worst-case card via effects due to stubs; WarpLink driver pre-emphasis control; and data rates running at 2.5 Gbd and 3.125 Gbd. Although the majority of the results focus on the differential interconnects for the active WarpLink device, simulation cases were completed for the clock interconnects as well. The data in this summary includes receiver input eye patterns as well as TDR impedance profiles of the various differentials interconnects.

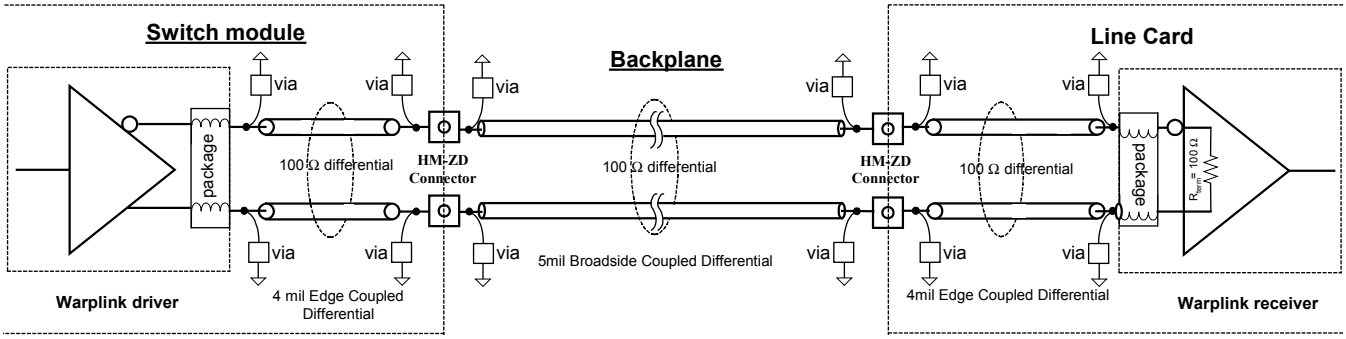

Figure 8. WarpLink Gigabit Interconnect

The Motorola WarpLink single-threaded HSPICE model was inserted as the driver and receiver into the interconnect environment constructed by NESA. Figure 8 details the interconnect topology for the gigabit links. Motorola provided the WarpLink device and package HSPICE models. Models for the ERmet® ZD connector were supplied by ERNI Components. The coupled, lossy transmission line and via models were created from the PCB parameters by NESA.

The construction of the HSPICE models was followed by the identification of the range of data rates and transmission line lengths needed to prove out the gigabit capabilities of WarpLink. Table 2 identifies the variables and ranges of the values for the simulation matrix. 0 lists the constants, which are fixed parameters within the HSPICE simulation matrix. Table 4 identifies the simulation matrix and the resultant eye pattern parameters. Note that simulations were run at 65°C to better emulate the case temperature (versus the typical 25°C).

Table 2. Simulation Matrix Values

Variables	Range
WarpLink Device	
Temperature	0°C, 25°C, and 100°C
Supply voltage (1.8 V ±10%)	1.6 V, 1.8 V, and 2.0 V
Driver pre-emphasis (transmit equalization)	off, on
Process	SS, TT, FF
Interconnect	
Backplane trace length	2, 10, and 20 inches
Line and switch card vias	One layer transition with stub and full board transition

Table 3. HSPICE Simulation Matrix Constants

Constants	Value
Device	
Gigabit Tx/Rx device	WarpLink 2 / WarpLink 2
Package	
324 PBGA (WL)	19 mm x 19 mm Body, 1 mm Ball pitch
Data	
Data rate	3.125 Gbps
Data pattern	K28.5
Etch	
Switch/mesh card trace length	6 inches
Line card trace length	4 inches
Board Parameters	
Driver location	Switch/mesh
Receiver location	Line
Backplane material	FR-4
Switch/uplink module material	FR-4
Line card material	FR-4
Backplane thickness	0.250 in
Switch/mesh card thickness	0.078 in
Line card thickness	0.078 in
Etch width (backplane) BC Diff	5 / 5 / 5 mils (w/s/w)
Etch width (line and switch) EC Diff	4 / 5.5 / 4 mils (w/s/w)
Switch/mesh card via count	2 - 1 @ package pad, 1 @ connector
Backplane via count	2 - 1 @ each connector
Line card via count	2 - 1 @ connector pad, 1 @ package
Backplane Connectors	
Gigabit nets - switch/mesh modules	8 - Row HM-ZD - GH Row
Gigabit Nets - line card	8 - Row HM-ZD - GH Row

Table 4. Simulation Matrix Eye Pattern Parameters

Case	Temperature (°C) 0 / 65 / 100	Supply Voltages ±10%	Process Corner	Pre-Emphasis On / Off	Data Rate (Gbps)	Backplane Length (in)	Via Type	Eye Height (mV)	Eye Width (ps)	Jitter (p-p) (ps)	Within XAUI Mask
1	0	2.0	FF	Off	3.125	2	Full	650	275	45	Y
2	0	2.0	FF	On	3.125	2	Full	740	280	40	Y
3	0	2.0	FF	Off	3.125	20	Full	300	220	100	Y
4	0	2.0	FF	On	3.125	20	Full	450	260	60	Y
5	100	1.6	SS	Off	3.125	2	Full	360	255	65	Y
6	100	1.6	SS	On	3.125	2	Full	450	270	50	Y
7	100	1.6	SS	Off	3.125	20	Full	170	220	100	N
8	100	1.6	SS	On	3.125	20	Full	220	240	80	Y
9	100	1.6	SS	Off	3.125	20	Stub	120	190	130	N
10	100	1.6	SS	On	3.125	20	Stub	190	240	80	N
11	65	1.8	TT	Off	3.125	10	Full	370	260	60	Y
12	65	1.8	TT	On	3.125	10	Full	480	280	40	Y
13	65	1.8	TT	Off	3.125	10	Stub	280	235	85	Y
14	65	1.8	TT	On	3.125	10	Stub	440	275	45	Y

The eye parameters measured at the receiver input include eye height, eye width, peak-to-peak jitter, and XAUI mask. The precise measurement point of the eye pattern is between the package and the silicon at the receiver input. A comparison of the odd-numbered cases (pre-emphasis off) in Table 4 to the even-numbered cases (pre-emphasis on) shows a great improvement for all eye metrics when pre-emphasis is engaged. The XAUI channel eye mask has been overlaid on the receiver input eye to illustrate compliance (refer to the eye pattern plots below).

