69 Introduction to Ethernet operation over electrical backplanes

69.1 Overview

69.1.1 Scope

Ethernet operation over electrical backplanes, also referred to as "Backplane Ethernet", combines the IEEE 802.3 Media Access Control (MAC) and MAC Control sublayers with a family of Physical Layers defined to support operation over a modular chassis backplane.

Backplane Ethernet supports the IEEE 802.3 MAC operating at 1000 Mb/s and/or 10 Gb/s. For 1000 Mb/s operation, the family of 1000BASE-X Physical Layer signaling systems is extended to include 1000BASE-KX. For 10 Gb/s operation, two Physical Layer signaling systems are defined. For operation over four logical lanes, the 10GBASE-X family is extended to include 10GBASE-KX4. For serial operation, the 10GBASE-R family is extended to included 10GBASE-KR.

Backplane Ethernet also specifies an Auto-Negotiation function to enable two devices that share a backplane link segment to automatically select the best mode of operation common to both devices.

Backplane Ethernet supports point-to-point topologies in the full-duplex mode of operation. Since there are no modifications to the IEEE 802.3 MAC or 1000BASE-X PCS, and the network radius is limited to the modular chassis backplane, the half-duplex mode of operation may also be supported at 1000 Mb/s.

69.1.2 Objectives

The following are the objectives of Backplane Ethernet:

a) Support the CSMA/CD MAC.
b) Provide for Auto-Negotiation among Backplane Ethernet physical layer signaling systems.
c) Meet or exceed CISPR/FCC Class A.
d) Support operation over links consistent with differential, controlled impedance traces on a printed circuit board with 2 connectors and total length up to at least 1m meeting the requirements of 69.3.
    i) a 1 Gb/s PHY
    ii) a 4-lane 10 Gb/s PHY
    iii) single-lane 10 Gb/s PHY
e) Support a BER of 10^{-12} or better.

69.1.3 Relationship of Backplane Ethernet to the ISO OSI reference model

Backplane Ethernet couples the IEEE 802.3 (CSMA/CD) MAC to a family of Physical Layers defined for operation over electrical backplanes. The relationships among Backplane Ethernet, the IEEE 802.3 MAC, and the ISO Open System Interconnection (OSI) reference model are shown in Figure 69–1.

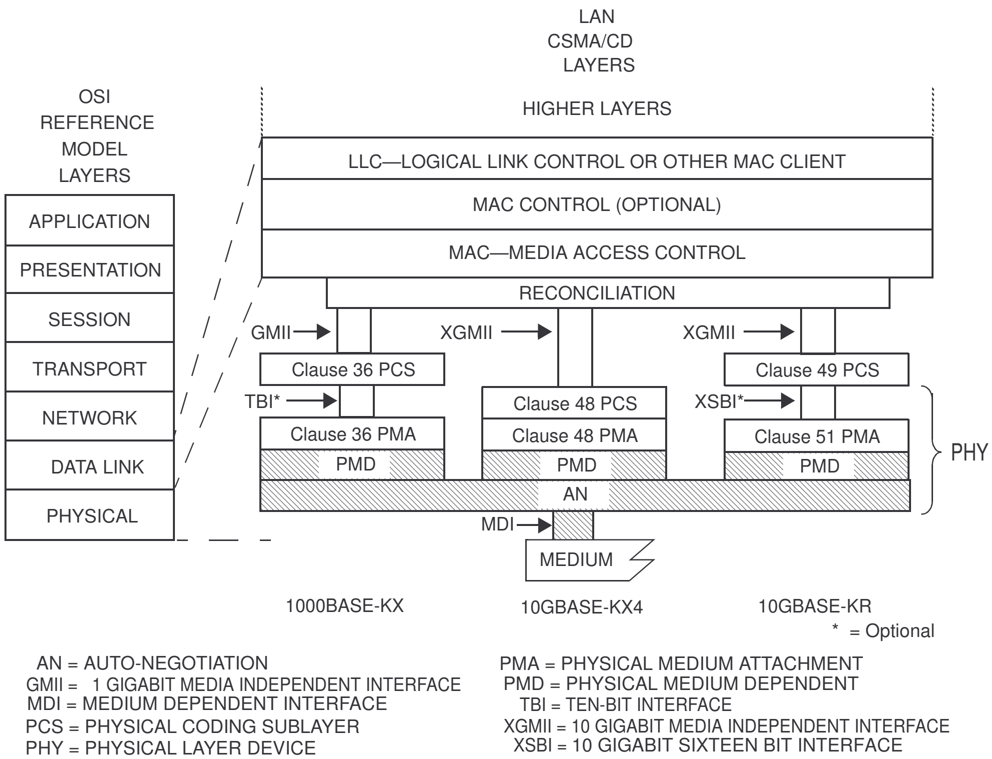

Figure 69–1—Architectural positioning of Backplane Ethernet

It is important to note that, while this specification defines interfaces in terms of bits, octets, and frames, implementers may choose other data-path widths for implementation convenience. The only exceptions are as follows:

