

Design Note:

HFDN-27.0

Rev 0; 09/03

**Equalize 10Gbase-CX4 and Copper InfiniBand Links
with the MAX3983**

MAXIM High-Frequency/Fiber Communications Group



Maxim Integrated Products



Equalize 10Gbase-CX4 and Copper InfiniBand Links with the MAX3983

1 Introduction

This discussion summarizes popular transmission line compensation techniques and presents a solution to compensate cable losses in "four-lane" InfiniBand (2.5Gbps) and 10Gbase-CX4 (3.125Gbps) transmission links. This information can be applied to Fibre Channel and Ethernet XAUI links as well.

2 Transmission Requirements

The physical characteristics of InfiniBand (IB) and 10Gbase-CX4 (CX4) links are similar. In fact, CX4 began its efforts with a copy of the IB transmission channel definition for cable and connectors. CX4 has since modified the requirements by adding more complete frequency-domain requirements for return loss, through loss and crosstalk. IB and CX4 both are targeting 15 meters using 24AWG (American Wire Gauge) cable and both use the same connectors scheme.

3 Loss Compensation and Signal Regeneration

Compensating for transmission line losses can take many forms with many names. The terminology has become more confusing lately as vendors attempt to differentiate product by coining new phrases for the same old thing. Compensation that is provided after encountering losses, by locating a circuit at the termination (receiver) end of the line, is known simply as "post-compensation." Compensation that is provided before encountering losses, by locating a circuit at the source (transmitter) end of the line, is "pre-compensation." In some situations, both pre- and post-compensation are used in the same link to correct for losses. Other popular terms synonymous with compensation are equalization and emphasis. Compensation can be achieved many ways. The simplest and oldest form is a high-pass analog filter. Such a filter is designed to compensate for the dispersive or frequency-dependent losses of the

transmission medium. By so doing, the original shape of the signal can be recovered, albeit smaller in amplitude. To regenerate the amplitude of the signal at the receiver, one or more gain elements must be introduced.

3.1 Post-Compensation

Compensation administered at the receiver can be fixed, manually adjusted or automatic.

3.1.1 Automatic Compensation

Automatic techniques detect some characteristics of the signal such as the power relationship between two, or more, spectral bands within the pass band of the system. By comparing the differences to what is expected for the bit rate and modulation scheme, signal restoration is accomplished by adjusting tunable elements or gain blocks in a regulated closed-loop system. The effectiveness of this automatic technique is fantastic when the signal characteristics are predictable, but the channel characteristics are sometimes unknown or known to change often. See the MAX3800, MAX3801 and MAX3980 for examples. Compensation range is often on the order of 30dB. Such systems require that certain prerequisites be met. The automatic sensing of the needed correction requires some time to acquire and settle. During this time, the data must be continuous, usually for a few microseconds or more. Once the correction is established and operating with closed-loop regulation, the data cannot be interrupted for more than a few hundred nanoseconds to maintain regulation. Longer and indeterminate intervals can be accommodated, but with increased complexity.

3.1.2 Manually Adjusted Compensation

Manually adjusted techniques are far simpler and need not depend upon closed-loop regulation to control the amount of compensation. As with the automatic version, the compensation range is on the

order of 30dB. The use of manually adjusted devices requires some more knowledge of the variations encountered in the transmission environment, but they remain very versatile for wide-band and intermittent data applications. See, for example, the MAX3804.

3.1.3 Fixed Compensation

Fixed techniques do not allow adjustment of the compensation. Fixed is the simplest of the three types; however, it is not as effective over the same 30dB compensation range. It is not that 30dB is not possible; it is not practical for general-purpose use. High compensation levels require very low noise environments not usually found in multi-channel backplane and cable links. A practical compromise is to offer less compensation to gain a wider range of use. Contrary to popular belief, a fixed equalizer is not relegated to a fixed channel length. Fixed compensation has been shown to be effective over very wide ranges in bit rate, distance and coding schemes while offering as much as 15dB of compensation in incredibly small packaging. See, for example, the MAX3785.

3.2 Pre-Compensation

Compensation administered at the transmitter can also be fixed, manually adjusted or automatic; however, automatic detection of the needed compensation often requires knowledge of the fidelity (error rate) at the receiver to adjust the


compensation at the transmitter. The most common technique for pre-compensation is to use fixed compensation for the best compromise between distance, receiver sensitivity and signal-to-noise ratio. The most common form of this is today is called “pre-emphasis.” Because the compensation is a high-pass filter, it suggests that there is a great deal of overshoot in the step response. This overshoot sounds the alarm to those worried about radiated emissions and crosstalk. The assumption is that the added compensation is the result of increasing the magnitude of the high-frequency content. A term intended to convey how the emphasis is generated is “de-emphasis.” This term has come to mean compensation delivered by the transmitter by reducing the low-frequency content rather than increasing the high-frequency content. The effect is the same, a high-pass filter. The difference is the gain. In the truest sense of the word, they are both “pre-compensation”; therefore, they are both “pre-emphasis.” You just have to know if the swing (gain) and compensation are appropriate for your application.

