TEST AND VERIFICATION REPORT

HM Gbps End-to-End Test Results and Design Recommendations

TVR Document: **A25022N0044**

Revision: DRAFT

B. Carlson, August 5, 2005

National Research Council Canada
Herzberg Institute of Astrophysics
Dominion Radio Astrophysical Observatory

P.O. Box 248, 717 White Lake Rd
Penticton, B.C., Canada
V2A 6J9

Table of Contents

List of Figures

1 Revision History

Revision	Date	Changes/Notes	Author
DRAFT	Aug. 5, 2005	Initial release	B. Carlson

2 Introduction

This document describes testing activities and test results of the HM Gbps end-to-end test. This test involves using the Timecode Generator Board (TGB) prototype with its Stratix GX FPGA and 1.024 Gbps transmit and receive capability, and the prototype Fanout Board to perform an end-to-end test of the HM Gbps transmission system. The HM Gbps transmission system forms the backbone of the station-to-baseline data transmission system in the correlator. Details of the physical attributes of this system are contained in [1], and details of the signaling protocol are contained in [2].

The basic test setup is shown in the following figure:

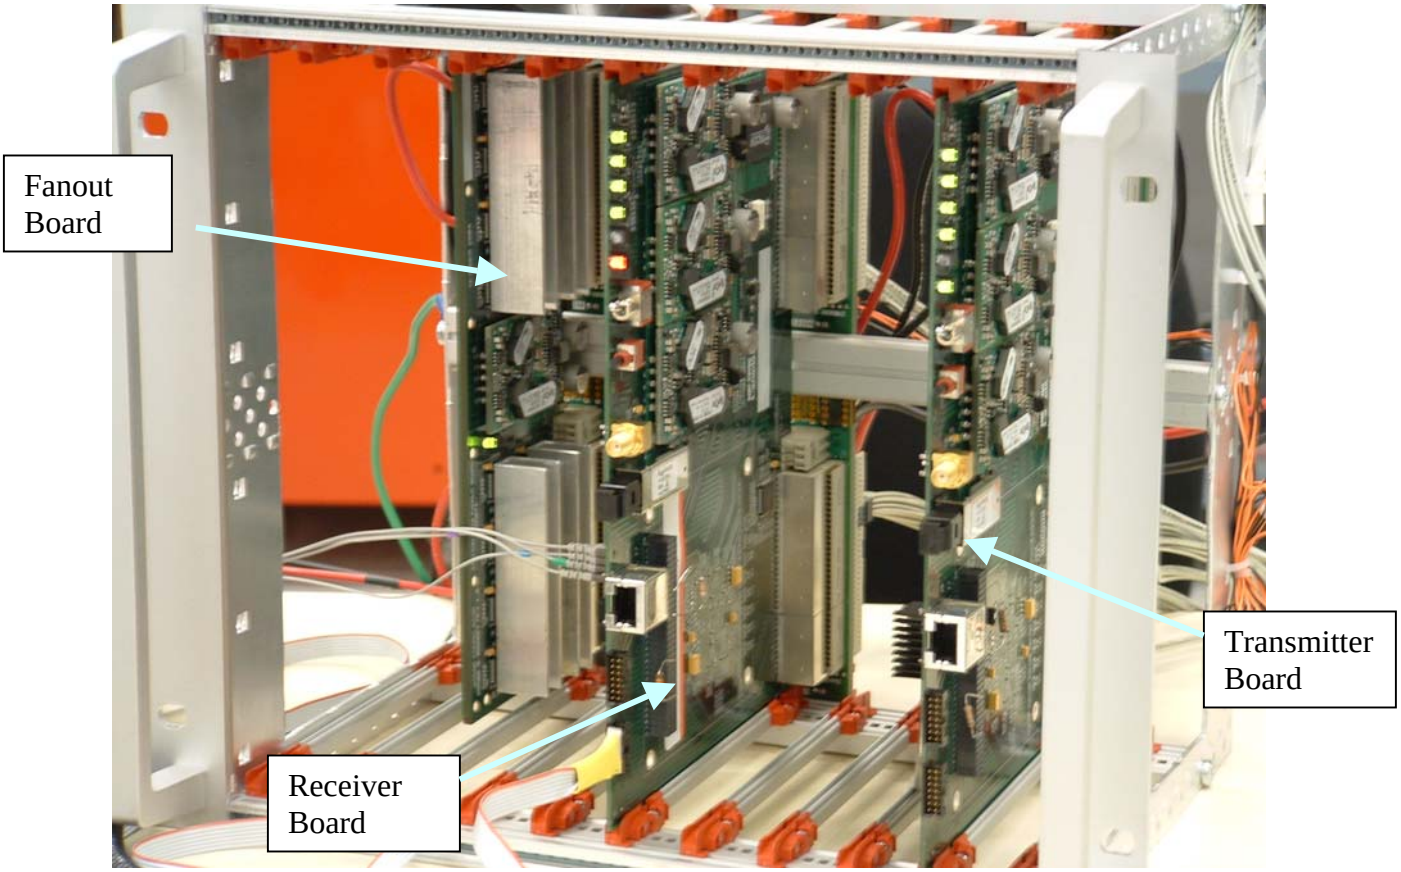

Figure 1 HM Gbps end-to-end test setup.

Due to clocking limitations in the FPGA, a separate transmitter TGB and receiver TGB were required. Both boards are the same design, but the FPGA program in each is different—one for 1.024 Gbps transmitter functions, and one for 1.024 Gbps receiver functions with Dynamic Phase Alignment. The receiver FPGA code is identical to the receiver code in the Recirculation Controller FPGA [3]. Details on the design and functionality of the transmit and receive FPGAs for this test can be found in [4].

Tests were run with and without the Fanout Board in place. A simplified diagram of the main configurations that were tested is shown in Figure 2. Note that each data path consists of multiple differential pairs and each transmission line shown in the figure represents multiple concatenated transmission lines (e.g. PCB—connector—cable—connector—PCB).