a) The GMII, which, when implemented at an observable interconnection point, uses an octet-wide data path as specified in Clause 35.
b) The XGMII, which, when implemented at an observable interconnection point, uses a 4-octet-wide data path as specified in Clause 46.
c) The management interface, when implemented as the MDIO/MDC (Management Data Input/Output, Management Data Clock) at an observable interconnection point, uses a bit-wide data path as specified in Clause 45.
d) The 1000BASE-X PMA service interface, when implemented at an observable interconnection point, uses the 10-bit-wide data path as specified in Clause 36.
e) The PMA service interface for 10Gb/s serial, when implemented at an observable interconnection point, uses the 16-bit-wide data path as specified in Clause 51.
f) The MDI as specified in Clause 70 for 1000BASE-KX, Clause 71 for 10GBASE-KX4, or Clause 72 for 10GBASE-KR.

69.2 Summary of Backplane Ethernet Sublayers

69.2.1 Reconciliation sublayer and media independent interfaces

The Clause 35 RS and GMII, and the Clause 46 RS and XGMII, are both employed for the same purpose in Backplane Ethernet, that being the interconnection between the MAC sublayer and the PHY sublayers.

69.2.2 Management interface

The MDIO/MDC management interface (Clause 45) provides an interconnection between MDIO Manageable Devices (MMD) and Station Management (STA) entities.

69.2.3 Physical layer signaling systems

Backplane Ethernet extends the family of 1000BASE-X Physical Layer signaling systems to include 1000BASE-KX. This embodiment specifies operation at 1Gb/s over two differential, controlled impedance pairs of traces (one pair for transmit, one pair for receive). This system employs the 1000BASE-X PCS and PMA as defined in Clause 36.

Backplane Ethernet also extends the family of 10GBASE-X Physical Layer signaling systems to include 10GBASE-KX4. This embodiment is based on XAUI with 10GBASE-CX4 extensions and specifies 10Gb/s operation over four differential paths in each direction for a total of eight pairs, or sixteen connections. This system employs the 10GBASE-X PCS and PMA as defined in Clause 48.

Finally, Backplane Ethernet extends the family of 10GBASE-R Physical Layer signaling systems to include the 10GBASE-KR. This embodiment specifies 10Gb/s operation over two differential, controlled impedance pairs of traces (one pair for transmit, one pair for receive). This system employs the 10GBASE-R PCS as defined in Clause 49 and the serial PMA as defined in Clause 51.

Table 69–1 specifies the correlation between nomenclature and clauses. A complete implementation conforming to one or more nomenclatures meets the requirements of the corresponding clauses.

Table 69–1—Nomenclature and clause correlation

Nomenclature	Clause						
	36 1000BASE-X PCS/PMA	48 10GBASE-X PCS/PMA	49 10GBASE-R PCS	51 Serial PMA	70 1000BASE-KX PMD, PCS/PMA	71 10GBASE-KX4 PMD	72 10GBASE-KR PMD
1000ASE-KX	M[a]				M		
10GBASE-KX4		M				M	
10GBASE-KR			M	M			M

[a]M = Mandatory

69.2.4 Auto-Negotiation

Auto-Negotiation provides a linked device with the capability to detect the abilities (modes of operation) supported by the device at the other end of the link, determine common abilities, and configure for joint operation.

Auto-Negotiation for Backplane Ethernet is based on the Clause 28 definition of Auto-Negotiation for twisted-pair link segments. Since the connection of twisted-pair and backplane physical layer signaling systems is not expected, Auto-Negotiation for Backplane Ethernet utilizes an extended base page and next page format and modifies the timers to allow rapid convergence. Furthermore, Auto-Negotiation does not utilize Fast Link Pulses (FLPs) for link code word signaling and instead uses a signaling more suitable for electrical backplanes.

Auto-Negotiation for Backplane Ethernet is defined in Clause 73.

