



Jamieson T. Olsen  
**Engineering Note**

**Date:** 11 January 2001  
**Rev Date:** 19 January 2001

**Project:** Digital Front End  
**Doc. No:** 2001-01-11a

**Subject:** LVDS link and DFE motherboard timing skew

## Introduction

Over the past several weeks there has been much discussion concerning the worst case skew between the LVDS links going to a particular board in the L1 trigger system. The purpose of this document is to explain the terms used in these discussions and also to quantify the skew numbers for various stages in the trigger system.

- ✓ Each LVDS link provides its own clock. This guarantees that if the link clock is used to capture the link data, setup and hold times will *always* be met. Of these constraints setup time is the most critical. Modern FPGAs require at least 2ns of setup time; minimum hold times are less important and are often quoted sub nanosecond.
- ✓ The term “clock domain” refers to all of the synchronous logic in a firmware design that is driven by a particular clock signal. Typically in the FPGA designs there is a clock domain for each incoming link, plus a larger “master” clock domain that drives the output data streams.
- ✓ In each FPGA there comes a point when information must cross from multiple clock domains to a “master” clock domain so that the data can be processed and sent out. This process is referred to as de-skewing the phase differences among the incoming links.
- ✓ Various schemes are used to de-skew the links. Some are more tolerant of skew than others. Often the skew tolerance is a directly proportional to the resources used and the latency through the de-skew logic.
- ✓ DFE motherboards are data driven. That is, they are clocked solely from the incoming data streams. There is no “experiment clock” or “global clock” common to all DFE motherboards. Information such as sync gaps, L1 accepts, resets, etc. must be encoded into the incoming data streams.

# Sources of skew

## AFE Boards

The following Sections were contributed by John Anderson

### LVDS Link Hardware

All data cables from AFE to DFE are commercially manufactured cables with a designed length of 6 meters. As delivered the cables vary no more than +/- 1 inch in length and vary in characteristic impedance by no more than +/- a few percent. The skew introduced by cabling will be well less than +/- 1nsec and is completely **ignorable**.

LVDS cables between DFE motherboards range from 2-10 meters. Skew generated on these cables is also **negligible**.

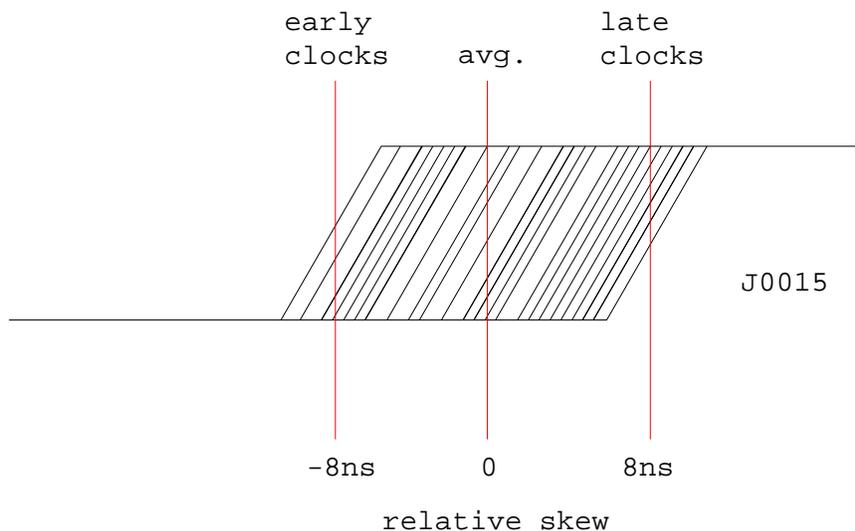
### AFE Skew Due to Detector Timing Constraints

The master D0 Clock generates a Beam Crossing marker over the Serial Command Link (SCL) which is then distributed by the SVX Sequencers to the AFE boards via flat ribbon cables. The Sequencers contain circuitry which allows for manual tuning of the timing of this 'XING' clock. The timing of the XING clock must be adjusted during installation of the detector such that the rising edge of XING occurs in each AFE at a known setup/hold time relative to the charge signal from the VLPCs. If this setup/hold time is not held to tight tolerances the charge will not fall completely within the charge integration window of the SIFT discriminators and the trigger is rendered useless.