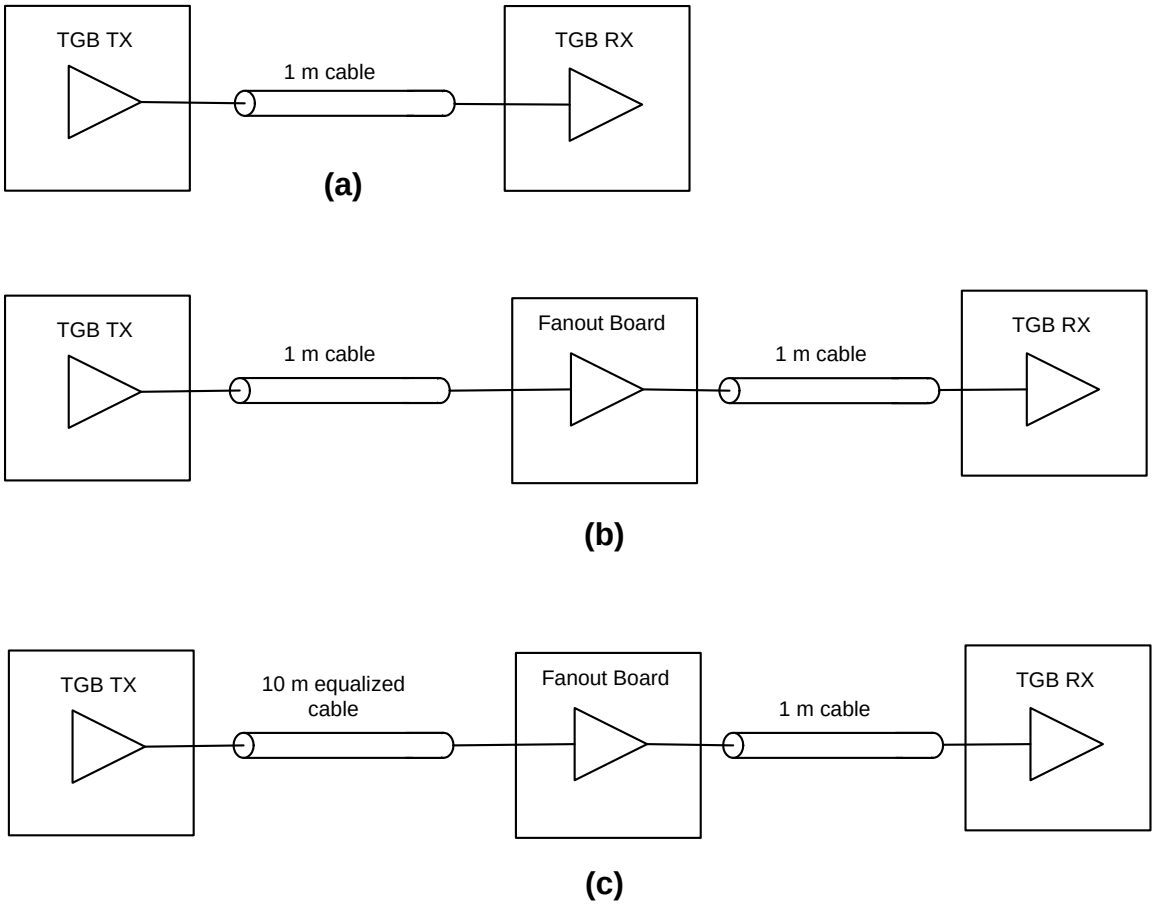

Figure 2 Gbit test configurations. (a) is TGB transmit to receive board only and is the simplest test. (b) includes the Fanout Board with a short cable from the transmitter. (c) includes the Fanout Board with a 10 m equalized cable, and is the most important test since it represents the actual worst-case cable configuration in the correlator.

Only 1 wafer (4 differential pairs) of the 10 m equalized cable is available for testing and so the Fanout Board is used to produce 4 wafers of data to go to the TGB receiver FPGA. The FPGA design is such that it requires only 1 wafer input to work—any unused inputs that are not toggling are ignored by the receiver[1].

[1] Although, on a differential pair if the input is not connected the output can still be toggling from noise, and this prevented full use of the functionality of the receiver to ignore inputs that are not toggling.

The transmit FPGA drives the on-board LED **green** if there is a clock and the PLL is locked. The clock for the transmit FPGA is provided by an off-chip, on-board, 128 MHz oscillator.

The receive FPGA drives the on-board LED **red** if it is receiving a clock and the PLL is locked, but the receiver has not acquired synchronization. The LED is **green** if the receiver has acquired synchronization and is not detecting errors. The LED is **yellow** if the receiver acquired synchronization but is currently detecting errors. Any time an error is detected, the yellow condition is driven for at least 0.25 seconds, and so it is possible to easily see synchronization errors, without software polling. If the receive LED is **off**, then there is no receive clock and/or the receive PLL is not locked.

For this test, the embedded PC/104+ and PCMC modules are not installed—the FPGAs are programmed from a PC using the Altera Quartus software, and a download cable via the TGB's JTAG interface. The four test pins that connect to FPGA pins were used extensively to probe chip internal signals.

To test PCB transmission line performance and attenuation, one of the 1.024 Gbps signals on one of the wafers follows a 17" (~2.9 nsec) serpentine path from the connector to the receiver FPGA. This also exercises the capability of the receiver FPGA to properly acquire synchronization on this long-path signal both in terms of DPA performance, and user logic that is able to absorb/compensate for +/-3 bits (~+/-12 nsec) of pair-to-pair delay within the same cable bundle.

3 Test Results and Design Recommendations

The tests were only partially successful, and clearly indicate that more conservative designs in the 1 Gbps system are necessary. The Stratix GX FPGA transmit and receive functions, with dynamic phase alignment were shown to successfully work as designed. The Meritec cable performed very well. It is a matter of getting the Gbps signals from the transmitter to the receiver with sufficient integrity including all PCB and connector effects that forms the focus of this report.

A number of eye diagrams at different points in the system were obtained using the 7 GHz, 20 Gs/s Tektronix TDS7704B digital storage oscilloscope with differential probe. These eye diagrams indicate why the Gbps signals failed to get to the receivers with sufficient signal integrity in many configurations.

A number of test configurations were devised that were not envisioned at the start of testing. These test configurations were used to nail down what works and what doesn't work and form the foundation of the design recommendations stemming from this report.

3.1 Design Recommendations

The design recommendations are as follows. Careful attention must be paid to every single segment of the 1.024 Gbps transmission line. We cannot afford to give up any signal integrity if at all possible. Failure to follow these design recommendations will likely result in a flaky and unreliable data transport system.

- The transmit chips and lines must be as close to the output connector as possible. The number of vias and impedance discontinuities must be absolutely minimized. If possible, lump the impedance discontinuities as close to the pins of the transmit chip as possible.

- Although not a result of these tests, DC-blocking capacitors on the outputs as well as the inputs must be used to avoid transmitter destruction from inadvertent cable short circuits. Locate these as close as possible to the device transmit and receive pins.

- Use flow-through, split-termination, where on-chip receiver differential termination is not possible (such as in the Motorola fanout buffers) instead of "conventional differential termination". This method has the smallest stub/impedance discontinuity. Flow-through termination vs conventional termination is shown in Figure 3. Split termination (50 ohm from each line to Vbb) has the advantage that there is a 25 ohm common-mode impedance to ground thus attenuating common-mode voltages. In the differential termination scheme common-mode signals see high impedance and would not be attenuated. There should be a 0.01 μF capacitor at the split termination as shown in the figure.

•

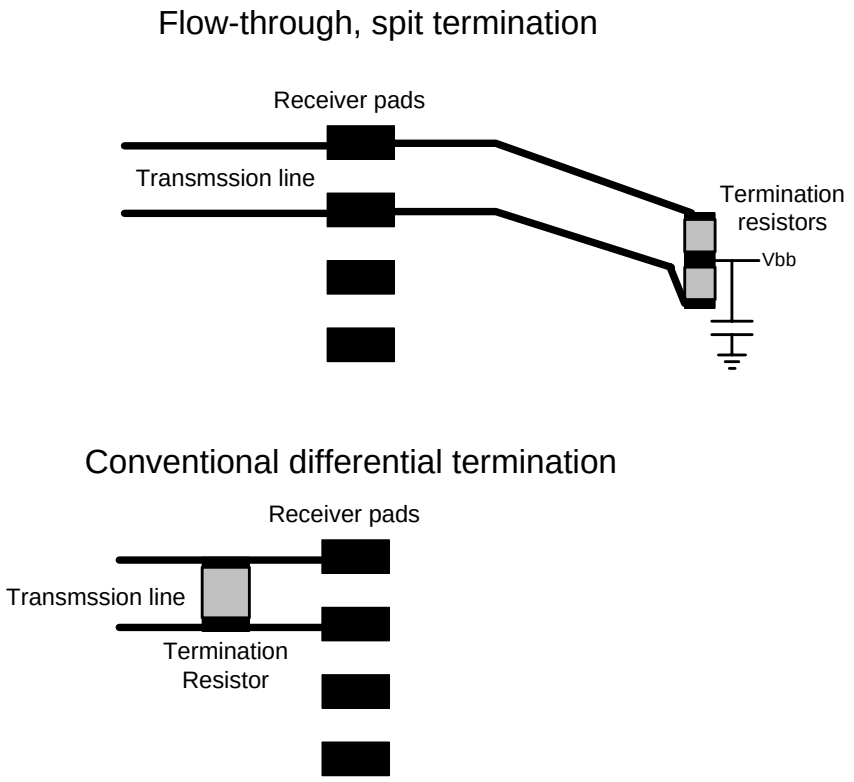

Figure 3 Flow-through termination and conventional termination.

Additionally, to prevent receivers from oscillating when there is no input, a 2 k resistor to ground on the positive input of the differential receiver will introduce a 30 mV bias if Vbb=1.2 V and for a 50 ohm termination resistor. This resister should be co-located with the termination resistor.

- The Fanout Board must be equipped with a re-timing FPGA. This can be the same FPGA as the Recirculation Controller FPGA on the Baseline Board—a Stratix-II EP2S15 in a 672-pin package. Based on the test results in this report it is evident that this re-timing FPGA is essential to obtain reliable end-to-end data transfer. This FPGA will be programmed by an on-board EEPROM (with in-system update interface via the front panel) and will operate without CPU interaction. It will remove jitter from the clock and the signal, and use automatic dynamic phase alignment to regenerate the signal. There are 256 Fanout Boards in the system, and one FPGA is required for each board. At ~$87 for each FPGA, it is a small price to pay for reliable operation.

- Before the re-timing FPGA on the Fanout Board, use high-speed buffers to clean-up the signal before entering the FPGA. The lines from the connector to the buffers must be absolutely minimum in length; impedance discontinuities must be absolutely minimized. Use the flow-through split termination on these inputs.