69.2.5 Management

Managed objects, attributes, and actions are defined for all Backplane Ethernet components. Clause 30 consolidates all IEEE 802.3 management specifications so that 10/100/1000 Mb/s and 10 Gb/s agents can be managed by existing network management stations with little or no modification to the agent code.

69.3 Interconnect Characteristics (informative)

69.3.1 Overview

Backplane Ethernet links are primarily intended as point-to-point interfaces of up to 1m using differential controlled impedance traces on low-cost printed circuit boards (PCBs) between integrated circuits residing on different boards residing in a backplane environment. The performance of such interfaces is highly dependent on the specific implementation.

69.3.1.1 Reference model

For purposes of this section, the backplane interconnect is defined between test points TP1 and TP4 as shown in Figure 69–2. The transmitter and receiver blocks include all off-chip components associated with the respective block. For example, external AC-coupling capacitors, if required, are to be included in the receiver block.

Informative characteristics and methods of calculation for the channel insertion loss, crosstalk, and ratio of channel insertion loss to crosstalk between TP1 and TP4 are defined in 69.3.3.3, 69.3.3.5 and 69.3.3.5.4.

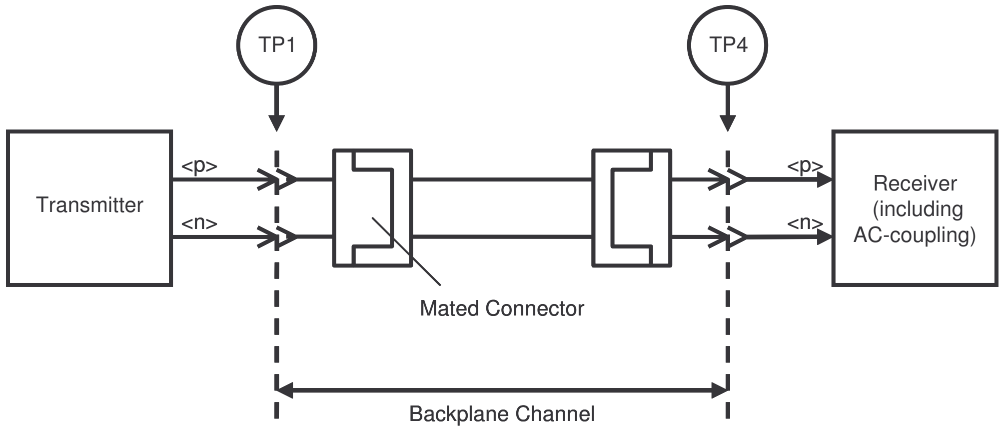

Figure 69–2— Interconnect reference model

69.3.2 Characteristic impedance

The recommended differential characteristic impedance of circuit board trace pairs is 100 Ohms +/- 10%.

The total differential skew from TP1 to TP4 is recomended to be no more than 20 ps.

69.3.3 Channel parameters

69.3.3.1 Overview

A series of informative parameters are defined for use in backplane channel evaluation.. These parameters address the channel insertion loss and crosstalk. The informative parameters for channel insertion loss are summarized in Table 69–1.

The maximum amount of attenuation (A_{max}) due to trace skin effect and dielectric properties is defined in 69.3.3.2. The maximum amount of insertion loss (IL_{max}) is defined in 69.3.3.3. The maximum deviation of insertion loss from the best-fit attenuation (ILD) is defined in 69.3.3.4. The limit on crosstalk in relation to insertion loss (ICR) is defined in 69.3.3.5.4.

To enable system tradeoffs for the designer a series of confidence curves have been created for the different parameters. All of the different parameters must be considered together in evaluating the overall channel performance.

.

Table 69–1—Insertion loss parameters

Parameter	1000BASE-KX	10GBASE-KX4	10GBASE-KR	Units
f_{min}	0.05			GHz
f_{max}	15.00			GHz
b_1	2.25E-5			
b_2	1.20E-10			
b_3	3.50E-20			
b_4	–1.25E-30			
f_1	0.125	0.312	1.000	GHz
f_2	1.250	3.125	6.000	GHz
Maximum insertion loss, IL_{max1}	Refer to Equation (69–7) and Equation (69–8).			
High-confidence maximum insertion loss, IL_{max2}	Refer to Equation (69–9) and Equation (69–10)			
Minimum insertion loss deviation, ILD_{min1}	Refer to Equation (69–12).			
High-confidence minimum insertion loss deviation, ILD_{min2}	Refer to Equation (69–13).			
Maximum insertion loss deviation, ILD_{max1}	Refer to Equation (69–14).			
High-confidence maximum insertion loss deviation, ILD_{max2}	Refer to Equation (69–15).			