Cases 9 and 10 and 13 and 14 contain the *stub* via model. In a typical backplane system, parasitic stubs are seen for vias where a signal traverse a via for less than the thickness of the backplane. For example, a differential pair routed on layers 5 and 6 will only pass through approximately 40 mil of the 250-mil backplane. The remaining 210 mils becomes a parasitic capacitive stub, reducing the trace impedance, and contributing to the degradation of the gigabit signal.

Figure 9 and 0 show examples of the simulated eye pattern for the WarpLink interconnect for cases 13 and 14. The measurement of the eye width and peak-to-peak jitter are made at the zero crossing of the pattern and the eye height is made at the center of the eye opening. Since this is a *zero average value* pattern, there should be no DC shift of the receiver input.

Figure 9 and 0 are typical of the improvement seen when the WarpLink transmit pre-emphasis feature is utilized. As shown in the simulation results of Table 4, the eye height was improved by an average of 26%, the eye width improved by 10%, and the peak-to-peak jitter improved by an average of 30%.

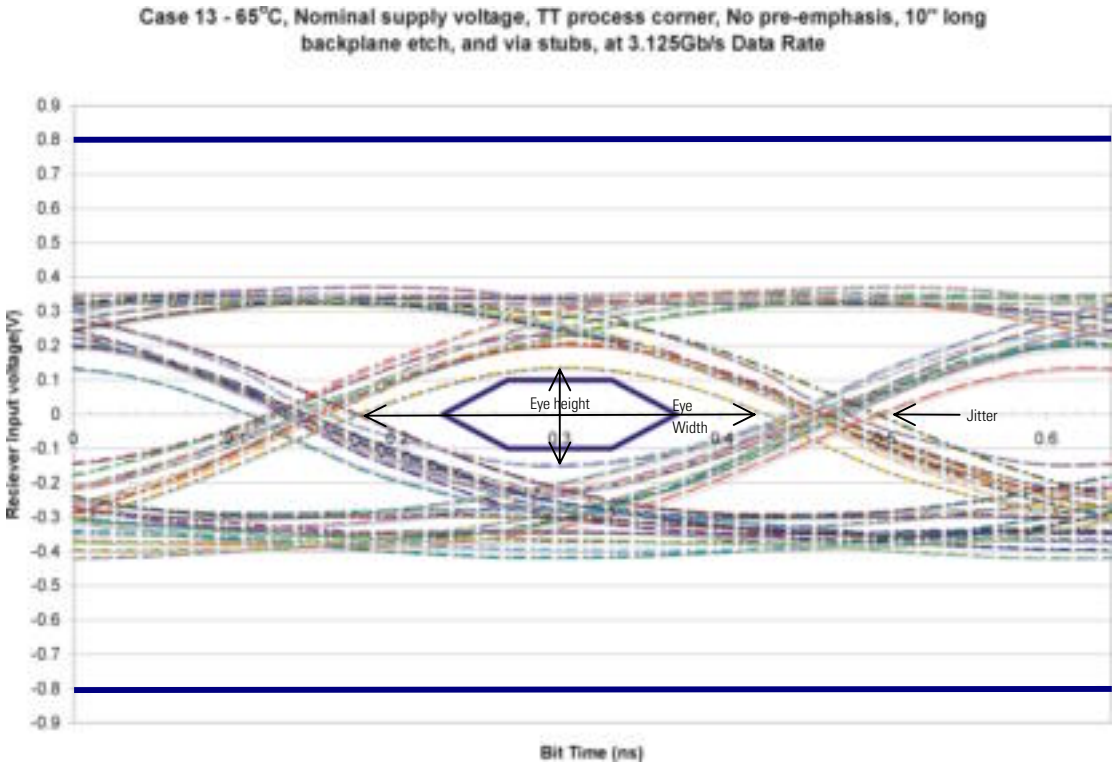

Figure 9. Case 13 with XAUI Mask Overlay

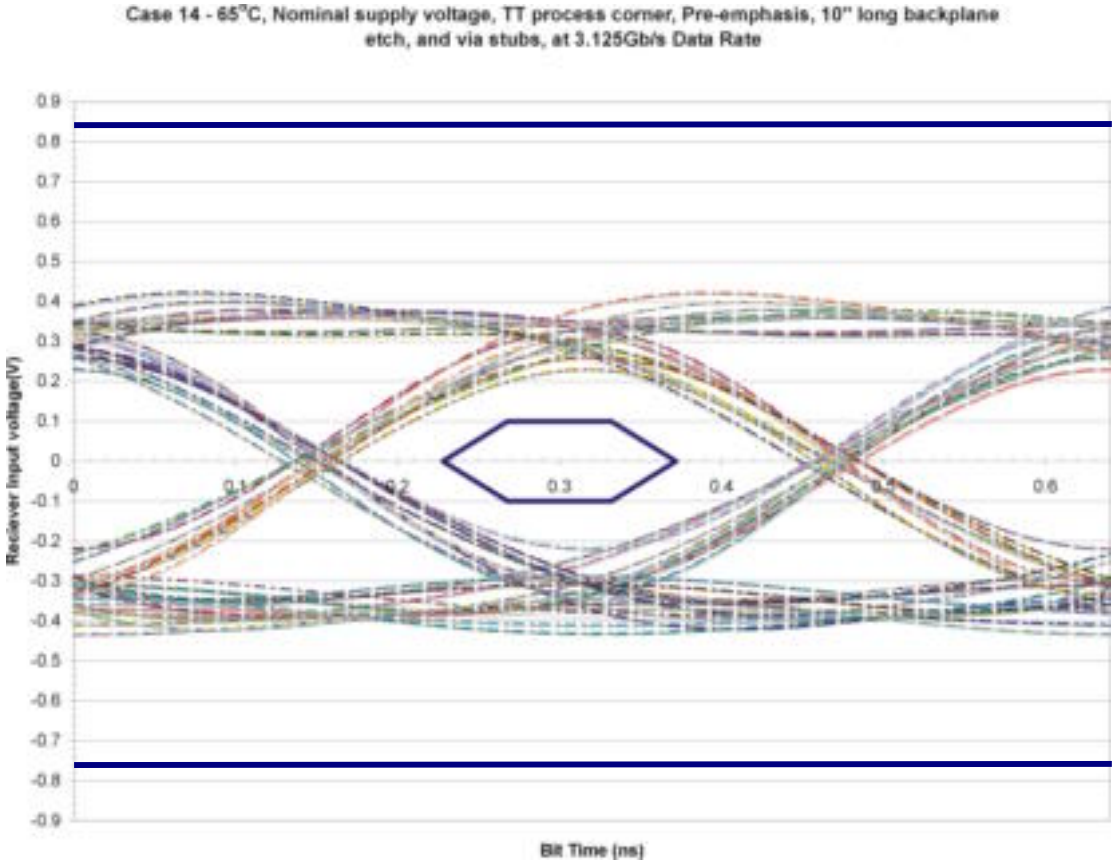