3.2.1 Expressing “Pre-Emphasis”

Table 1 below is a compilation of four different ways to express the quantity of pre-emphasis. With the widespread use of digitally generated pre-emphasis has come new ways of expression. Whereas the effect of an analog, high-pass filter can be seen in decibels (dB), other ratios that relate to specific implementations have become popular. Table 1 is provided as a translation tool.

Table 1. Compilation and comparison of different pre-emphasis expressions

Pre-emphasis Translation			
Ratio	α (alpha)	10Gbase-CX4	in dB
$\frac{V_{HIGH_PP}}{V_{LOW_PP}}$	$\frac{V_{HIGH_PP} - V_{LOW_PP}}{V_{HIGH_PP} + V_{LOW_PP}}$	$1 - \frac{V_{LOW_PP}}{V_{HIGH_PP}}$	$20 \left[\log \left(\frac{V_{HIGH_PP}}{V_{LOW_PP}} \right) \right]$
1.41	0.17	0.29	3dB
2.00	0.33	0.50	6dB
2.82	0.48	0.65	9dB
4.00	0.60	0.75	12dB



3.2.2 Technique versus System Constraints

CX4 has decided to simplify the transmitter and receiver by defining a link that uses a prescribed amount of compensation at the transmitter and receiver. CX4 is intended to transmit Ethernet XAUI data continuously. This being the case, some form of automatic equalization could be used by a CX4 receiver, but the meager compensation requirement of the receiver doesn't warrant the added complexity. Whether the receiver is automatic or not, compensation at both the transmitter and at the receiver is necessary to meet the CX4 distance goal of 15 meters while maintaining interoperability with other CX4 devices.

On the other hand, InfiniBand does not operate with a continuous signal. There are various standby modes, during which no data are transmitted. Additional signaling contained in the envelope of timed bursts called "beaconing" is used to activate a link that is in standby. Unless equipped with some storage mechanism, an automatic equalizer will not be able to maintain closed-loop control for the long periods in which the data are absent. The intermittent nature of the beaconing poses regulation problems because the data are not stable and the bursts are modulated, single tones (i.e. 1010....). This beaconing behavior reduces the compensation

choices to a fixed form in the receiver and either fixed or adjustable in the transmitter.

To best satisfy both of these applications, a modest level of fixed compensation in the receiver and a wide range of adjustable compensation in the transmitter are recommended.

4 The MAX3983 Transceiver

The approach used by the MAX3983 is to rely on both pre- and post-compensation to correct for dispersive losses in the transmission line. The cable transmitter offers four selections for pre-emphasis as controlled by two digital control lines (i.e. two bits). The transmitter compensation ranges from 3dB to 12dB at 1.6GHz.

The cable receiver section of the MAX3983 uses a fixed compensation stage providing approximately 6dB of boost at 1.6GHz. Alone, this is sufficient to compensate 5meters of 24AWG cable. Combined with the maximum compensation provided by the transmitter, spans of 20meters can be achieved.

The host side of the link is supported by less aggressive transmitter pre-emphasis and receiver equalization designed to compensate for up to 0.5 meters of FR4 ($\tan\delta = 0.022$) board material.