Use the FPGA on-chip termination to minimize components from the buffers to the FPGA.

- When routing high-speed lines from the connector pins to the receivers, use the bottom stripline layer to minimize stubs at the end of the connector. This is shown in Figure 4.

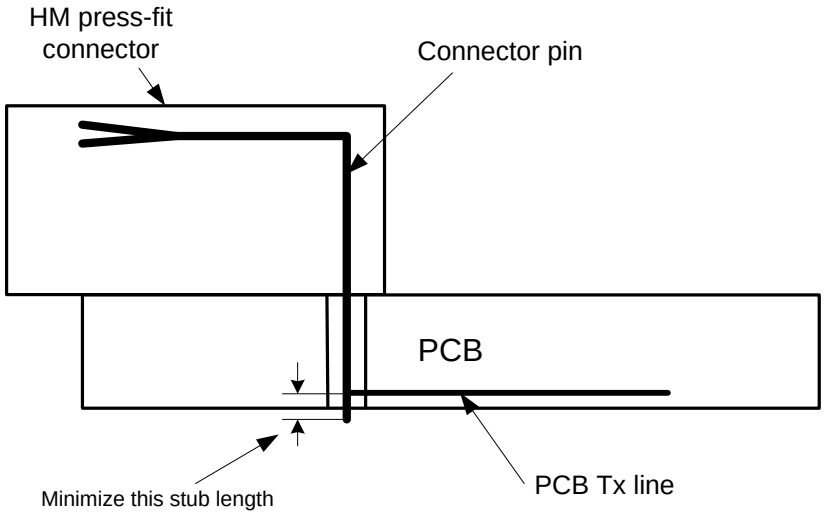

Figure 4 Connector pin-to-PCB Tx line transition for minimum stub length.

- Put signal receive buffers on the Baseline Board as close to the connectors as possible. This will minimize the PCB trace length driven after the cable that, according to Figure 7 has a severe detrimental effect on the signal. Use flow-through split terminations. For the FPGA receivers on the Baseline Board, use on-chip differential terminations. If possible, install series resistors right at the output of the buffer chips so that some mitigation of board impedance problems can be performed if necessary.

- Use 0.006" traces for the 1 Gbps differential pairs. There is some thought that the impedance problems on the TGB were due to smaller traces and therefore larger relative tolerances (i.e. the impedance of the Fanout Board traces were fine and they were 0.006" traces, with 0.009" spacing).

- In all cases, lump impedance discontinuities as close to the receiver pins as possible, and absolutely minimize the number of vias used.

3.2 Test Results

This sub-section contains a number of eye diagrams of the system under test. Due to the difficulty in probing points of interest, it is necessary to devise various transmission line configurations to demonstrate signal integrity at various points in the system. The transmission line configuration is described in each diagram caption. For all test plots shown, the Tektronix TDS7704B 20 Gs/s digital storage oscilloscope and 7 GHz

bandwidth differential probe was used. Before the signal leaves the transmitter board, it has already traveled over ~4 inches of FR-4 PCB and through 1 connector and backplane.

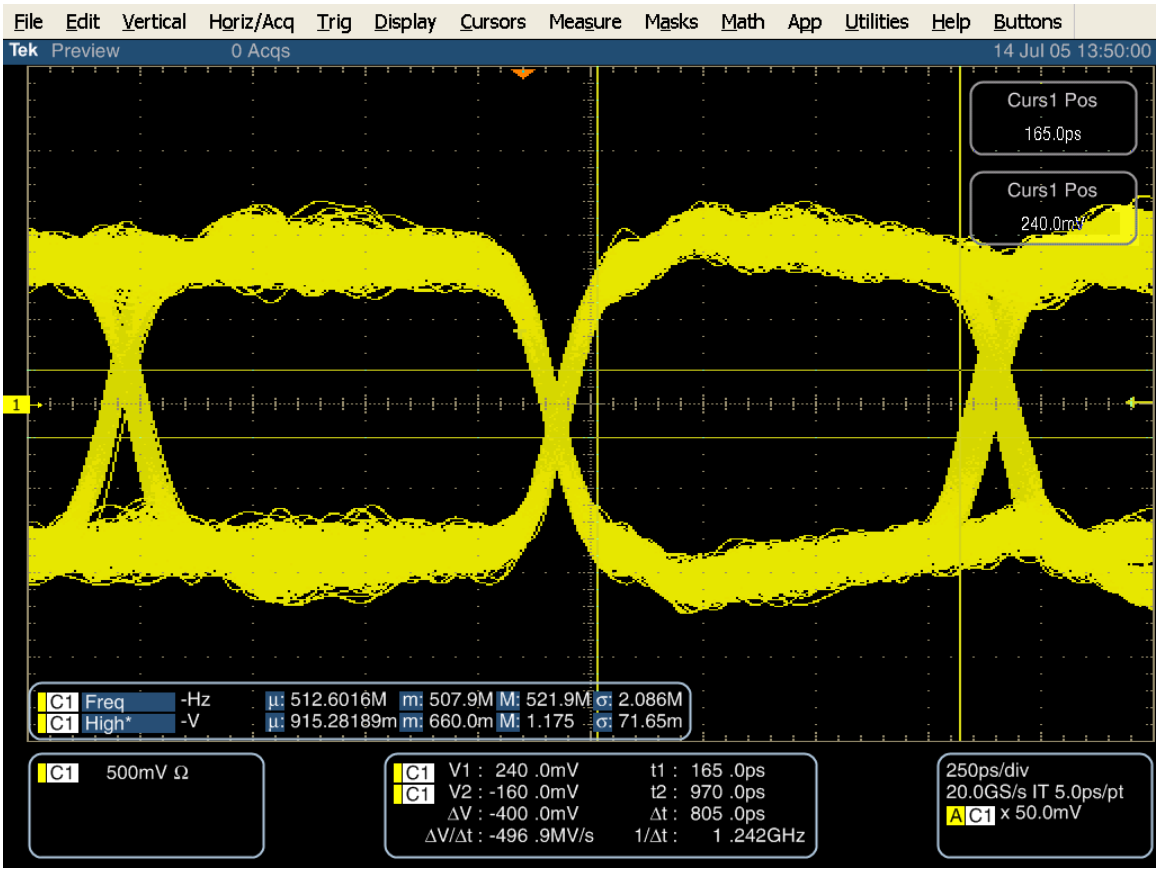

Figure 5 Transmitter—1 m cable—100 cable termination. This eye diagram demonstrates that the transmitter, through the connector, the backplane, and through 1 m of cable operates very well. The eye opening is about 800 psec. For this test, 100 ohm radial-lead termination resisters were plugged into the end of the cable. The signal was measured at the termination resistors with 1 wafer of signals active.