Figure 10. Case 14 with XAUI Mask Overlay

WARPLINK INTERCONNECT IMPEDANCE PROFILE

Time Domain Reflectometry (TDR) interconnect simulation shown in Figure 11 and Figure 12 illustrates the impact of a via stub. A differential TDR source (NESA Inc. proprietary simulation model) replaces the WarpLink driver while leaving the WarpLink package in the model. The receiver model is also removed and the interconnect is left open-circuited, which offers a clear indication of the end of the path.

Table 5. WarpLink Gigabit Interconnect Variations

Case	Backplane Length (in)	Via Type
1	2	Full
2	2	Stub
3	10	Full
4	10	Stub
5	20	Full
6	20	Stub
7	40	Full
8	40	Stub

Table 5 details the six variations of the WarpLink gigabit interconnect. Figure 11 and Figure 12 illustrate the difference of the impedance profile for a *stub* via and a *through* via, respectively.

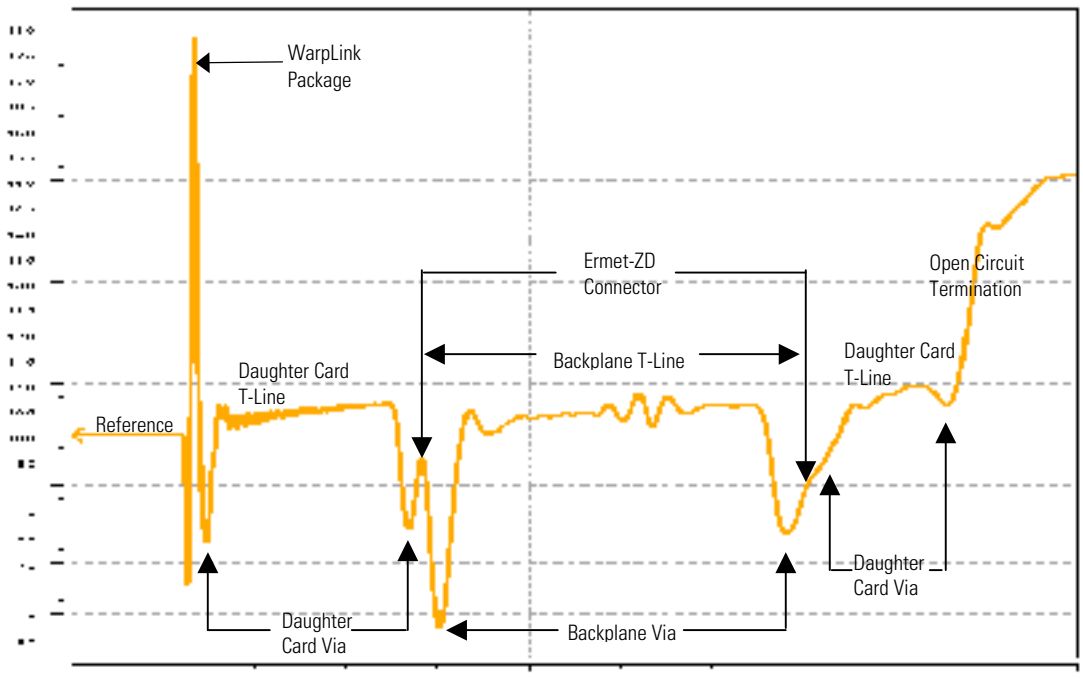

Figure 11. TDR of 10-inch Backplane with Via Stubs

The TDR highlights a number of issues for the interconnect. One of the goals of a high-frequency digital interconnect is to maintain as constant a characteristic impedance as possible. The vias and package impedance profiles of Figure 11 and Figure 12 show significant excursions from the reference impedance of 100 ohms.

