Cross-talk in bi-directional, single wavelength, single fiber Gigabit Ethernet links

by
Vipul Bhatt, Finisar Corporation
vipul.bhatt@finisar.com
408-857-1973
June 28, 2001

Abstract

This paper analyzes the performance considerations for fiber optic links that deploy Gigabit Ethernet (1.25 Gb/sec) over a single fiber, supporting full duplex, bi-directional transmission using one wavelength. Couplers that are not wavelength specific are desirable because they are environmentally robust, permit broad spectrum 1310 nm lasers to be used, and eliminate some logistical problems associated with using WDM in single fiber links. However, reflections of optical signal at the transceiver-cable interface are a major source of performance degradation in single-wavelength links, because they cause cross-talk. This degradation can be overcome by paying a corresponding power penalty, i.e., by increasing transmitter output power to regain an acceptably low error rate. We analyze the extent of this cross-talk penalty, and show how its severity is linked to permissible channel insertion loss. We show that this cross-talk penalty can be limited to a small value by adjusting the receiver threshold. A model link power budget is suggested and some design recommendations are made.

Introduction

The IEEE 802.3 EFM (Ethernet in the First Mile) Study Group is evaluating the feasibility of bringing Gigabit Ethernet (1.25 Gb/sec) to the residential and office communities.[1] Several types of physical layer solutions are being examined, including one that uses a single fiber to carry full duplex communications. While a proposal to use just one fiber is attractive because it reduces the cost of deployment, it also brings with it some new design considerations.

Such a link is shown conceptually in Figure 1. The link length L is expected to be up to 10 kilometers. There needs to be some device – we will call it a coupler – that combines and separates the transmit and receive signals (traveling over the single fiber) within a transceiver module. These couplers can be classified in two categories – those that use Wavelength Division Multiplexing (WDM) and those that don't.

WDM couplers assign one nominal wavelength (or a band) to the transmit path and another to the receive path. For EFM application, there are several difficulties with their use. First, it does not allow the transceivers at the two ends of a link to be identical – both ends cannot transmit at the same wavelength. This limitation increases the logistical and deployment costs. Second, the choice of a wavelength plan becomes restricted, affecting cost and performance. If a 1310/1550 nm combination is used, the dispersion at 1550 nm may limit the maximum permissible spectral width of a laser, affecting the yield of Fabry Perot (FP) lasers or requiring the use of more expensive DFB lasers. Also, there is some concern that such links may have to accommodate an overlay of analog video systems operating at 1550 nm. If two wavelength windows in 1310 nm region is used, it becomes challenging to simultaneously support a wide temperature range and use low cost FP lasers. (The operating wavelength range of a link around 1310 nm is limited by singlemode cutoff wavelength on one side, and the water absorption attenuation peak wavelength on the other side. Fitting two wavelength windows within that range is difficult, especially when you consider that we need to support a wide temperature range, allow manufacturing tolerances for laser wavelengths to keep yields high, and also allow a guard band between two windows.) The upshot is that either DFB lasers or temperature-controlled FP lasers may have to be used, each having a cost implication.