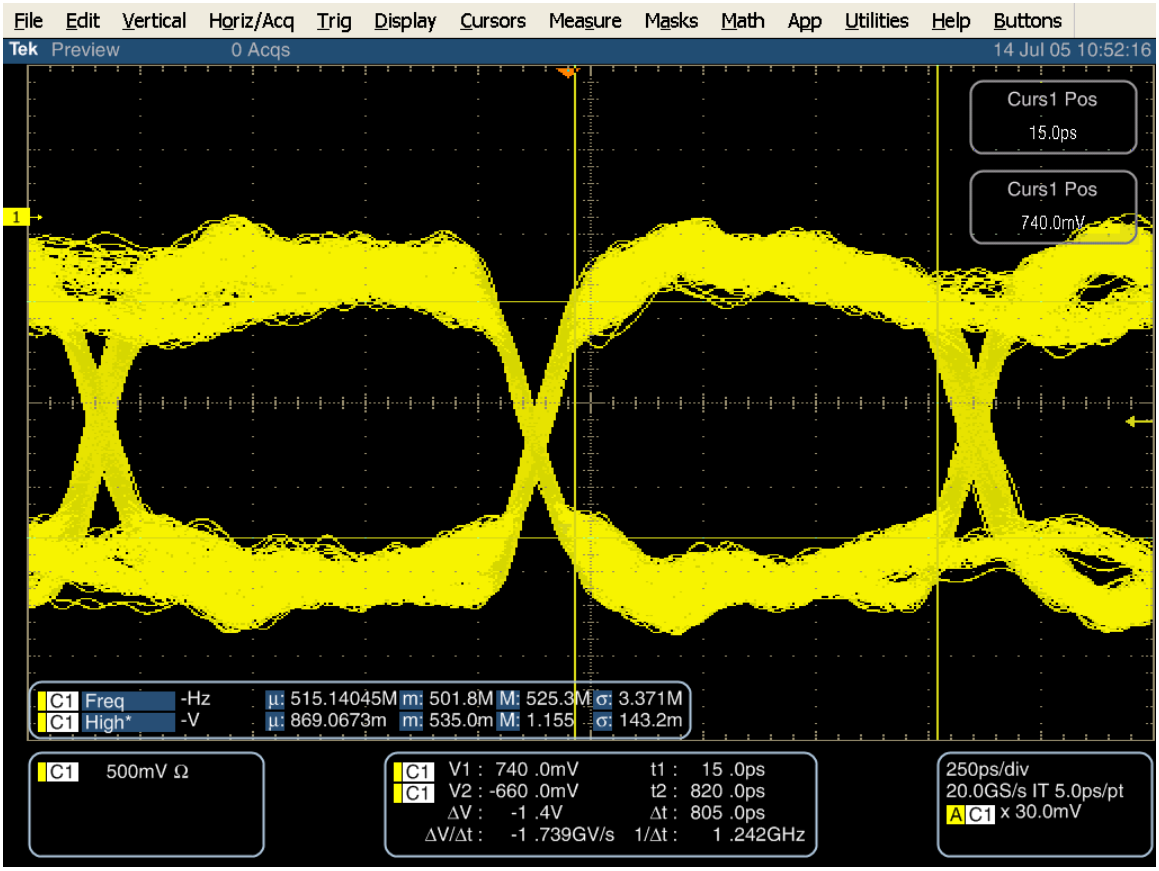

Figure 6 Transmitter—1m cable—backplane—connector—100 ohm termination. For this test, a backplane and right-angle hard-metric connector at the far end were used. 100 ohm termination resistors were soldered on the connector tails for 2 active adjacent wafers. The connector was not press-fit into a circuit board. This test shows some, but not significant, signal degradation over the previous figure. It indicates that there is some degradation going through the connector and backplane at the far end, but it is not significant.

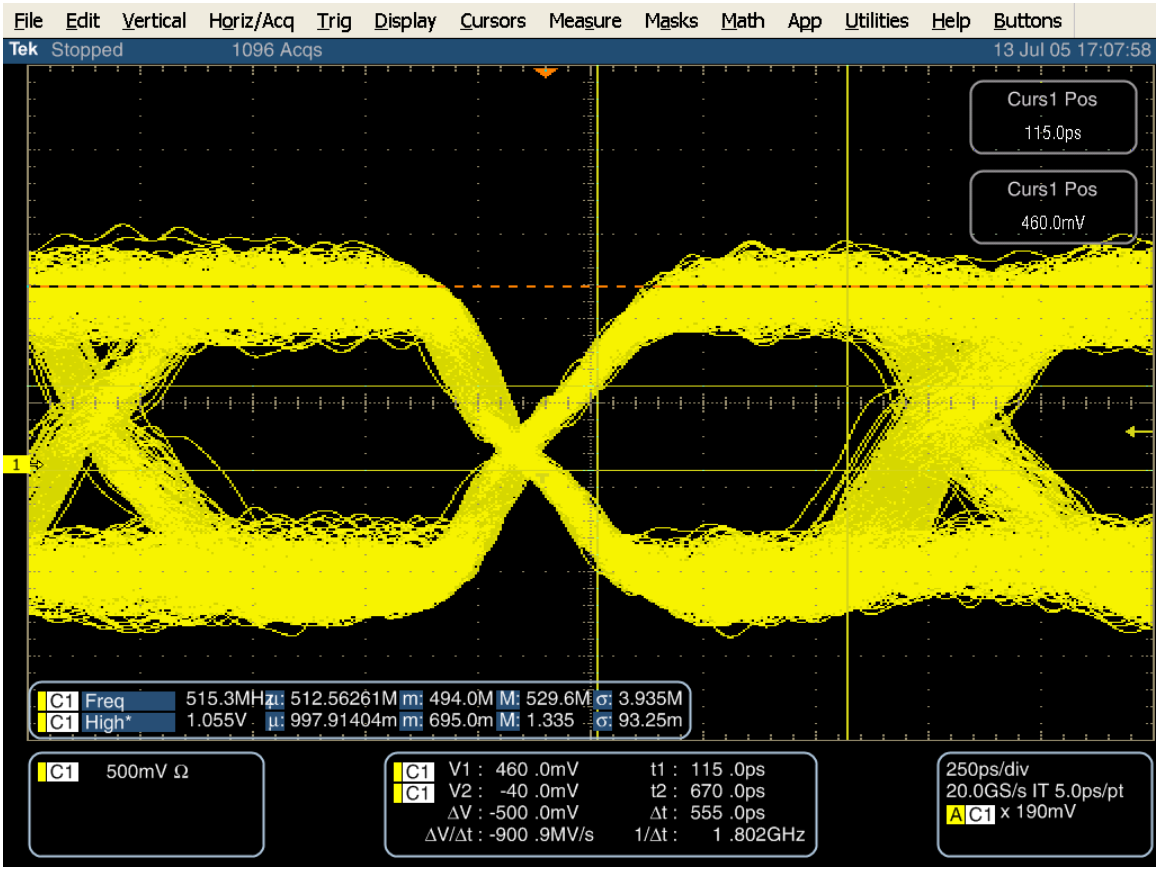

Figure 7 Transmitter—1 m cable—Fanout Board @ buffer input. This eye diagram shows the signal after it has traveled over the 1 m of cable and ~4 inches of FR-4 PCB trace. The increase in jitter, and high-frequency attenuation of the signal compared to the previous two figures is clear. Here, the eye opening is just barely large enough for the Altera DPA receiver. This eye diagram clearly shows the degradation of the signal as it travels over FR-4 at the far end; this is higher than the degradation at the near end, and is believed to be a result of the higher source impedance that the traces see coming from the cable. This result indicates that the length of FR-4 that the signal travels over at the far end (receiver) should be absolutely minimized.