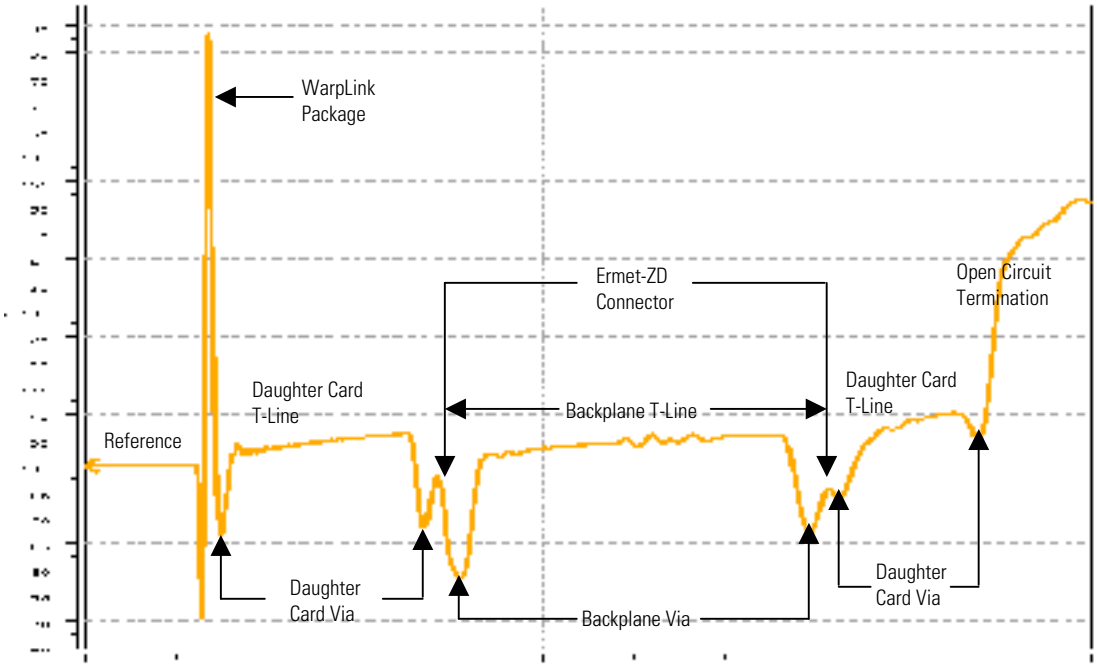

Figure 12. TDR of 10-inch Backplane with Through Vias

A comparison of the impedance profiles show that the *through* via produce smaller impedance discontinuity, Figure 12, than the stub via case, Figure 11. The package model presents a significant discontinuity in the simulation of the WarpLink interconnect, for both via cases. The large inductive spike is likely due to the lumped-element package model used in the simulations.

The discontinuity of the ERmet® ZD connector reflected its ability to transmit 5Gbps signals. The peak impedance level within the connector region was approximately 95-98 ohms. The effects of the vias on both the daughter card and backplane influence this level due to risetime roll-off of the TDR pulse. The lossy characteristic of the transmission line is evident as seen in the increasing impedance along the length of the TDR profile.

WARPLINK REFERENCE SYSTEM CLOCK SIMULATIONS

The clock architecture for the WarpLink Reference Design Platform is point-to-point. The clock operates at a digital frequency of 156.25 MHz. This data rate allowed the clock to be connected to the daughter card thought a standard 2 mm hard metric connector, the ERmet® 2 mm-HM. The interconnect path models used for the clock simulations are almost identical to that of the gigabit WarpLink models. The main differences were the replacement of the ERmet® ZD with the ERmet® 2 mm-HM connectors and the inclusion of the clock driver and receivers in place of the WarpLink devices.

Table 6. Simulation Matrix Values

Variables	Range
Clock Buffer Device	
Temperature	0°C, 25°C, and 100°C
Process	SS, TT, FF
Supply voltage (3.3 V ±10%)	3.0 V, 3.3 V, and 3.6 V
Interconnect	
Backplane trace length	2, 10, and 20 inches
Line card and switch card vias	One layer transition with stub and full board transition

The variables and range of the simulation matrix for the clock are shown in Table 6. The results of the simulations, exemplified by the waveform in Figure 13, show little distortion and a well-controlled clock signal.

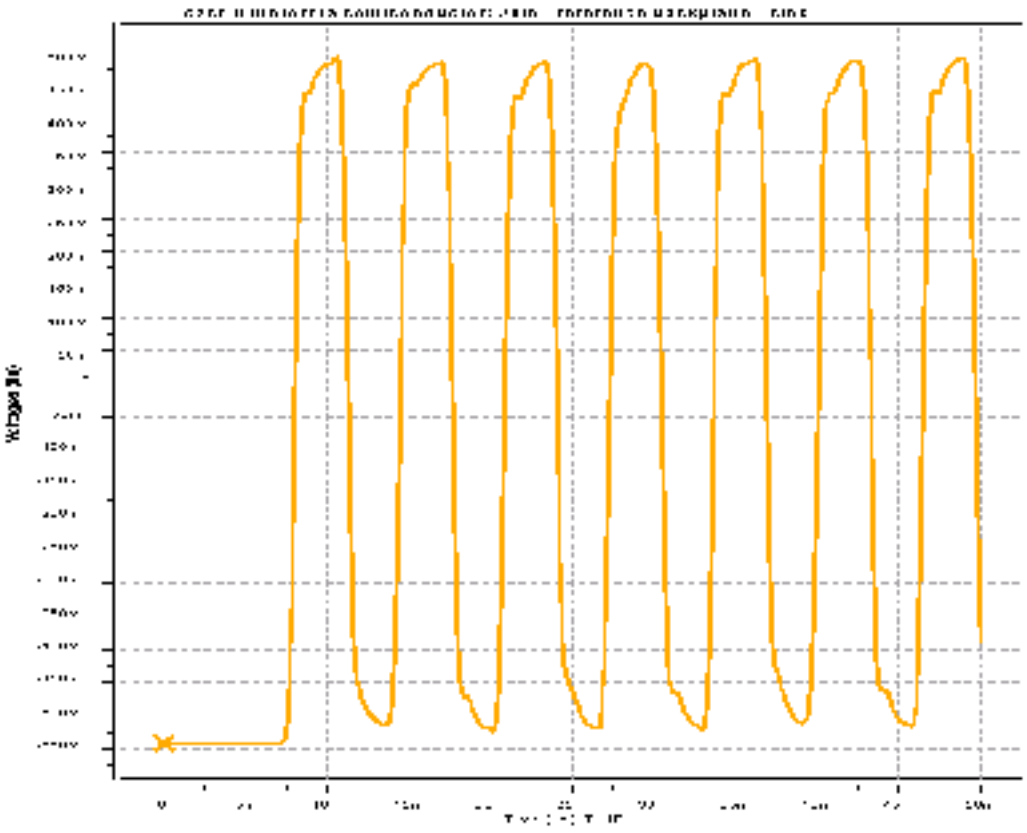

Figure 13. Typical Clock Receiver Input Simulation

DESCRIPTION OF PASSIVE SIGNAL INTEGRITY MEASUREMENTS

TIME DOMAIN REFLECTOMETRY

TDR was used to measure the impedance profile of the WarpLink Reference Design Platform's backplane interconnects. The impedance of an interconnect is important to determine since impedance mismatches due to surrounding structures can cause reflections, which in turn cause aberrations in signal fidelity. To measure the differential impedance, two signals of a pair are driven simultaneously with step sources of equal amplitude and opposite polarity. The far ends of the lines are either open circuit or match terminated and the reflections are sampled at the TDR heads. The impedance is then determined using both the incident and reflected wave components as an automatic function of TDR/DSO mainframe.