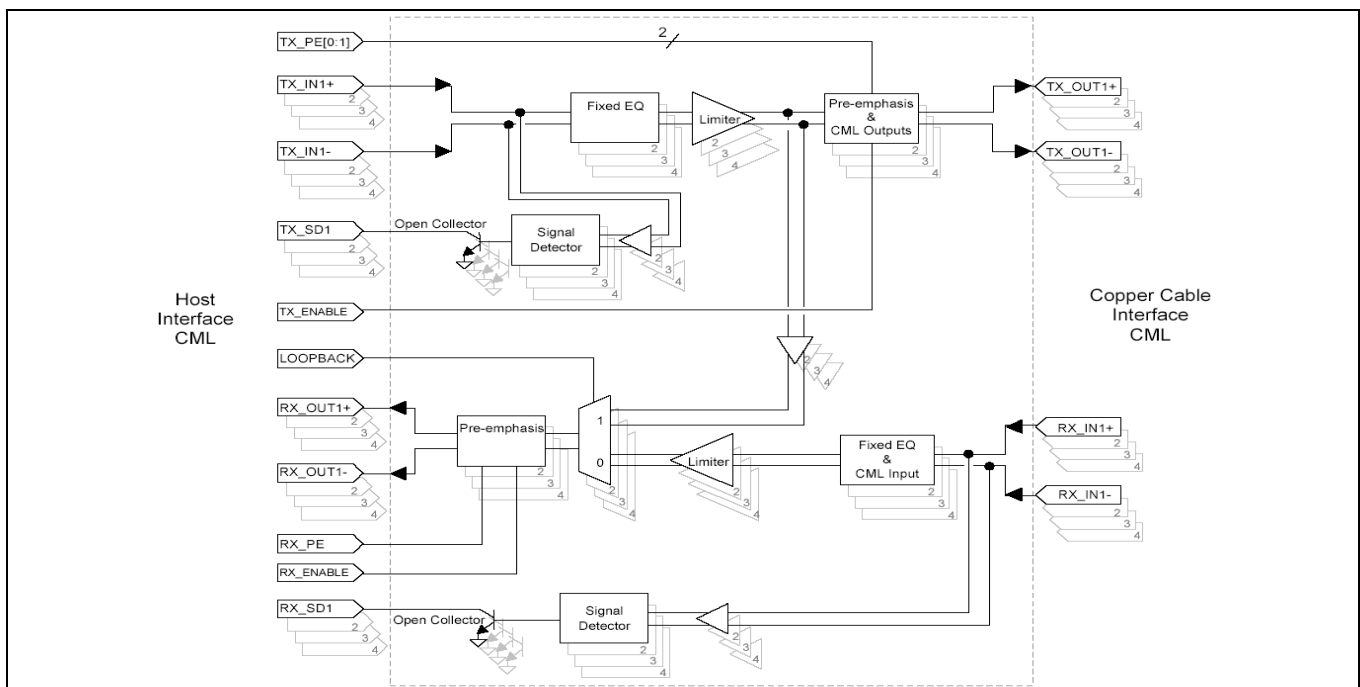


Figure 1. Block diagram of MAX3983, four-channel copper cable signal conditioner.

4.1 Analog Pre-Emphasis

Today many serial transmitters implement a pre-emphasis that is generated using a simple digital technique. By delaying a copy of the signal (usually by one clock cycle), attenuating it and subtracting it from the original signal, some excellent results can be achieved over modest distances. The transmitted signal emphasizes (increases amplitude) the first bit following a one-to-zero or a zero-to-one transition. If this is your first experience with line compensation, it might seem a peculiar and mysterious way to correct for the cable loss. It becomes easy to appreciate this idea if you think of the transient response as an approximation of what a high-pass filter produces. It creates overshoot that helps to keep the eye open at the receiver. It is a very practical approach given the pervasiveness of very large-scale, digital chip integration; however, it remains an approximation of the ideal response. The MAX3983 relies on a different approach. It uses an analog filter that is a much better approximation of the needed response. Using a few well-placed "zeros" in the frequency response, a near perfect compensation can be achieved for a given cable length. With this filter placed in the transmitter, the resulting pre-emphasis is clearly "analog." With the pre-emphasis being generated by an analog filter, there is no dependence on a clock. The MAX3983 can operate from 1Gbps to well over 4Gbps. Its simple analog filter can create long time constants for low frequency compensation that are considerably complex and power hungry in digital implementations.

Figure 2 shows the frequency response of a 15meter, 24AWG cable with approximately 16dB of loss at 1.6GHz. This includes the through loss of two mated pairs of connectors. A filter (Rs, Ls and Cs) was designed to compensate for the loss. The curve in Figure 2 labeled "Combined" is the total response of the filter and the cable. This response exhibits a gentle roll-off without ripple or peaking. This filter response shown in Figure 2 has been incorporated into the transmitter section of the MAX3983 and has been made adjustable to optimize the compensation for shorter lengths.




Figure 2. Frequency response of 15meter, 24AWG cable, the compensating filter and the combined response of the filter and the cable.

Figure 3 shows an overlay of the four different pre-emphasis levels. The data pattern is 5-ones followed by 5-zeros. Note that the peak-to-peak of the emphasis is constant while the long-term logic levels of ones and zeros are varied to achieve the compensation ratio (dB). The emphasis is nearly 1.6Vpp immediately at the output pins of the MAX3983. After a short, 5cm span of FR4 board material, the emphasis is reduced to approximately 1.3Vpp while the long-term logic levels of a one and zero remain unchanged.




Figure 3. MAX3983 transmitter output. The ratio of emphasis to vertical eye opening is shown in dB corresponding to the four available pre-emphasis settings.