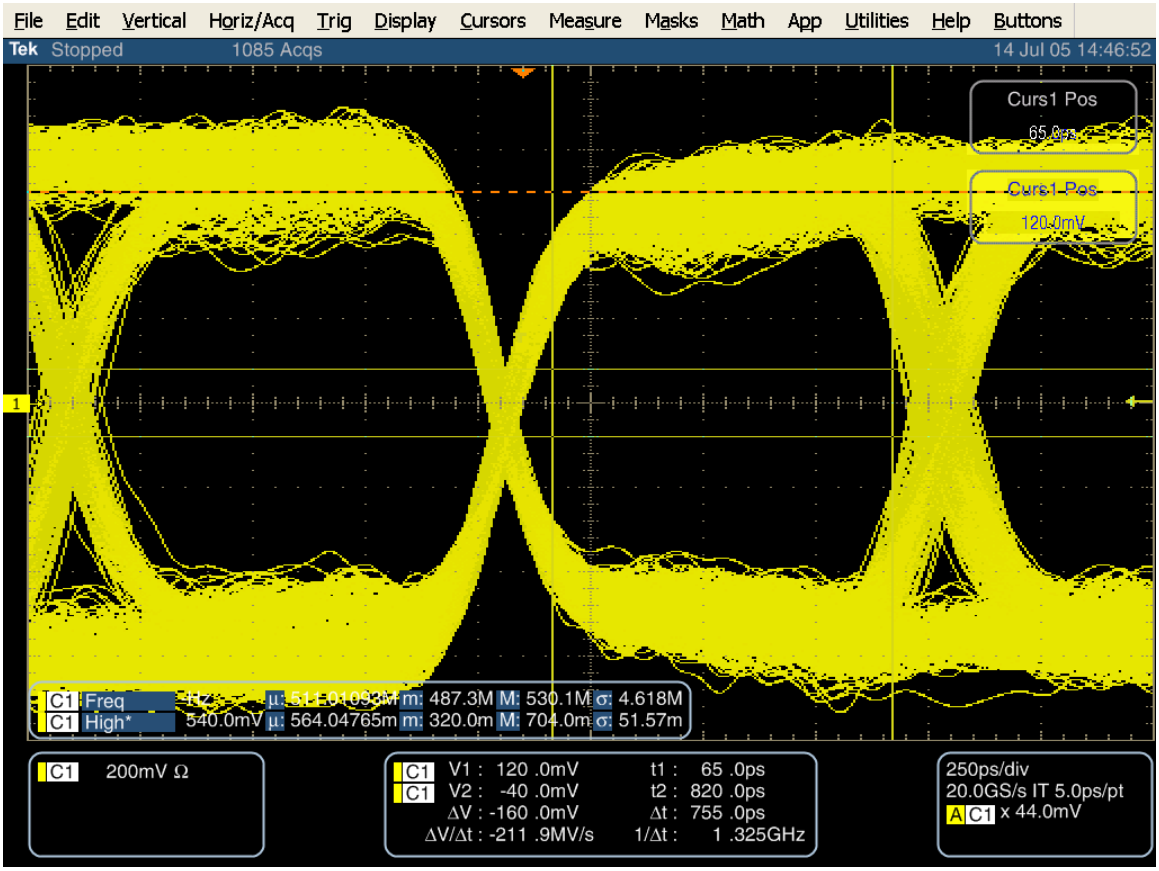

Figure 8 Transmitter—1m cable—Fanout Board—1 m cable—100 ohm term. The buffer on the Fanout Board has cleaned up the signal quite a bit, however the jitter accumulation is clear compared to Figure 5. If the receiver FPGA were seeing this signal, it would likely have no problem with synchronization since the eye is 755 psec. Note that this diagram is not necessarily on exactly the same signal path as the eye of Figure 7, and so the improvement may not be actually real (i.e. a buffer cannot remove jitter, which it appears to have done).

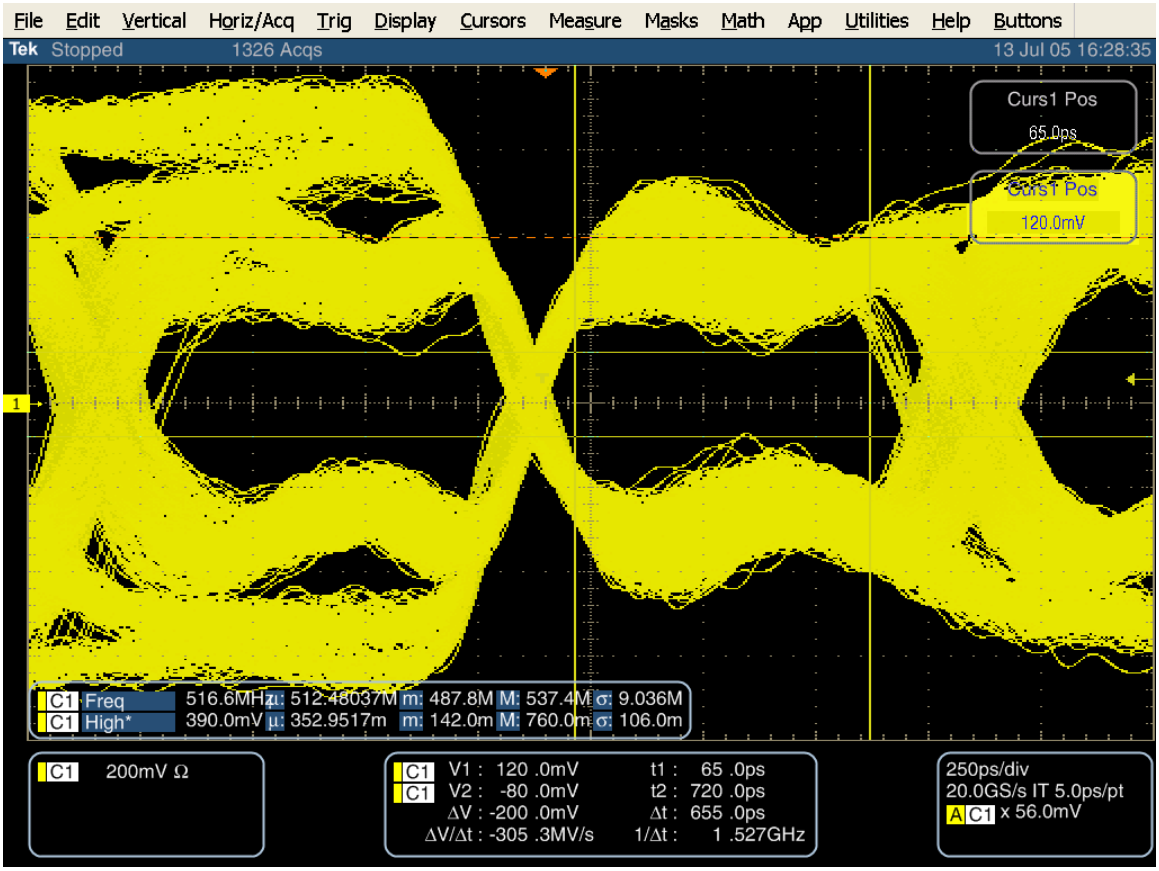

Figure 9 Transmitter—1 m cable—Fanout Board—1 m cable—measured @R185—Receiver on-chip termination. This is the signal of Figure 8 after it has entered the receiver board, measured about half-way down a 4" transmission line. The degradation of the signal from transmission line reflections is clear, and was finally attributed to poor impedance matching of the PCB transmission line. The receiver should be able to handle this kind of signal, but in the test the receiver could not lock on all input signals because some were worse than this (especially the 17" of serpentine transmission line). Nevertheless, this is an unacceptable eye diagram for a reliable system.

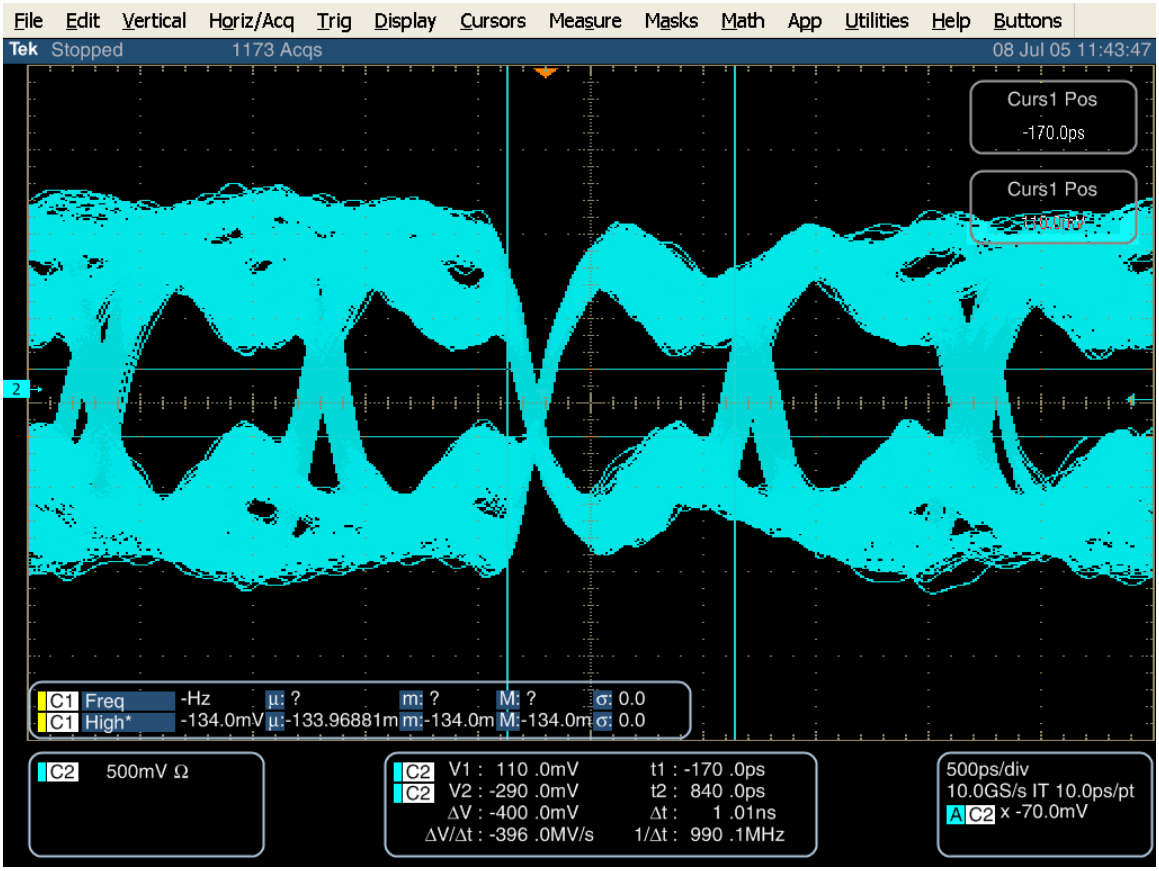

Figure 10 Transmitter—1m cable—measured @ R185—Receiver. The Altera receiver chip is able to lock onto this signal. Reflections on the receiver PCB are clearly evident.