DIFFERENTIAL TIME DOMAIN CROSSTALK

Crosstalk measurements were completed using the TDR step generator to perturb an aggressor line while monitoring a victim line via the differential input channels. Crosstalk measurements are basically a form of transmission measurement, where instead of measuring at the output ports of a given excited pair, one measures at the near- or far-end ports of a neighboring pair. Any remaining open ports are terminated, ideally in their characteristic impedance. Such measurements will illustrate undesired signal coupling occurring between neighboring lines or differential pairs. The crosstalk can be expressed both in terms of the voltage amplitude as well as in percentage of the driver or receiver voltage.

Coupling the stimulus and measurement techniques of the crosstalk measurement with the signals used to generate eye patterns can produce some interesting effects. By using a clock aggressor signal and a standard PRBS signal through the victim line, distortion can be observed when the added signal crosstalk is present.

EYE DIAGRAMS

Eye diagrams can demonstrate the *goodness* of a component or connection path by passing a random bit stream through and monitoring the output over a span of time. Where the *hits* take place in time and voltage, and the distribution of the hits, can determine whether the device under test (DUT) is transmitting data that can be properly interpreted when compared to a *mask* or other criteria set in advance.

To measure an eye pattern across the WarpLink backplane, the HP8133A pulse generator was used in the PRBS23 data pattern mode, with the differential pulse voltage characteristics set to ± 500 mV level at 2.5 Gbps and 3.125 Gbps. Once the equipment was set up, the procedure consisted of introducing the bit pattern into a line or differential signal pair and monitoring the output at the far end on the HP54750A.

TIME DOMAIN TEST EQUIPMENT

The equipment used for the passive measurements included NESA's in-house Time Domain Reflectometer/Digital Storage Oscilloscope (TDR/DSO) and 3 GHz pulse generator with PRBS capabilities. A list of the hardware used to complete the relevant tasks follows:

Test Equipment

1. Hewlett-Packard (HP) 54750A DSO Mainframe
2. HP54754A 18 GHz Differential TDR/2 channel plug-in module
3. HP54751A 20 GHz 2 channel plug in module
4. HP8133A (w/option 2) 3 GHz Pulse Generator Fixturing
5. W. L. Gore Phaseflex Cables with APC 3.5 mm connectors (BW: DC-26.5 GHz)
6. SMA, and APC loads, adapters, probes, and miscellaneous fixturing

The TDR/DSO equipment required for these investigations had to be very capable in terms of bandwidth, accuracy, and software support. NESA's in-house HP54750A TDR/DSO is capable of completing differential TDR and Crosstalk, and can be used as a four-channel sampler for measurements such as bit-pattern data rate tests.

The HP8133A 3 GHz pulse generator was selected for its high bandwidth and many high-end features including two sets of differential outputs. The HP8133A can supply standard pulses, 32-bit pattern data sequences for data rate testing, and Pseudo-Random Bit Streams (PRBS), useful for such measurements as eye diagrams. This unit allowed for testing up to 3.125 Gbps. These measurements were collected and stored electronically using Hewlett-Packard's Benchlink Lightwave™ software.

PASSIVE MEASUREMENT RESULTS

A passive signal integrity study was performed as part of the design of the WarpLink Reference Design Platform. This study was performed in support of the active measurement and demonstration testing. The evaluation of the backplane establishes the quality of the impedance control of the PCB construction as well as setting a baseline for the eye pattern data before any equalization is used to recover the gigabit signal.