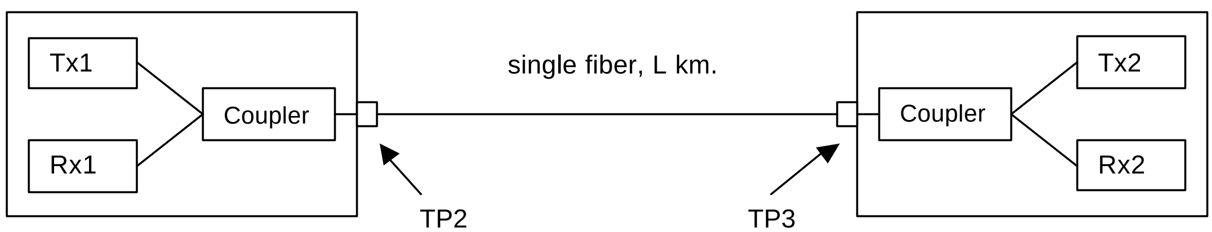

Fig. 1. Block diagram of a single-fiber link

An alternative is to use a coupler that is not wavelength selective. Instead of discriminating along wavelengths, such a coupler combines or splits optical power. By operating the link in the 1310 nm band, and by permitting high yield FP lasers with a fairly broad center wavelength range, the disadvantages of a WDM link described above are eliminated. But prima facie, such links have two disadvantages – their deployment makes the link performance sensitive to reflections, as we will describe shortly. And they have a higher insertion loss. However, as this paper shows, they can be successfully used in EFM links because with suitable design, a link length of 10 kilometers can be supported. We will assume the use of such couplers in this paper because the cross-talk problem is relevant only if such couplers are used.

Now let's review the cross-talk problem. Points marked TP2 and TP3 in Figure 1 are the locations where the transceivers interface with the fiber optic cable plant, which consists of a combination of indoor and outdoor type fiber cables, splices and connectors. Typically, this interface is the mating point between a fiber optic connector and the transceiver receptacle. If a straight-face fiber end encounters an air gap, about 4% of light is reflected. This is characterized by a specification called the return loss, expressed in dB. The return loss specification for 802.3z links is 12 dB.

A part of the signal sent by Tx1 will be reflected back at TP2. Assuming that the coupler ratio is 50:50, half of this reflection will be received by Rx1. If there were no reflections, Rx1 would be receiving signal only from Tx2. Instead, now it is also receiving this unwanted cross-talk, a portion of the energy of the signal sent by Tx1. The question is, how badly does this cross-talk degrade performance, and what can be done to overcome that degradation?

At first glance, the scenario may indeed look bleak, as described by the following example. Assume that the channel insertion loss (a total of cable attenuation, connector loss and splice loss) is 7 dB, the coupler insertion loss is 3.5 dB, the return loss at TP2 (and TP3) is 12 dB, and the transmitter output power is 0 dBm. Then the optical power sent by Tx2 and received by Rx1 will be:

Received signal power = Tx power – 2*coupler loss – channel insertion loss
= 0 – 7 – 7 = -14 dBm

And the received cross-talk power (Tx1 signal power arriving at Rx1) will be:

Received cross-talk power = Tx power – 2*coupler loss – return loss
= 0 – 7 –12 = -19 dBm

This would give a "signal to cross-talk ratio" of 5 dB, an apparently low value that can degrade performance. However, such a conclusion is misguided. What really matters is the ability of the receiver to make correct decisions in the presence of this crosstalk and noise. A closer examination will reveal that the link can continue to meet its target specification of bit error rate of 10^-12 or less.

Bit Error Rate in presence of cross-talk and noise

In this section, we will derive an expression for the bit error rate as a function of various link parameters. The following variables will be used:

return_loss: Return loss at TP2 and TP3, expressed as a fraction.
coupler_loss: Insertion loss of a coupler, expressed as a fraction.
channel_loss: Channel insertion loss, expressed as a fraction
T1: Transmitter signal output power, in mW, binary 1
T0: Transmitter signal output power, in mW, binary 0
R1: Receiver signal input current, in mA, binary 1
R0: Receiver signal input current, in mA, binary 0
rms_noise: Receiver noise current, rms value, in mA
Threshold: Receiver decision threshold, in mA
responsivity: Photodetector responsivity, in A/W
C0: Total crosstalk current as seen by the receiver, binary 0, in mA
C1: Total crosstalk current as seen by the receiver, binary 1, in mA
Pe: Probability of error
E : extinction ratio

Signal current as seen by the receiver can now be expressed as:

R1 = T1*(coupler_loss^2)*channel_loss*responsivity

R0 = T0*(coupler_loss^2)*channel_loss*responsivity

Cross-talk current as seen by the receiver can be expressed as:

C0 = C0_near_end + C0_far_end
C1 = C1_near_end + C1_far_end

where
C1_near_end = T1*(coupler_loss^2)*return_loss*responsivity
C0_near_end = T0*(coupler_loss^2)*return_loss*responsivity
C1_far_end = T1*(coupler_loss^2)*(channel_loss^2)*return_loss*responsivity
C0_far_end = T0*(coupler_loss^2)*(channel_loss^2)*return_loss*responsivity

For simplicity, we make several assumptions. We assume that signals add incoherently at the detector – even if signals add coherently and lead to interference noise, we assume that the analysis of that effect and the resulting power penalty can be treated as if it were an independent issue. We assume that shot noise can be neglected, otherwise we would be forced to attribute a higher value of noise to binary level 1 than to binary level 0. We also assume that all noise can be referred to the transimpedance amplifier input as one equivalent noise current value, with a Gaussian probability density function. Noise will add to the various signal and cross-talk currents. The receiver will proceed to make decisions - about whether a binary 0 or 1 was transmitted - based on whether the total receiver current falls below or above the threshold, respectively. This is shown in Figure 2. For clarity, Gaussian noise is shown superimposed on only two values (R0+C1 and R1).

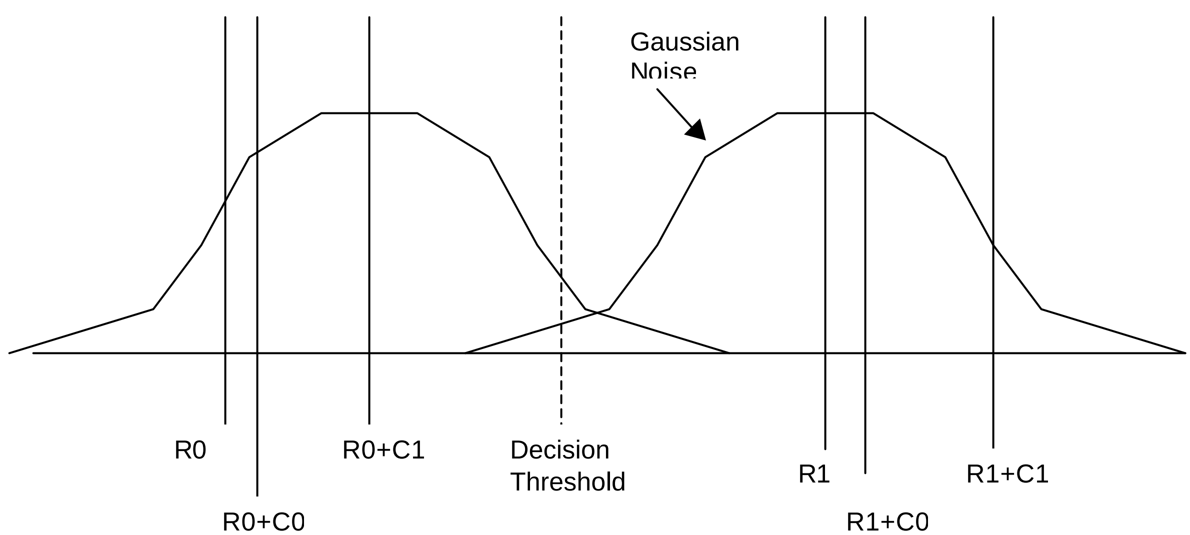

Fig. 2. A representation of signal, cross-talk, noise and threshold levels

What is the probability of error in the presence of cross-talk and noise? We assume that this is a random binary data stream, with equally likely occurrence of 1's and 0's.

Therefore,
Pe = (0.5* Pe01) + (0.5*Pe10)

where

Pe01: Probability that 1 was sent but was declared as 0. This is equal to the probability that signal plus crosstalk plus noise fell below threshold. This can happen when crosstalk bit was either 0 or 1.

Pe10: Probability that 0 was sent but was declared as 1. This is equal to the probability that signal plus crosstalk plus noise exceeded threshold. This can happen when crosstalk bit was either 0 or 1.