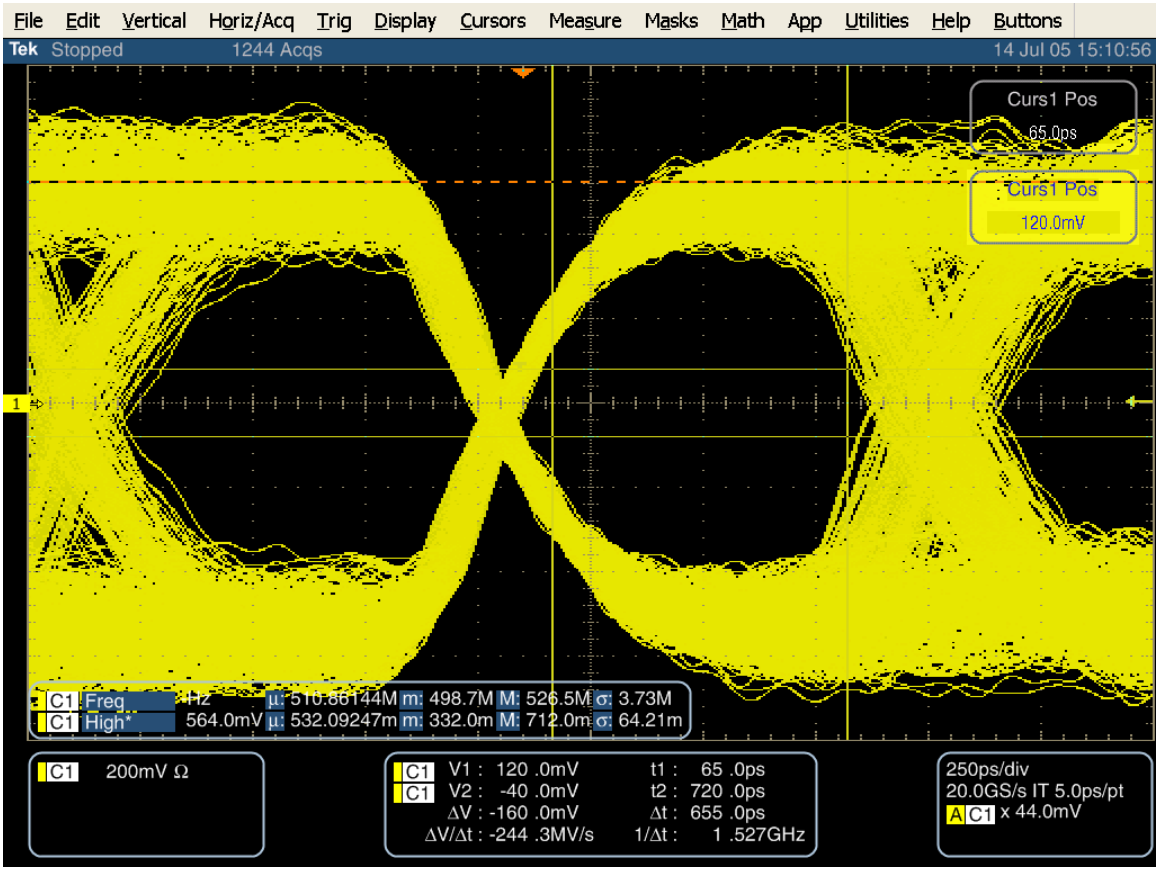

Figure 11 Transmitter—10 m cable—100 ohm termination. Same as Figure 5, only the signal has traversed 10 m of Meritec cable, with equalization built into the cable header. The eye opening is 655 psec, very close to the 560 psec required by the FPGA receiver. The nominal signal amplitude is about 800 mV, and when compared to Figure 5 with a nominal signal amplitude of about 2 V, indicates that the signal has been attenuated by about 8 dB. This is consistent with the attenuation specification of ~9dB @ 500 MHz for the Meritec cable (no equalization), bearing in mind that the fundamental frequency of a 1 Gbps signal is 500 MHz. With the minimum transmitter voltage of 525 mV, the receiver voltage would be about 210 mV.

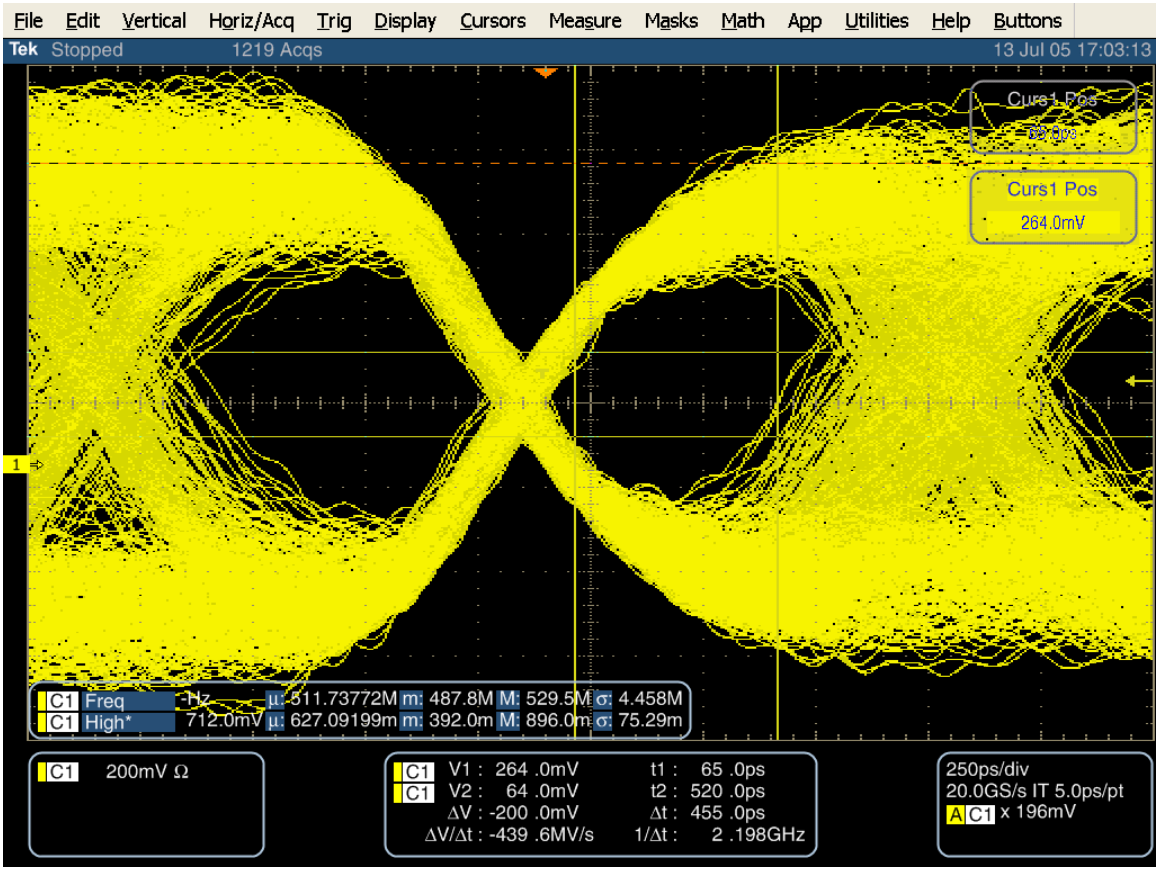

Figure 12 Transmitter—10 m cable—Fanout Board @ input. This is the signal after about 4" of PCB transmission line on the Fanout Board after traveling 10 m. The high-frequency attenuating effect of the PCB trace is clear (i.e. slower rise time). The decreased eye compared to the Figure 11 is evident, and further emphasizes the need to minimize the signal path on the PCB at the receiver end.