TDR RESULTS

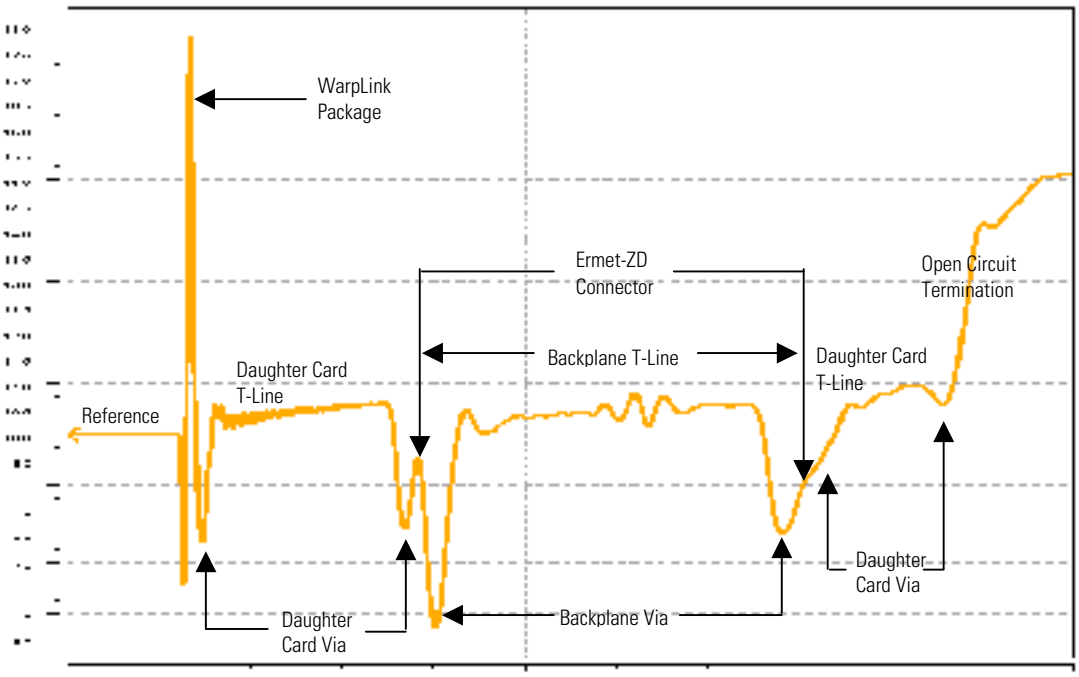

Figure 14. Simulated TDR of 10-Inch Backplane with Via Stubs

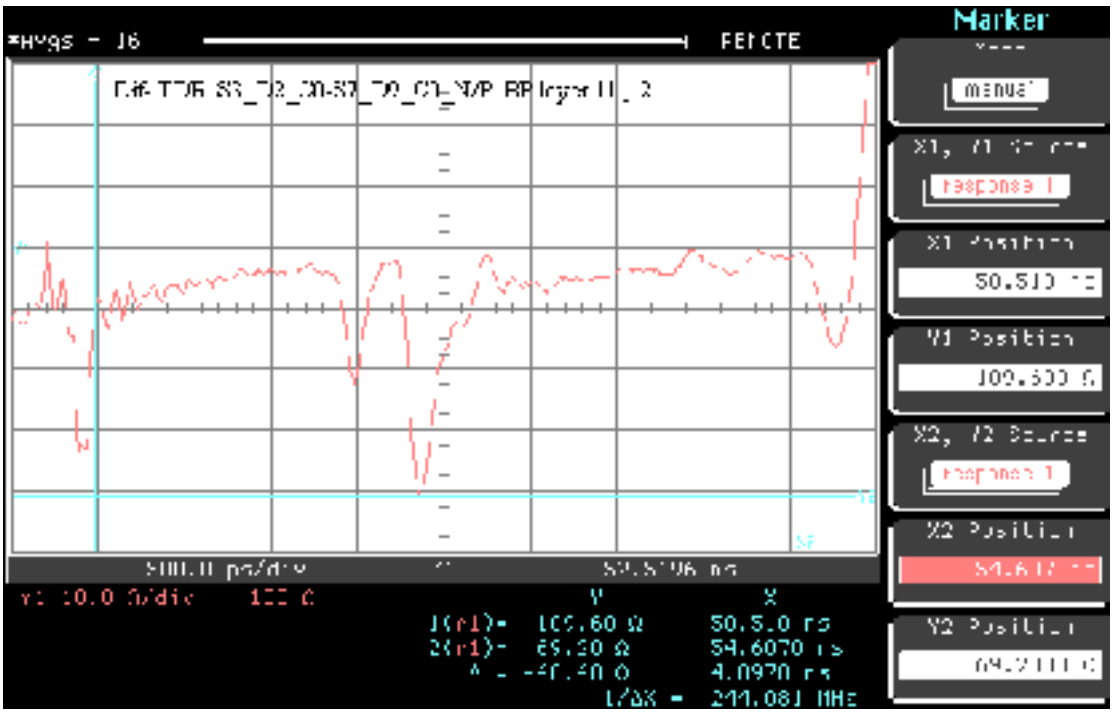

Figure 15. Measured TDR of 4.6-Inch Test Card and 6.9-Inch Backplane

Figure 14 and Figure 15 show the simulated and measured TDR impedance profiles for the WarpLink gigabit interconnect. The measured TDR profile was taken using the Test Card designed as part of this project, to access the channels on the backplane. The far end of the backplane interconnect was left open-circuited to ease detection. Although the impedance profiles are on different time scales the features of the interconnect do match. Using this type of evaluation we can confirm that the simulation result, shown in the section *WarpLink Signal Integrity Simulation Program*, uses an interconnect model that is representative of the physical interconnect. Note that the actual measurement does not include the package or the second daughter card in the impedance profile.

The TDR measurement was also used to confirm the quality of the PCB fabrication. The broadside couple differential interconnects routes were specified to be 100 ohms ±10%. This is the case for the WarpLink Reference Backplane and Test Card. The model and the measurement also similarly model the increasing impedance along the length of the interconnect due to accumulated losses.

EYE DIAGRAM MEASUREMENT RESULTS

As part of the passive signal-integrity study, eye diagrams were taken for a variety of route lengths on the WarpLink backplane. Table 7 contains a full set of data measurements for both 2.5 Gbd as well as 3.125 Gbd. As expected, the longer the interconnect got and the higher the data rate got, the smaller the eye opening got. These results proved the need for equalization in gigabit signaling over long interconnects.

Table 7. 2.5 / 3.125 Gbd Data Measurements

Net Name	Total Length (in)	Data Rate (Gbps)	Eye Height (mV)	Eye Width (ps)	Jitter (p – p) (ps)
Reference Signal	Cables Only	2.5	850.0	358.0	40.7
Reference Signal	Cables Only	3.125	840.0	255.5	61.1
S4_D1_D0-S4_D1_D0_[N|P]	8.1	2.5	537.8	295.8	101.5
S4_D1_D0-S4_D1_D0_[N|P]	8.1	3.125	480.2	183.1	133.5
S7_D10_D0-S8_D10_C0_[N|P]	11.3	2.5	501.3	291.6	107.1
S7_D10_D0-S8_D10_C0_[N|P]	11.3	3.125	342.3	215.0	105.0
S2_D1_C0-S3_D1_B0_[N|P]	18	2.5	376.9	252.3	145.9
S2_D1_C0-S3_D1_B0_[N|P]	18	3.125	163.2	170.0	150.0
S1_D2_C0-S7_D9_A0_[N|P]	29.0	2.5	253.0	185.5	212.7
S1_D2_C0-S7_D9_A0_[N|P]	29.0	3.125	39.8	83.0	237.0

Figure 16(b), demonstrates the distortion seen for the differential pair S4_D1_D0-S4_D1_D0_P|N, a moderately long interconnect, 8.1 inches in total length, when compared to a reference path measurement, shown in Figure 16(a). The direct comparison allows us to attribute the PCB interconnect of the WarpLink backplane and Test Card as the cause of the distortion in the eye opening.