The variant fiber lengths within the CFT detector result in the charge being delivered to the inputs of the AFE boards at varying times relative to the actual beam crossing. Thus, the XING clocks delivered to each AFE must also vary in the same way. This torturous manual tuning process is done once during detector installation and should never be changed.

**The skew between XING clocks introduced by this physical constraint is estimated to be +/- 8nsec.**

In other words, if the rising edges of all the AFE board clocks could be superimposed on top of one another, the picture would look like this:



## Skew Due to 53MHz Clock Generation on AFE

As no other continuously running clock other than XING is presented to the AFE board by the system, the only way to synthesize higher frequencies is to use a phase locked loop. In order to serialize the data and transfer it to the DFE boards each AFE multiplies the XING clock by seven to generate a 53MHz data clock. This data clock is used by all links coming out of any given AFE board. A single PC board trace is used for the link clock from PLL to all four LVDS drivers and the **skew between links due to the input link clock is well under 1nsec, probably within +/- 100 psec.**

**Thus, the AFE preserves whatever skew it is given by the Sequencers in the LVDS clocks, but adds no more of its own accord.**

## AFE Skew Summary and Thoughts on Additional Fine Tuning

The AFE clock generation circuitry provides an adjustment mechanism whereby the 53MHz LVDS link clock may be phase adjusted relative to the XING clock on an AFE by AFE basis. The PLL provides a couple of programming bits allowing an adjustment of +/- 4 "time units" in 1 "time unit" steps, where a "time unit" is proportionate to the speed at which the PLL runs. For the LVDS link clock, a "time unit" works out to be 1.18 nsec, providing an adjustment range of +/- 4.72 nsec in each AFE.

Given the complex way in which AFEs route to the DFEs this adjustment should provide compression of the skew in LVDS data. A conservative estimate would be a reduction from +/- 8nsec to +/- 6nsec, sufficient to guarantee a window of valid data across all links at least 4ns to 5ns wide every 18.9 nsec. Use of a fast gate combining all of the copies of the XING clock in each link to create the across-all-links sync pulse, a T-flop and a PLL would then be sufficient to create a master reception clock to remove this skew using a two-stage pipeline. This circuitry should be trivial to implement in today's modern FPGAs that incorporate PLLs directly on the chip.

Any further adjustment to reduce skew would require manual intervention to either the AFE boards or the cables after installation in the detector. Either method is fraught with risk. The cables connecting AFEs to DFEs are fragile and expensive, and the tooling necessary to modify them does not exist on site.

It is possible to introduce changes into the hardware of a given AFE such that the SIFT discriminator clocks are all moved relative to the XING clock. This would allow a compensatory change in the XING clock from the Sequencer, allowing a given AFE's LVDS links to be shifted in 5nsec steps. This could be combined with the software approach noted above to potentially compress the skew present at a DFE board by another few nsec. However, the costs of this approach almost certainly outweigh the benefits.

- 1) Design labor would have to be provided to redesign the programmable logic of the selected AFE.
- 2) Testing the new design cannot be done by simply downloading a new file from the comfort of one's terminal. AFEs can only be reprogrammed by a physical access to the detector.
- 3) Such changes create a whole new tracking scenario where all AFEs are no longer the same, requiring database entry and validation.
- 4) New routing in a programmable logic device always risks creating different delay paths which can result in spikes or glitches that would seriously compromise the noise performance of the SIFT, and thus deteriorate or destroy the trigger efficiency.

# DFE Motherboards

## Board Clock Distribution

Each DFE motherboard selects one of the incoming LVDS link clocks to be the **board clock**. LVDS link #2 or link #7 can be selected as the board clock; this selection can be made remotely.

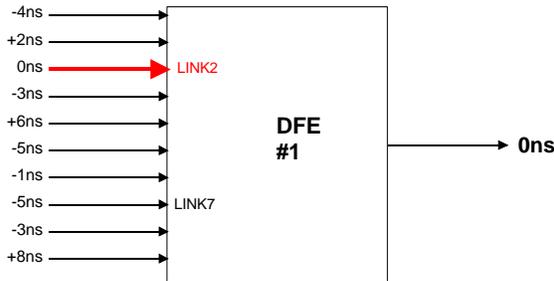
The board clock is re-distributed to all FPGAs on the DFE motherboard using a PLL-based clock driver with very low (<100ps) skew. Data flowing out of the DFE motherboard is synchronous to this board clock.