69.3.3.2 Attenuation

The attenuation, A(f), is defined to be least mean squares line fit to the insertion loss computed over the frequency range f_1 to f_2. Given that insertion loss is measured at N frequencies, f_n, from f_1 to f_2, the least mean squares line fit procedure is defined by Equations (69–1) through (69–5).

$$f_{avg} = \frac{1}{N}\sum_{n} f_n \quad (69\text{–}1)$$

$$IL_{avg} = \frac{1}{N}\sum_{n} IL(f_n) \quad (69\text{–}2)$$

$$m = \frac{\sum_n (f_n - f_{avg})(IL(f_n) - IL_{avg})}{\sum_n (f_n - f_{avg})^2} \quad (69\text{–}3)$$

$$b = IL_{avg} - mf_{avg} \quad (69\text{–}4)$$

$$A(f) = mf + b \quad (69\text{–}5)$$

It is recommended that attenuation of the channel be greater than the worst-case attenuation limit described by the equation:

$$A(f) \le A_{max}(f) = 20\log(e) \times (b_1\sqrt{f} + b_2 f + b_3 f^2 + b_4 f^3) \quad (69\text{–}6)$$

where f is expressed in Hz and the coefficients b_1 through b_4 are given in Table 69–1.

In addition, it is recommended that the insertion loss also satisfy the insertion loss limit defined in 69.3.3.3, the insertion loss deviation limit defined 69.3.3.4, and the insertion loss to crosstalk ratio limit defined in 69.3.3.5.4. The attenuation limit for each port type is illustrated in Figures 69–3, 69–4 and 69–5.

69.3.3.3 Insertion loss

The insertion loss is defined as the magnitude, expressed in decibels, of the differential response measured from TP1 to TP4. It is recommended that the insertion loss magnitude, IL(f), be greater than the lower limit defined by Equation (69–7) and Equation (69–8).

$$IL(f) \le IL_{max1}(f) = A_{max}(f) + 1.1 + 9.0\times10^{-10} f \quad (69\text{–}7)$$

for $f_{min} \le f \le f_2$

$$IL(f) \le IL_{max1}(f) = A_{max}(f) + 1.1 + 9.0\times10^{-10} f_2 + 2\times10^{-8}(f - f_2) \quad (69\text{–}8)$$

for $f_2 < f \le f_{max}$

The values of f_{min}, f_2, are f_{max} are given in Table 69–1 and the attenuation limit, $A_{max}(f)$, is given in Equation (69–6).

A high confidence acceptance region for IL(f) is constrained by Equation (69–9) and Equation (69–10).

$$IL(f) \le IL_{max2}(f) = A(f) + 0.8 + 2.0\times10^{-10} f \quad (69\text{–}9)$$

for $f_{min} \le f \le f_2$

$$IL(f) \le IL_{max2}(f) = A(f) + 0.8 + 2.0\times10^{-10} f_2 + 1\times10^{-8}(f - f_2) \quad (69\text{–}10)$$

for $f_2 < f \le f_{max}$

In addition, it is recommended that the insertion loss also satisfy the attenuation limit defined in 69.3.3.2, the insertion loss deviation limit defined 69.3.3.4, and the insertion loss to crosstalk ratio limit defined in 69.3.3.5.4. The insertion loss limit is illustrated in Figures 69–3, 69–4 and 69–5.