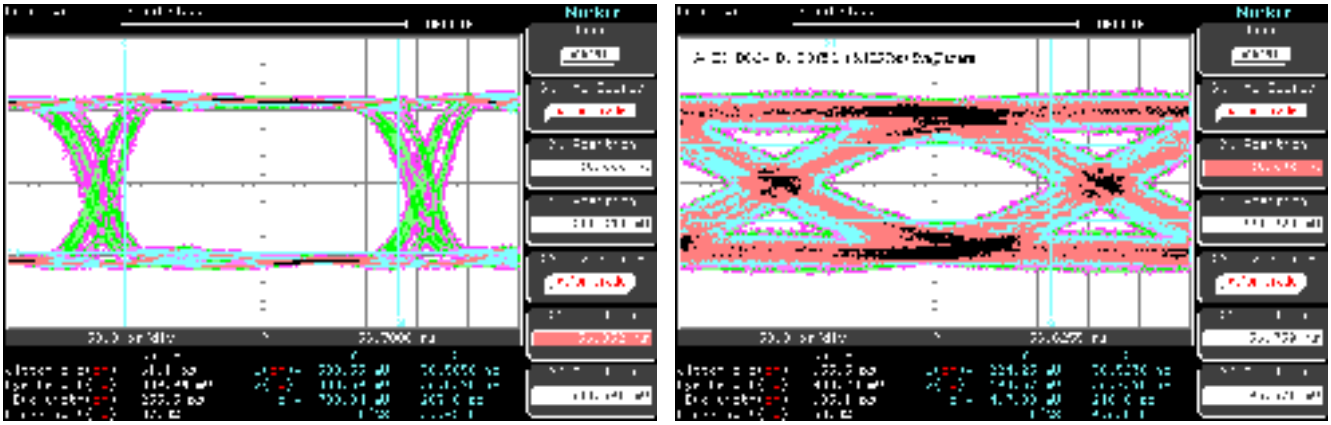

(a) Reference Path **(b) S4_D1_D0-S4_D1_D0_[P|N]**

Figure 16. Passive EYE Patterns

ACTIVE MEASUREMENT RESULTS

While the passive measurements quantify the quality of the passive transmission channels, the active measurements show how WarpLink 2.5 Quad performs over the same transmission channels. Active eye patterns driven by a WarpLink transmitter through a variety of transmission channels of the system's backplane were obtained. XAUI compliance was confirmed by overlaying the measured eye patterns to the eye mask defined in the XAUI specification. The following sections describe the test setup and the various channels chosen for the active measurements presented in the paper.

TEST SETUP

Eye diagrams from two different transmission channels are presented. The transmission channel connecting slot 8 and slot 1 was chosen as an example of the longest, and thus worst case, backplane channel that was not artificially lengthened. The second transmission channel chosen connects slot 7 and slot 1. This channel was artificially lengthened to present a condition that exceeds an interconnect length defined by the XAUI standard.

A Line Card was installed in either slot 7 or 8. A Test Card from which the eye diagrams were obtained was in slot 1. The transmitter was configured for the Built-in Self-Test (BIST) mode, configured for PN Equation 2 of the 23-bit PN generator (please refer to the WarpLink 2.5 Quad Product Brief for more details of the BIST mode). The full data path from a WarpLink transmitter on a Line Card in slots 7 or 8 to the Test Card in slot 1 consisted of:

1. WarpLink Device
2. Package via
3. Line Card FR-4 trace
4. Line Card Connector Via
5. Line Card - Backplane ERmet® ZD Connector
6. Backplane Connector Via
7. Backplane FR-4 trace
8. Backplane Connector Via
9. Backplane Connector
10. Test Card - Backplane ERmet® ZD Connector
11. Test Card FR-4 trace
12. SMA Via
13. SMT - SMA Connector
14. SMA Right-Angle adapter
15. 36 inches of Rosenberger rf-coaxial cable
16. AC-coupling SMA adapter
17. TDS-8000 scope.

The WarpLink device is clocked via the clock distribution networks originating from slot 7. (Please refer to the Clock Distribution section for more details). The clock trigger of the TDS-8000 is also clocked with the same clock distribution network, by way of the Test Card in slot 1.

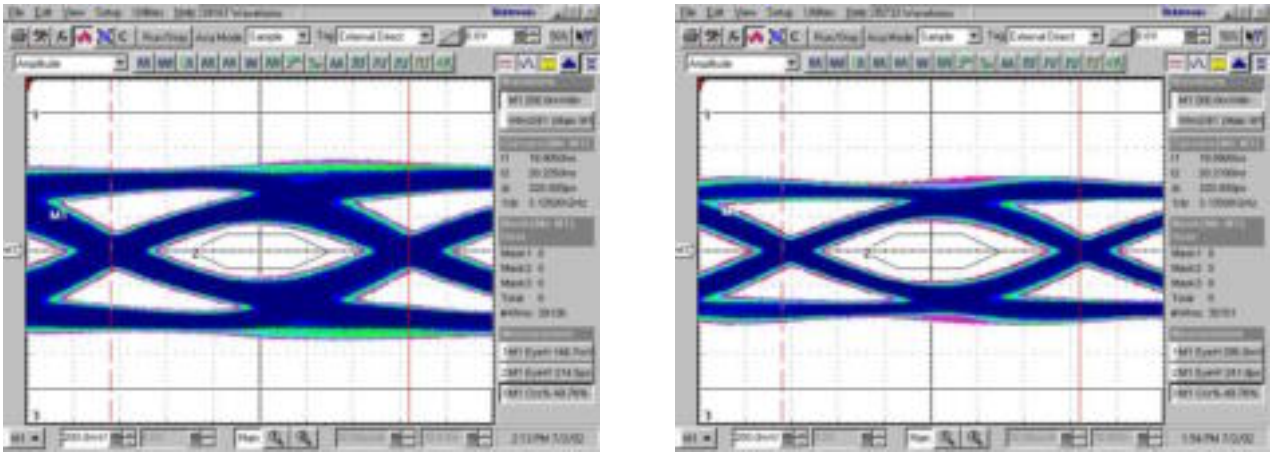

(a) Without Equalization **(b) With Equalization**

Figure 17. Active Eye Patterns Tx – Slot 8 to Rx – Slot 1

EYE DIAGRAMS FROM SLOT 8 TO SLOT 1

The interconnect between slots 8 and 1 are the longest channel in the 19" rack WarpLink Design Reference system that is not artificially lengthened. The total length of the transmission channel in FR-4 is 18.7". This includes 1.8" of FR-4 on the Line Card, 12.2" on the backplane, and 4.7" on the Test Card. Without pre-emphasis/equalization, WarpLink exceeds the XAUI mask with a 5% margin. With pre-emphasis/equalization, the WarpLink gigabit link satisfies the XAUI mask with a 9% margin. The comparison of Figure 17(a) and Figure 17(b) shows a clear benefit of pre-emphasis/equalization to the signal fidelity of the eye diagram.