Therefore,

Pe01 = 0.5*(Probability that R1+C0+noise < Threshold)
 + 0.5*(Probability that R1+C1+noise < Threshold)

= 0.25*erfc((R1+C1-Threshold)/(rms_noise*sqrt(2)))
 + 0.25*erfc((R1+C0-Threshold)/(rms_noise*sqrt(2)))

Pe10 = 0.5*(Probability that R0 + C0 + noise > Threshold)
 + 0.5*(Probability that R0 + C1 + noise > Threshold)

= 0.25*erfc((Threshold - R0 - C0)/(rms_noise*sqrt(2)))
 + 0.25*erfc((Threshold - R0 - C1)/(rms_noise*sqrt(2)));

The symbol erfc stands for Complementary Error Function.

The receiver decision threshold is normally kept in the middle of R1 and R0:

Threshold = (R0 + R1)/2

Based on this threshold, the results are plotted in Figure 3. The example here assumes that the coupler loss is 3.5 dB, connector loss is 2 dB, link attenuation is 0.5 dB/km, receiver noise RMS value is 0.4 micro-amperes, and the photodetector responsivity is 0.8 A/W. For simplicity, the value of link penalties was assumed to be zero.

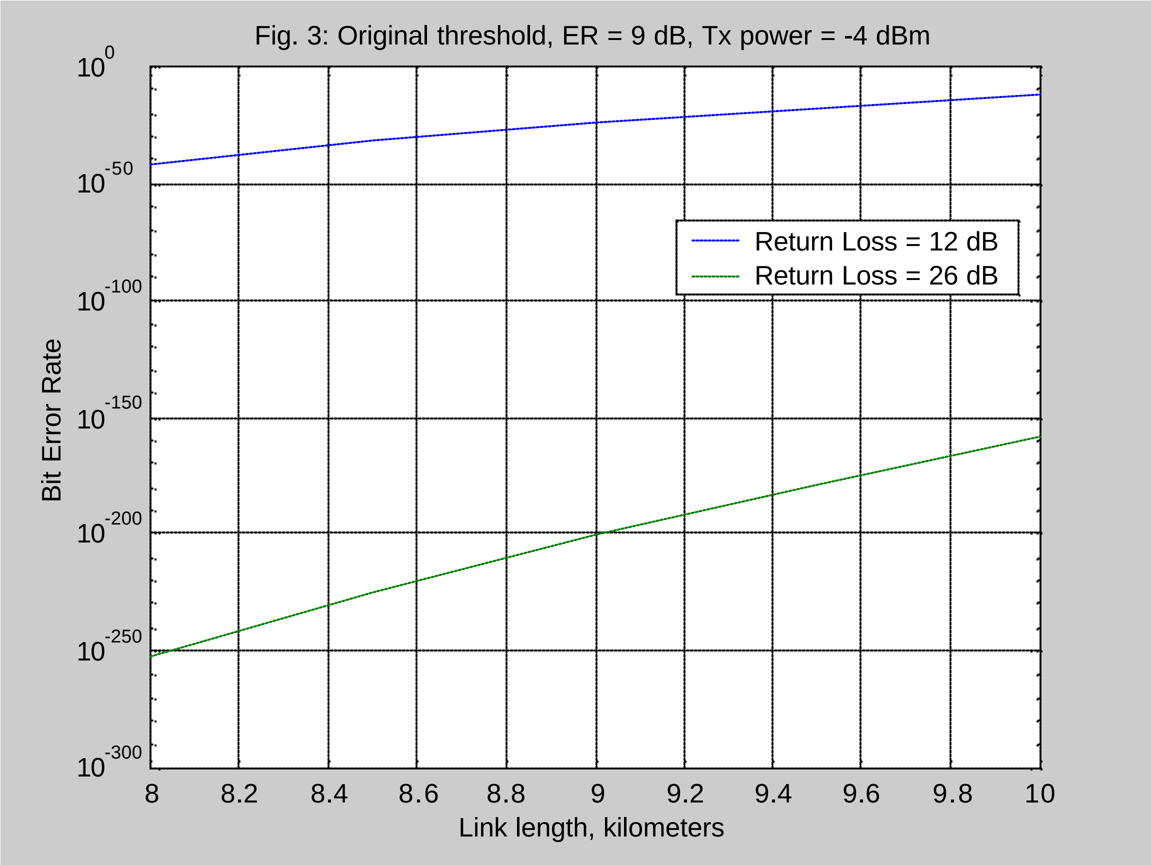

Clearly, the link performance is limited in the presence of low return loss when the link length is 10 kilometers. To achieve acceptably low BER we must increase transmitter output power to a perhaps unreasonably large value – or to specify a high return loss transceiver design.