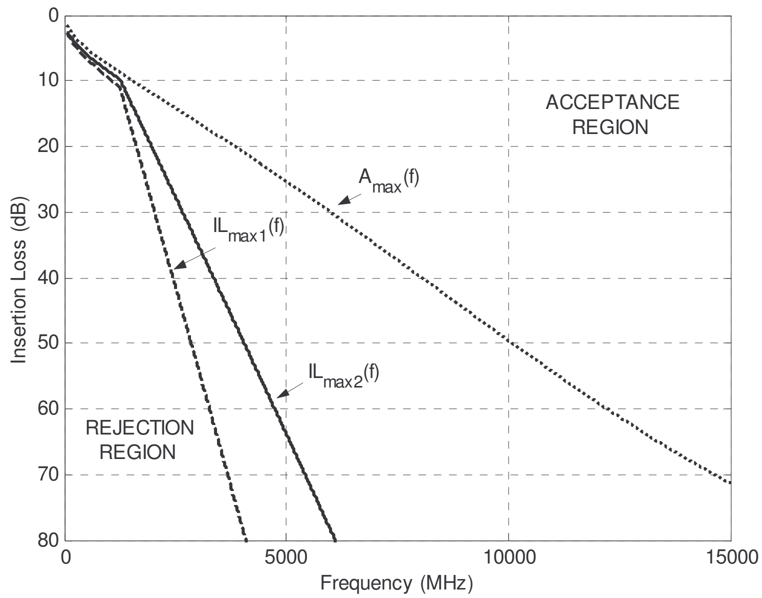

Figure 69–3—Insertion loss and attenuation limits for 1000BASE-KX

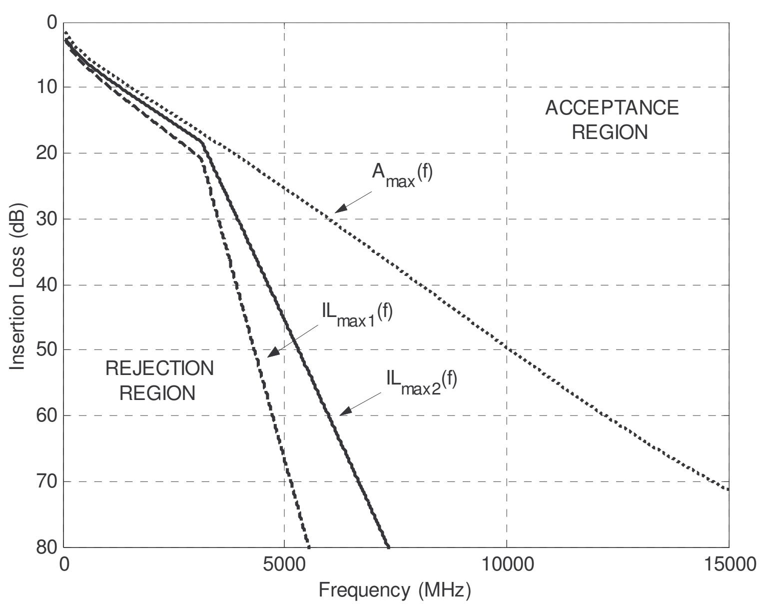

Figure 69–4—Insertion loss and attenuation limits for 10GBASE-KX4

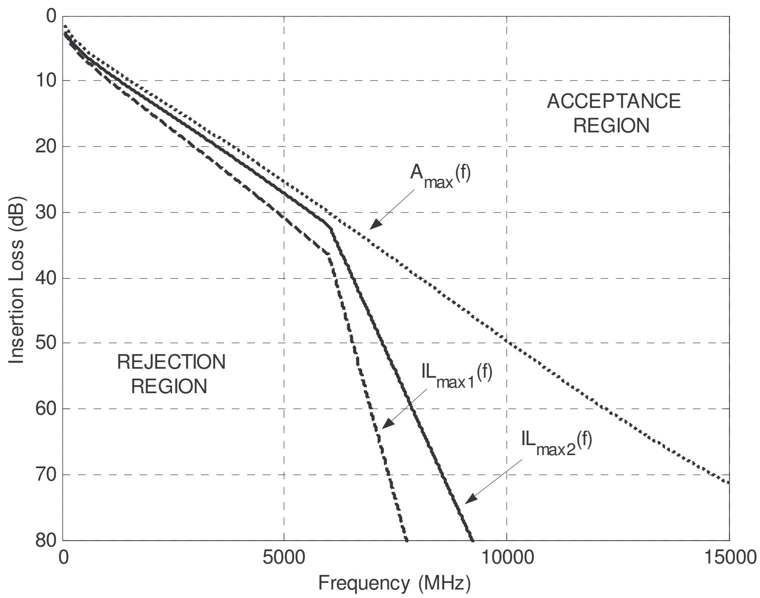

Figure 69–5— Insertion loss and attenuation limits for 10GBASE-KR

69.3.3.4 Insertion loss deviation

The insertion loss deviation, as defined by Equation (69–11), is the difference between the insertion loss and the least mean squares fit defined in 69.3.3.2.

$$ILD(f) = IL(f) - A(f) \quad (69\text{–}11)$$

The insertion loss deviation, ILD, is recommended to be constrained within the limits defined by Equation (69–12) and Equation (69–13):

$$ILD(f) \geq ILD_{min1}(f) = -2.1 - 0.9\times10^{-9}f \quad (69\text{–}12)$$

$$ILD(f) \leq ILD_{max1}(f) = 1.5 + 1.0\times10^{-9}f \quad (69\text{–}13)$$

for $f_1 \leq f \leq f_2$.