## Board to Board Skew and Fine Tuning

Since the DFE motherboard board clock is generated from a link clock, skew between DFE motherboards is a function of the skew between the two LVDS links that are selected to be the board clocks.

Due to the very low skew PLL-based clock distribution, DFE motherboards will impose no additional skew into the system.

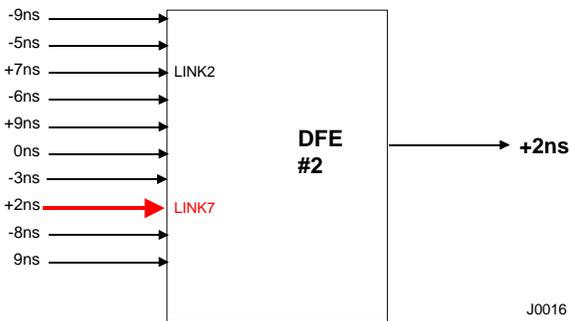
In the example below DFE #1 is configured to use LINK2 for the board clock. DFE #2 uses LINK7 as the board clock. The skew at the outputs of the DFE motherboards is 0 and +2ns; thus the skew between the motherboards is 2ns.



If DFE #1 had been configured to use LINK7 as the board clock the skew between these motherboards would be worse: 7ns.

If DFE #1 uses LINK7 and DFE #2 uses LINK2 the skew between motherboards would be 13ns.

Link selection is not automatic – it requires that someone measures the skew on LINK2 and LINK7 for each motherboard and informs the online systems personnel who then configure the DFE motherboards remotely.



Link selection may not always be able to de-skew motherboards. Consider the situation where:

DFE #1 LINK2 = +8ns    DFE #2 LINK2 = -8ns

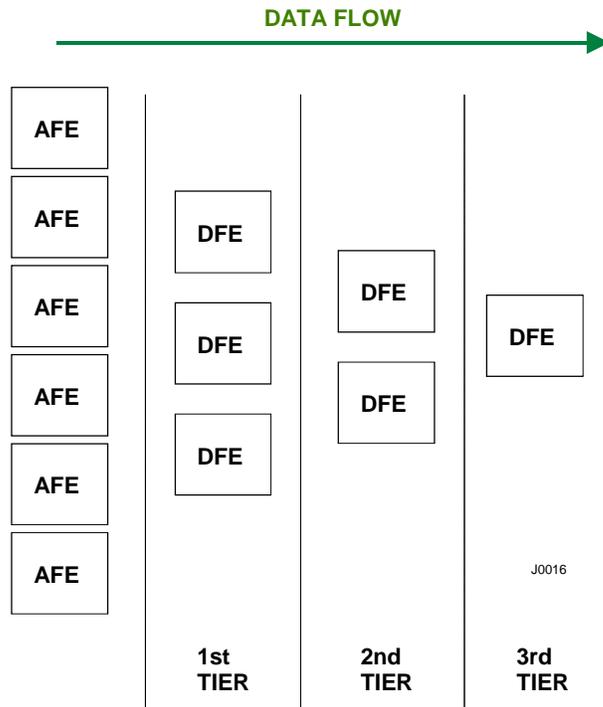
DFE #1 LINK7 = +8ns    DFE #2 LINK7 = -8ns

No matter which links are selected for the board clock, the skew between motherboards will be 16ns.

J0016

# Skew Summary

Keep in mind that these numbers are the **worst case link skew estimates**. A conservative estimate of  $\pm 1$ ns for the skew generated by the LVDS links (Transmitter, Cable, Receiver) is used.



location	worst case link to link skew (ns)
AFE outputs	$\pm 8$
DFE inputs (first tier motherboards)	$\pm 9$
DFE inputs (2nd tier motherboards)	$\pm 10$
DFE inputs (3rd tier motherboards)	$\pm 11$