Table 2. Pre-emphasis setting (TX_PE1, TX_PE0) versus cable output. Horizontal full scale = 3UI. Vertical full scale = 1200mV. \updownarrow = vertical eye opening. \leftrightarrow = horizontal eye opening.

	TX_PE1	TX_PE0	1meter 24AWG 3.125Gbps	5meter 24AWG 3.125Gbps	10meter 24AWG 3.125Gbps	15meter 24AWG 3.125Gbps	20meter 24AWG 3.125Gbps
Minimum Pre-emphasis	0	0	 \updownarrow 960mV \leftrightarrow 300ps	 \updownarrow 570mV \leftrightarrow 280ps	 \updownarrow 240mV \leftrightarrow 200ps	 \updownarrow 0V \leftrightarrow 0s	 \updownarrow 0V \leftrightarrow 0ps
		1	 \updownarrow 720mV \leftrightarrow 300ps	 \updownarrow 600mV \leftrightarrow 290ps	 \updownarrow 290mV \leftrightarrow 240ps	 \updownarrow 70mV \leftrightarrow 130ps	 \updownarrow 0V \leftrightarrow 0ps
		0	 \updownarrow 450mV \leftrightarrow 300ps	 \updownarrow 480mV \leftrightarrow 300ps	 \updownarrow 340mV \leftrightarrow 300ps	 \updownarrow 140mV \leftrightarrow 240ps	 \updownarrow 20mV \leftrightarrow 70ps
		1	 \updownarrow 220mV \leftrightarrow 290ps	 \updownarrow 240mV \leftrightarrow 270ps	 \updownarrow 240mV \leftrightarrow 290ps	 \updownarrow 190mV \leftrightarrow 290ps	 \updownarrow 100mV \leftrightarrow 270ps

5 Link Performance

5.1 Transmitter Contribution

Table 2 reports the relationship between the step response and eye at the end of various cable lengths and the two transmit pre-emphasis control bits (TX_PE1 and TX_PE0). This table shows only the contribution by the transmitter. Further compensation and gain needed to restore the signal are provided by the receiver section of the MAX3983. The horizontal eye opening is denoted by the symbol \Leftrightarrow . The vertical eye opening is denoted by the symbol \Uparrow . All of the waveforms are displayed with the same scaling. The horizontal full scale is 3UI (unit or bit intervals). The vertical full scale is 1200mV. For brevity, the data rate was 3.125Gbps and the pattern was a 2^7-1 , PRBS. All waveforms were captured using the Tektronix CSA8000 oscilloscope equipped with the 80E03 sampling head.

Finding the correct pre-emphasis setting for a given cable is a straightforward exercise. Simply look for the best eye opening, but don't use more pre-emphasis than you need. Increasing the pre-emphasis level will extend the reach with very little jitter penalty at short distances; however, the vertical eye opening is small at all lengths! The signal-to-noise ratio at the receiver must be considered.

5.2 End-to-End Performance

When using the maximum pre-emphasis setting with short cables, far-end (relative to the receiver) reflections and far-end crosstalk do not get attenuated significantly and comprise a significant source of interference. Alternately, long cables require more emphasis to compensate for loss resulting in small voltage levels (e.g. vertical eye opening on the order of 100mVpp). This signal level needs to be considered in the presence of near-end coupling of the outgoing transmission adjacent to the receive channel. In such a case, a receiver might have the sensitivity to accurately regenerate the signal providing the crosstalk is less than 20mVpp. In the presence of four offending, 1.6Vpp transmitters, this is nearly 40dB of isolation. Such isolation is not likely to be found in the CX4 and IB connector scheme. CX4 has a specification of

approximately 30dB as the minimum isolation. Using less pre-emphasis to achieve a larger signal to noise ratio will mitigate crosstalk interference.

For custom or non-standard interfaces, it is possible to achieve sufficient isolation to use the maximum capability of the MAX3983 and achieve longer spans.

6 Conclusion

To get the best compensation for a given transmission channel, it is important to understand the advantages and prerequisites for different compensation techniques. Although it is tempting to use fully automatic techniques for everything, certain restrictions apply. In the case of InfiniBand, the beaconing burst will prove to be a challenge and perhaps a barrier to automatic equalization systems. In the case of 10Gbase-CX4, compensation in the form of pre-emphasis is required. This in turn gives some relief to the receiver and allows for less complicated techniques. With the receiver simplified and greater demand for transmitter pre-emphasis, the MAX3983 emerges as the logical choice for compensating cable links such as these. The use of adjustable "analog" pre-emphasis serves the requirements of CX4 and IB while offering superb performance for custom cable assemblies up to 20meters in length and operating between 1Gbps and 4Gbps.