EYE DIAGRAMS FROM SLOT 7 TO SLOT 1

The artificially lengthened interconnect between slots 7 and 1 presents a condition that is worse than that defined by the XAUI standard. This is a worst case routing condition that exemplifies a routing implementation on a large backplane. The total length of the transmission channel in FR-4 is 26.2". This includes 1.8" of FR-4 on the Line Card, 19.7" on the backplane, and 4.7" on the Test Card. With pre-emphasis/equalization on, WarpLink passes the XAUI mask with a 3% margin. As expected via simulation the resultant eye does not conform to the XAUI mask with pre-emphasis/equalization disabled. Again, the comparison of Figure 18(a) and Figure 18(b) shows a clear benefit of pre-emphasis/equalization to the signal fidelity of the eye diagram.

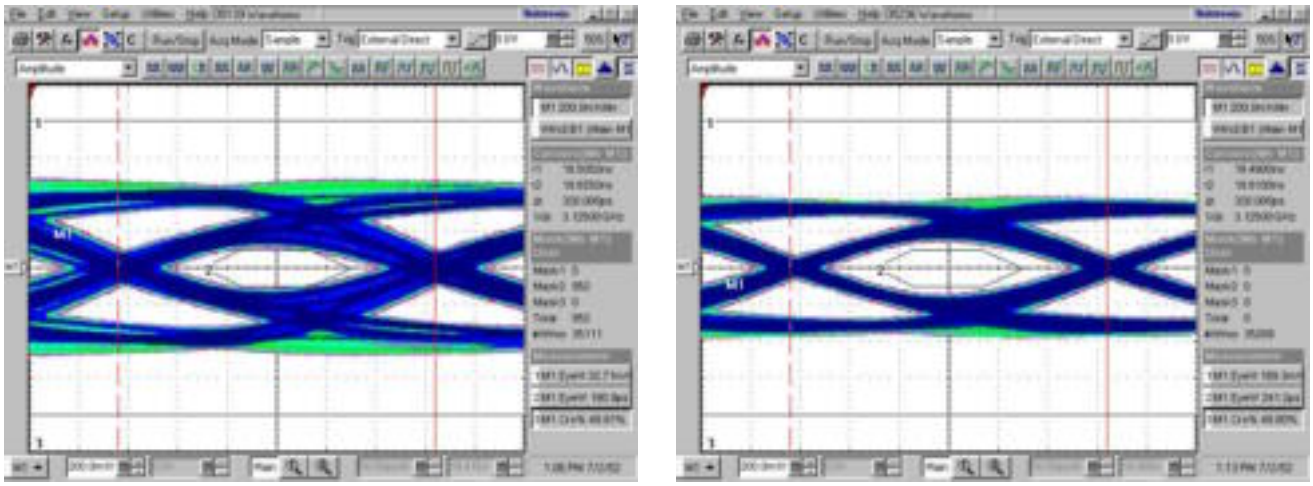

(a) Without Equalization **(b) With Equalization**

Figure 18. Active EYE Patterns Tx – Slot 7 to Rx –Slot 1

SUMMARY AND CONCLUSIONS

The Motorola WarpLink Reference Design Platform demonstrates the capabilities of the WarpLink 2.5 Quad SERDES device. The device was operated at wire speed of 3.125 Gbd to provide an aggregate speed of 10Gbps per device. The device's pre-emphasis/equalization feature proved indispensable for data transfer at gigabit wire speeds over long interconnects. In particular, WarpLink 2.5 Quad has demonstrated the ability to operate at, and beyond the XAUI specification of 3.125 Gbd over 50cm long channels.

The WarpLink Reference Design Platform is built with common design practices, components, and PCB material (FR-4). Unrealistic and expensive gigabit design requirements - such as blind vias, via drill-back, bottom-layer-only routing, and low dielectric constant PCB materials - were not required.

Using a fully matured simulation and design methodology, NESA implemented a full end-to-end gigabit simulation matrix of cases. The currently available SERDES model for WarpLink 2.5 Quad was verified to emulate actual measured results. Using this modeling strategy, areas of risk for the gigabit links were identified and properly addressed while still in the design phase.

The Motorola WarpLink Reference Design Platform has shown that 3.125 Gbd SERDES technology is a reality, with readily available support. Proven Reference Design Platforms and confirmed WarpLink I/O simulation models are available from Motorola.

ACKNOWLEDGEMENTS

This white paper was prepared as a joint effort by Quang Xuan Nguyen of the Motorola WarpLink design team and Dr. Edward Sayer 3rd from North East Systems Associates Inc. (NESA, www.nesa.com). NESA was contracted by Motorola to do the design, perform simulations and define the layout rules for the reference backplane. Their expertise lies in the area of the practical design and manufacturing techniques required for high performance systems and backplanes.

Additional thanks go to Michael Baxter, NESA and Thecla Chomicz, Motorola for their technical expertise and assistance in performing the extensive passive and active measurements presented in this paper.

NOTES

HOW TO REACH US:

USA / EUROPE / Locations Not Listed:

Motorola Literature Distribution
P.O. Box 5405
Denver, Colorado 80217
1-303-675-2140
1-800-441-2447

JAPAN:

Motorola Japan Ltd.
SPS, Technical Information Center
3-20-1, Minami-Azabu, Minato-ku
Tokyo, 106-8573 Japan
81-3-3440-3569

ASIA/PACIFIC:

Motorola Semiconductors H.K. Ltd.
Silicon Harbour Centre
2 Dai King Street
Tai Po Industrial Estate
Tai Po, N.T. Hong Kong
852-26668334

TECHNICAL INFORMATION CENTER:
1-800-521-6274

HOME PAGE:
http://www.motorola.com/semiconductors/

BR1570/D, REV 0

Freescale Semiconductor, Inc.