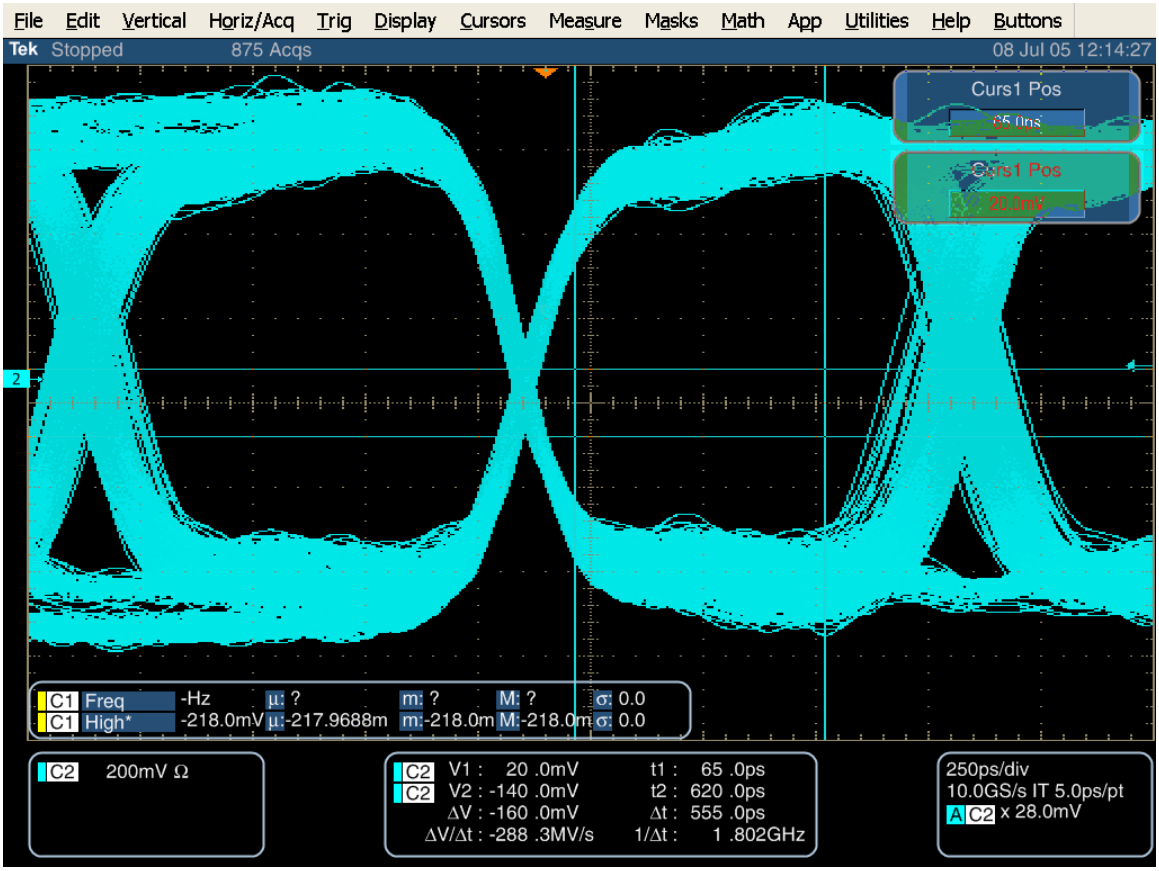

Figure 13 Transmitter—10 m cable—Fanout Board—1 m cable—100 ohm termination. This is the signal after the 10 m of cable and through 1 m of cable after the buffers on the Fanout Board. The increased jitter from 10 m of cable compared to 1 m of cable as shown in Figure 8 is clear, although the eye is much better than Figure 12. If this signal is present at the receivers, they would likely be able to lock onto it reliably.

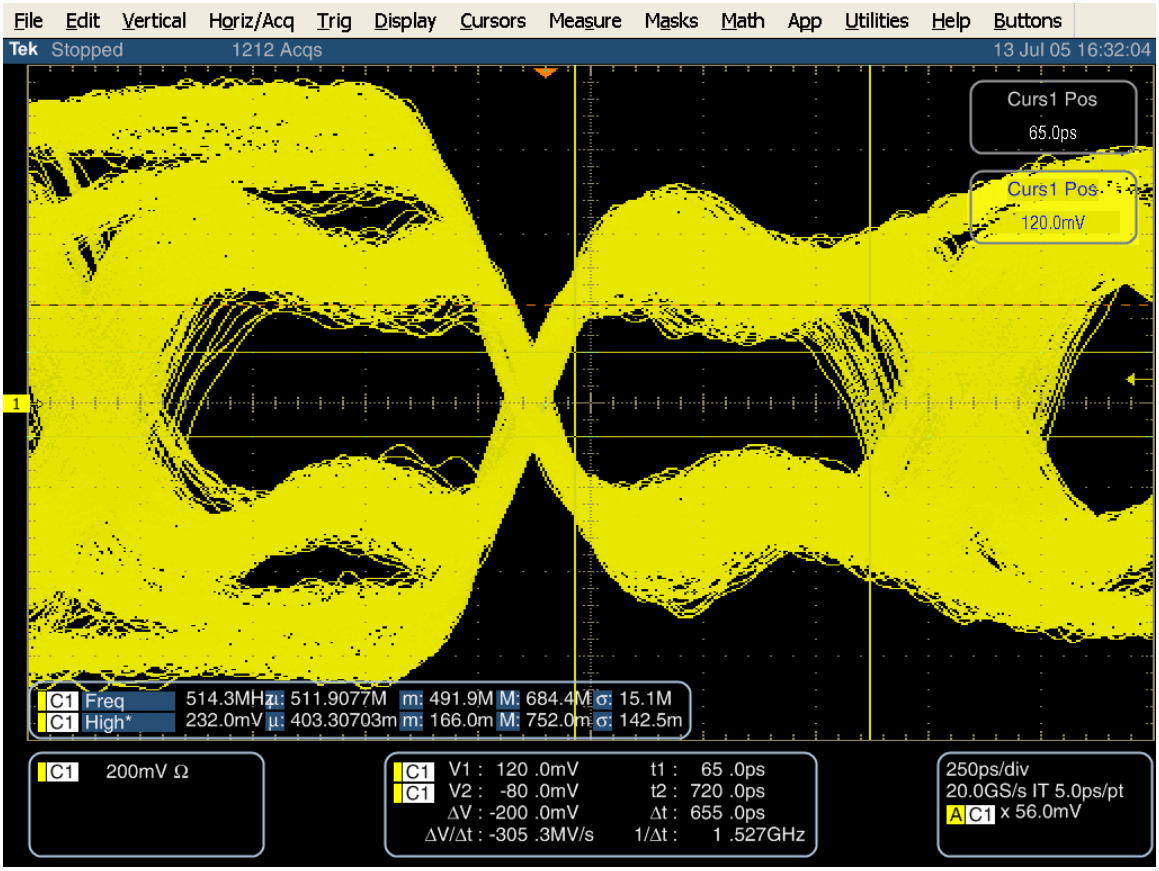

Figure 14 Transmitter—10 m cable—Fanout Board—1 m cable—measured @ R185—Receiver on-chip termination. This eye is quite a bit worse than the eye of Figure 9, and there is likely no way for the receiver to lock onto this signal. The primary degradation factor here is the 4" or so of PCB transmission line on the receiver board.

3.3 Additional Design and Testing Notes

3.3.1 Fanout Board Unused Output Terminations

Testing confirmed that unused outputs of the LVPECL buffer chips on the Fanout Board need to be terminated otherwise other channels are affected. Six 4-wafer outputs should be un-terminated on the board and rely on plugging wafers with built-in resistors into the backplane for termination if they are unused. The additional two 4-wafer outputs should be terminated in such a way that the terminations can be removed or installed after the board is built (i.e. not underneath the connector shield as they are now). This can be accomplished with line-to-line 100 ohm terminations at the output of the chip, followed immediately by DC blocking capacitors that are only installed if the output is used—requiring the removal of the 100 ohm resistors. The eye diagram of one terminated output but with all other outputs un-terminated is shown in Figure 15.