Threshold Adjustment

A simple but effective solution to this problem is to adjust the receiver threshold. By adjusting threshold to a higher value that accounts for crosstalk, substantial improvement in performance can be achieved. This can be reasoned as follows. As demonstrated by Figure 2, if threshold is maintained at a mid-point of R0 and R1, it will be well below what it needs to be in order to make decisions with minimum probability of error. Due to crosstalk, both 0 and 1 levels have shifted higher.

Let us consider the effect of setting the threshold to a new value:

Threshold_adjusted = threshold + Offset, where
Offset = (C0_L+C1_L)/2

C0_L and C1_L are fixed values, corresponding to the C0 and C1 at full link length, respectively. In other words, the offset added to the threshold is a known fixed value, based on link parameters. In terms of circuit implementation, it should be simple enough to add a fixed voltage offset to one of the inputs pins of a differential amplifier. This offset addition should be done before the gain of the amplifier stages reaches a limiting value.

All other parameters are assumed to be the same as used for Figure 3, except that transmit power is now reduced in anticipation of much better performance. As expected, the results are remarkably better even after allowing the transmitter power to be 3 dB lower. Even in the presence of a 12 dB return loss, it is possible to use non-WDM couplers and achieve a satisfactory link performance.

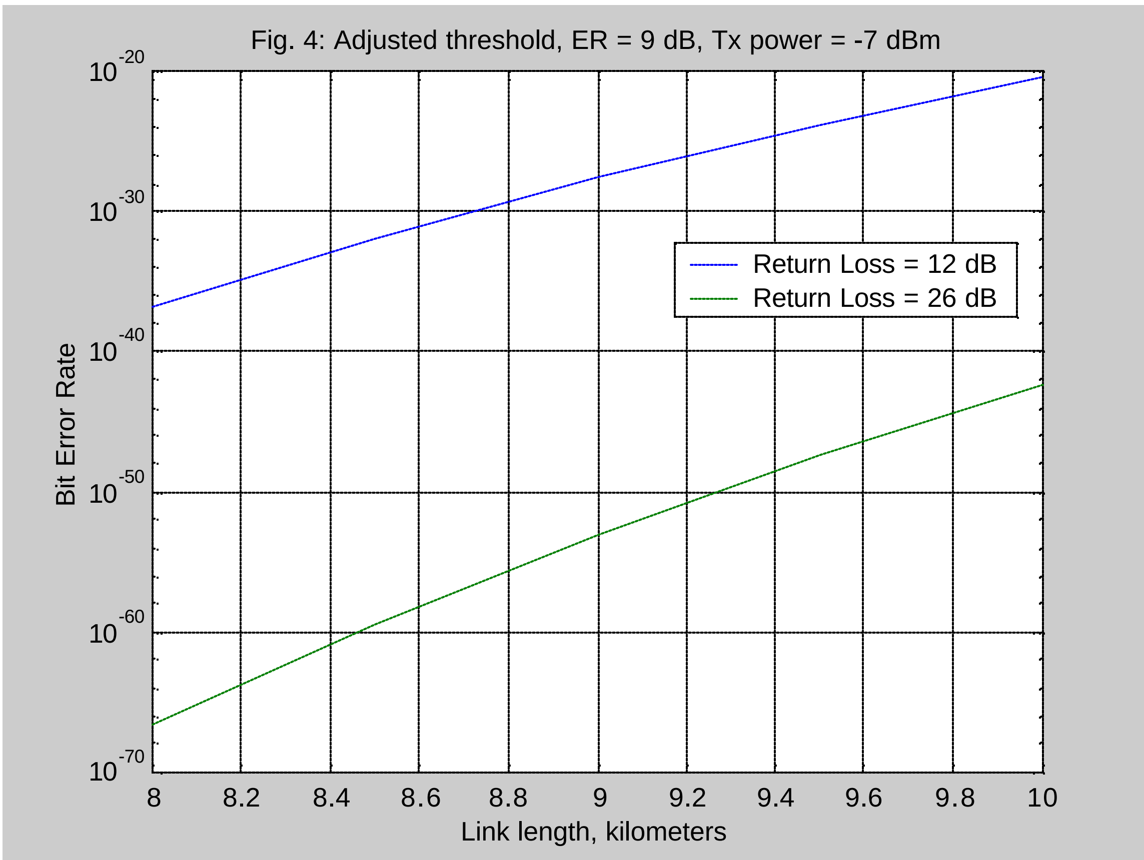

It may be possible to further optimize performance by dynamically adjusting the threshold to an optimum value depending on link length, but the incremental benefit is small, and not worth the complexity of the receiver design.

Cross-talk Penalty

We now define a performance measure that can be used to compare various design choices. We define Crosstalk Penalty as the additional transmit power required to offset the performance degradation due to crosstalk. An approximate expression for it can be derived as follows.