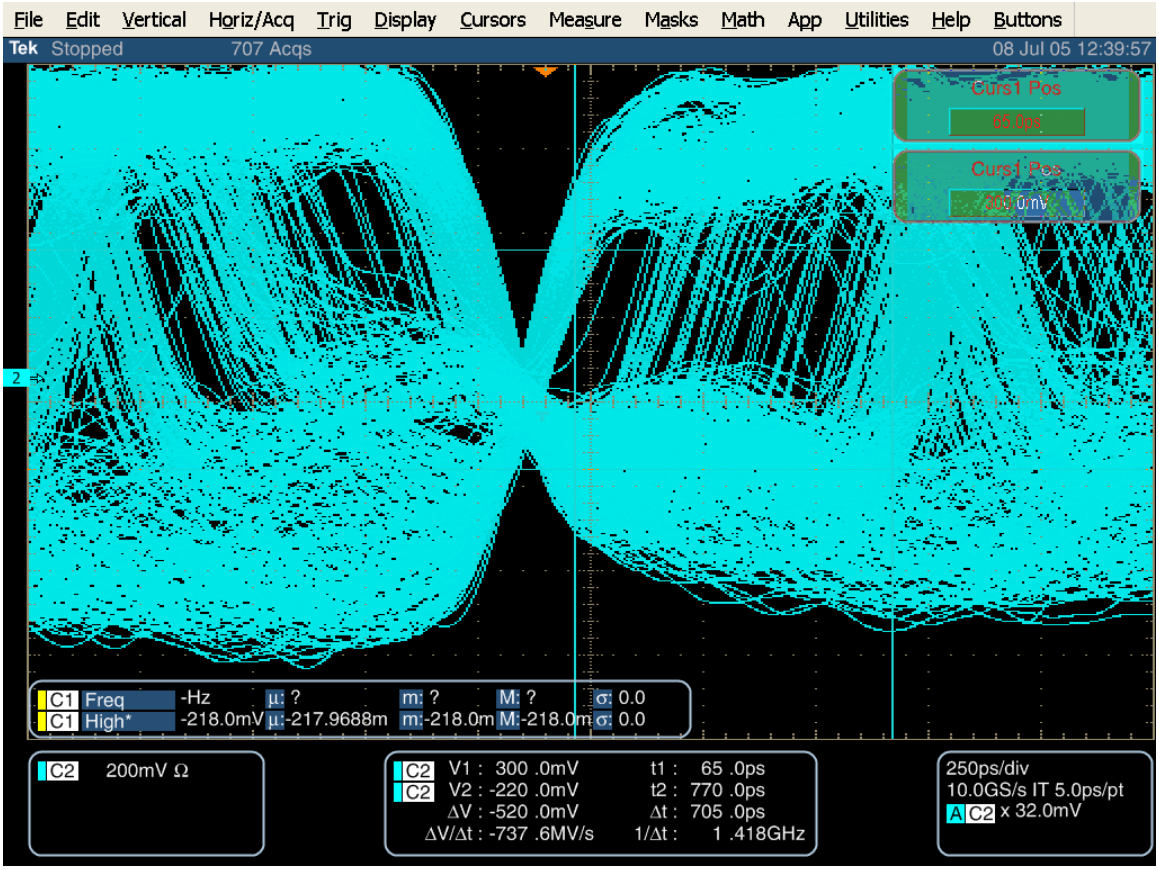

Figure 15 Eye diagram of output of Fanout Board with all unused outputs un-terminated. The eye is completely destroyed in this case.

3.3.2 Fanout Board Output Cross-talk

During testing with a single-ended probe, it became evident that there is cross-talk between output channels on the Fanout Board. This cross-talk was attributed to the close spacing of the PCB traces running from the 1:10 buffers to the connectors. Further investigation using a differential probe indicated that the cross-talk was mostly common-mode and therefore should not affect the receivers. Nevertheless, cross-talk did occur and it is likely prudent to minimize this effect by increasing pair-to-pair spacing. Xilinx recommends that the pair-to-pair spacing[2] is 5X the trace-to-trace spacing. For the Fanout Board that was tested, the pair-to-pair spacing is only 2.2.

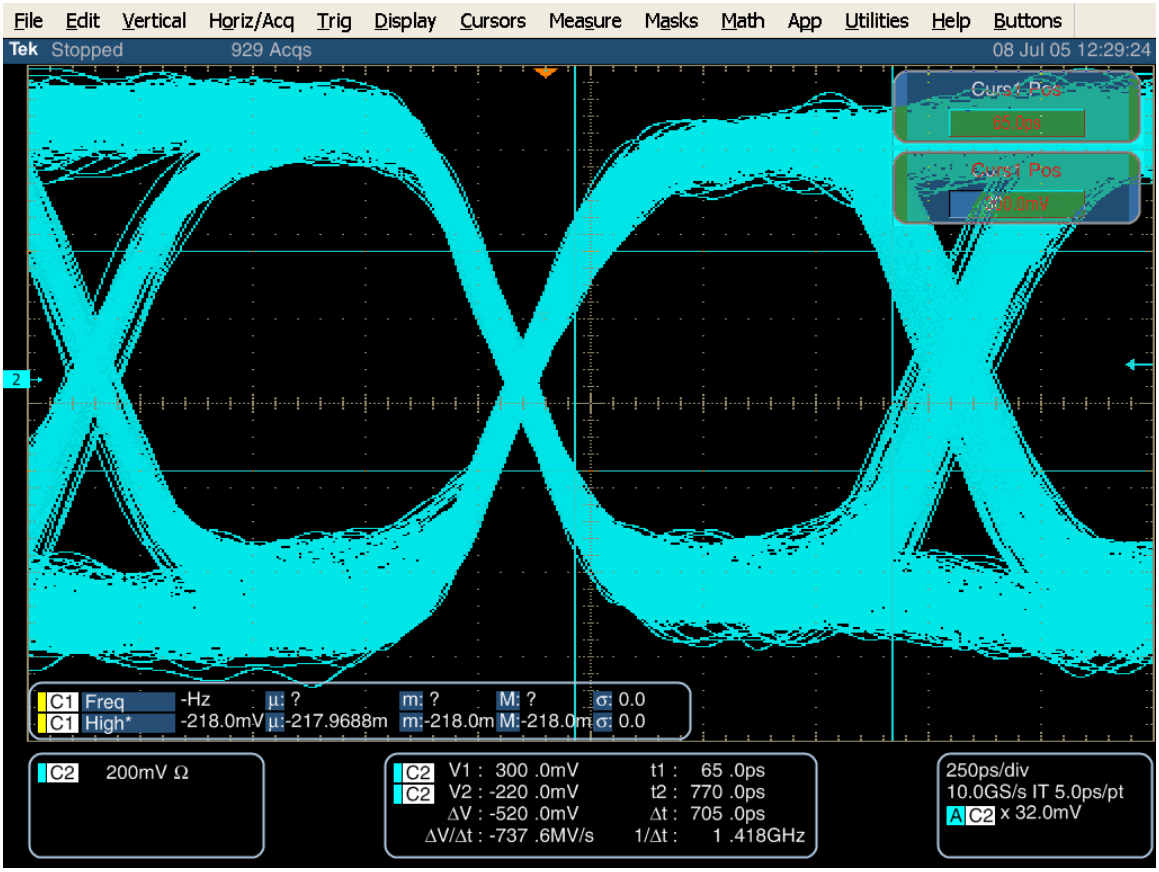

Figure 16 Transmitter—1 m cable—Fanout Board—1 m cable—100 ohm termination; all 4 wafers into the Fanout Board are active. The effect of differential cross-talk appears to minimal. Nevertheless, for robust operation and minimal cross-talk the pair-to-pair spacing should be increased from the current 2X to 5X.

[2] The pair-to-pair spacing is the dimension between the closest traces of different pairs.

4 References

[1] Heng, Zhang, INTERFACE CONTROL DOCUMENT: HM Gbps Cable Physical Specification, ICD Document: A25022N0040, Revision Draft, April 13, 2004.

[2] Carlson, Brent, PROTOCOL SPECIFICATION: HM Gbps Cable Signaling Specification, RFS Document: A25022N0041, Revision 1.1, February 24, 2005.

[3] Carlson, Brent, REQUIREMENTS AND FUNCTIONAL SPECIFICATION: Recirculation Controller FPGA, RFS Document: A25090N0000, Revision 2.2, September 10, 2004.

[4] Carlson, B., Application Note: Timecode Gbit test FPGAs for EVLA Correlator Testing, Version: DRAFT, May 5, 2005.