Let's say that in the absence of crosstalk, an acceptable value of BER is achieved for a certain distance D between the threshold and R0. For example, it is 7.04 times the RMS value of noise for a BER of 10^-12.

D = threshold – R0 = (R1 – R0) /2
= 0.5*T0*(E-1)*(coupler_loss^2)*channel_loss*responsivity

In the presence of crosstalk, we want to maintain the same distance D. But now this distance would be approximately equal to the new threshold minus R0', offset by C0' and C1' with equal probability. Here C0' and C1' are crosstalk values, and R0' represents the value of received binary 0 level resulting from the higher transmit power T0'. For simplicity of analysis, we assume that the threshold is dynamically adjusted, though in terms of power penalty, a value of fixed threshold as described above does nearly as well. Since C0' will be much smaller than C1', we can ignore that for our approximation. We also ignore the small changes in D needed to account for BER changing by a factor of 2 when C0' is ignored. The variables channel_loss and return_loss are expressed as fractions.

D = Threshold_adjusted – (R0' + C1')
= 0.5*(R0' + R1' + C0' + C1') - (R0' + C1')
= 0.5*(R1'- R0') – 0.5*(C1'-C0')
= 0.5*T0'*(E-1)*(coupler_loss^2)*channel_loss*responsivity
- 0.5*T0'*(E-1)* (coupler_loss^2)*return_loss*(1+channel_loss^2)*responsivity

Now, per our definition, Crosstalk Penalty = 10*log(T0'/T0)

By equating terms for D, and by substitution and rearrangement of terms, it can be shown that the Crosstalk Penalty is approximately given by

Crosstalk Penalty (dB) = 10*log(channel_loss/(channel_loss – return_loss))

For example, if channel insertion loss is 7 dB, then the variable channel_loss is equal to 10^(-0.1*7), or 0.1995. And for a return loss of 12 dB, the variable return_loss is equal to 0.0631. Crosstalk Penalty then is equal to 1.65 dB, a value that is reasonable to accept as an upper limit in the link budget planning. This suggests a thumb rule: Keep the return loss (in dB) at least 5 dB greater than the channel insertion loss (in dB).

Figure 5 shows a plot of Crosstalk Penalty as a function of return loss.

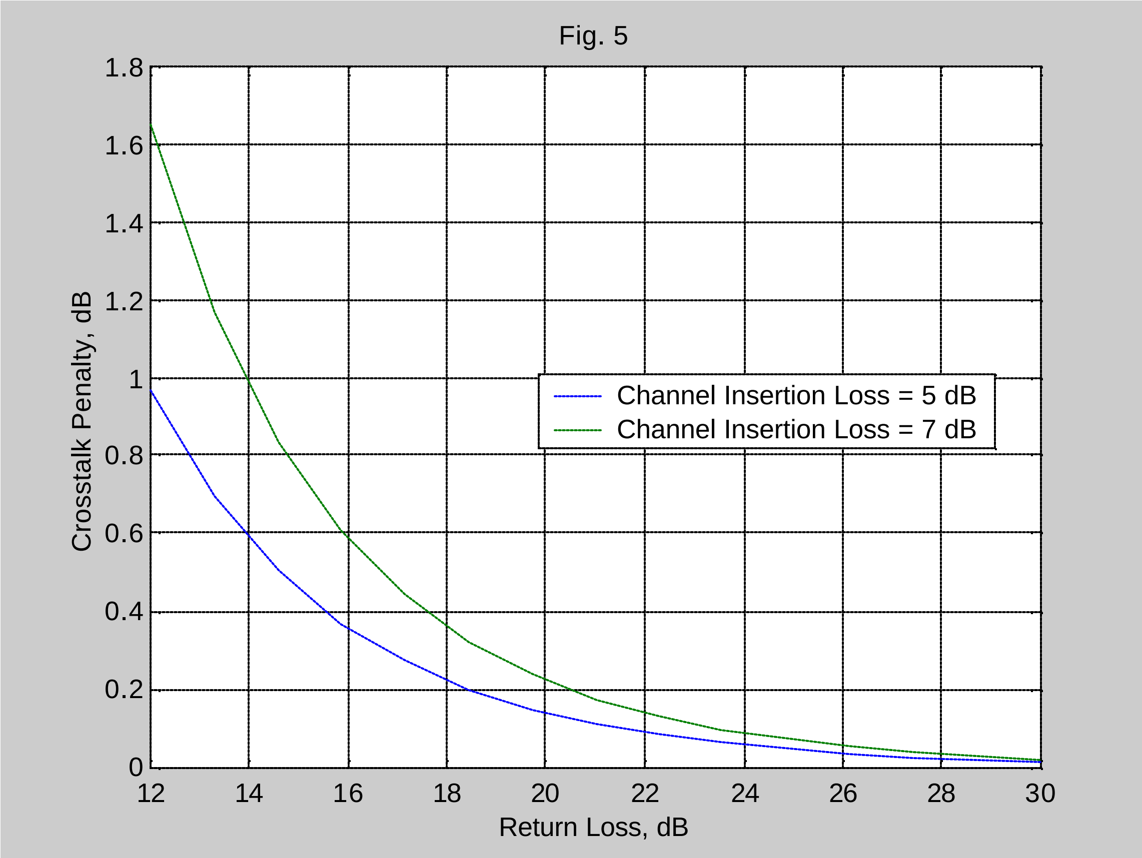

Design Recommendations

Modify the receiver design to add an offset to the receiver threshold. Keep the return loss of transceiver-cable interface to at least 5 dB higher than the total channel insertion loss, so as to keep the Crosstalk Penalty to well under 2 dB.

Example link power budget

The use of passive couplers – splitters/combiners that are not wavelength selective – can be supported in a 10 kilometer long EFM single fiber link with the following specifications, for example.

Transmitter power, average: -2 dBm
Receiver sensitivity: -22 dBm
Coupler loss: 3.5 dB*2 = 7 dB
Channel Insertion Loss = 7 dB (0.5 dB/km*10 km plus 2 dB connector loss)
Crosstalk Penalty = 1.60 dB
Other penalties: 3.5 dB
Return loss: 12 dB

Conclusion

It appears feasible to develop a low cost single-fiber EFM transceiver that can support a link length of 10 kilometers. This can be done with a single wavelength of operation. Single wavelength links overcome some of the disadvantages of using WDM links. Performance degradation due to crosstalk resulting from reflections at the transceiver-cable interface can be overcome by making a simple threshold adjustment at the receiver, and by paying a small power penalty.

References

[1] The IEEE 802.3 EFM Study Group. See http://www.ieee802.org/3/efm/