A high confidence acceptance region for ILD is constrained by Equation (69–14) and Equation (69–15):

$$ILD(f) \geq ILD_{min2}(f) = -1.0 - 0.5\times10^{-9}f \quad (69\text{–}14)$$

$$ILD(f) \leq ILD_{max2}(f) = 1.0 + 0.5\times10^{-9}f \quad (69\text{–}15)$$

for $f_1 \leq f \leq f_2$.

The values of f_1 and f_2 are dependent on port type and are given in Table 69–1.

In addition, it is recommended that the insertion loss also satisfy the attenuation limit defined in 69.3.3.2, the insertion loss limit defined in 69.3.3.3, and the insertion loss to crosstalk ratio limit defined in 69.3.3.5.4. The insertion loss deviation limit for each port type is illustrated in Figure 69–6.

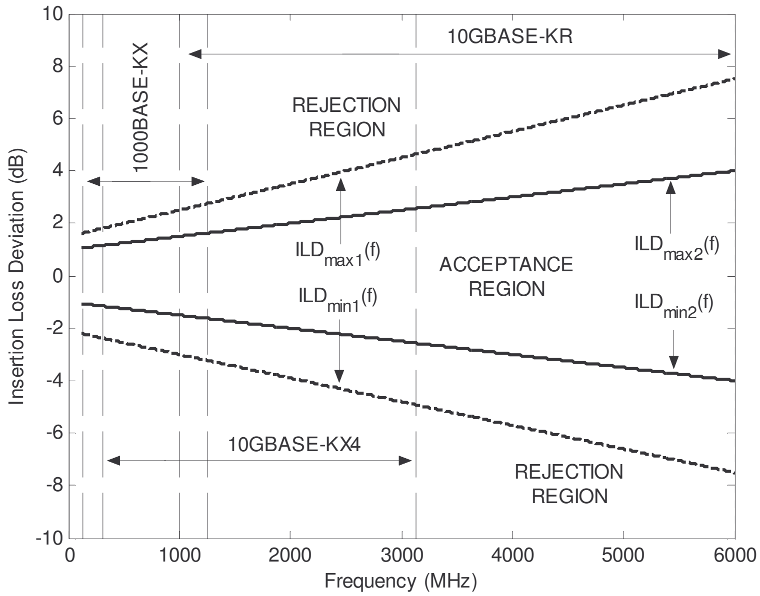

Figure 69–6— Insertion loss deviation limits

69.3.3.5 Crosstalk

In order to limit the crosstalk at the receiver of the link segment, the total differential crosstalk due to NEXT and FEXT aggressors has been specified in order to meet the BER objective stated in 69.1.2.

69.3.3.5.1 Total Differential Near-End Crosstalk (TNEXT)

The total differential near-end crosstalk at the receiver (in dB with f in MHz) is calculated as the power sum of the individual NEXT aggressors. For cases of a single NEXT aggressor, the total differential near-end crosstalk at the receiver will be equal to the individual aggressor (in dB with f in MHz).

$$TNEXT(f) = 10\log\left(\sum_n \left|NEXT_n(f)\right|^2\right) \quad (69–16)$$

69.3.3.5.2 Total Differential Far-End Crosstalk (TFEXT)

The total differential far-end crosstalk at the receiver (in dB with f in MHz) is calculated as the power sum of the individual FEXT aggressors. For cases of a single FEXT aggressor, the total differential far-end crosstalk at the receiver will be equal to the individual aggressor (in dB with f in MHz).

$$TFEXT(f) = 10\log\left(\sum_n \left|FEXT_n(f)\right|^2\right) \quad (69–17)$$

69.3.3.5.3 Total Differential Crosstalk

The amount of differential crosstalk at the receiver from NEXT and FEXT aggressors (in dB with f in MHz) is calculated as the power sum of the total differential aggressor near-end crosstalk, Equation (69–16), and the total differential far-end crosstalk, Equation (69–17).

$$PSXT(f) = 10\log(10^{TNEXT(f)/10} + 10^{TFEXT(f)/10}) \quad (69\text{–}18)$$

69.3.3.5.4 Insertion Loss to Crosstalk Ratio (ICR)

The ratio of the insertion loss, IL, measured from TP1 to TP4 to the total differential crosstalk (in dB with f in MHz) is calculated using Equation (69–19).

$$ICR(f) = IL(f) - PSXT(f) \quad (69\text{–}19)$$

The ICR at the receiver (in dB with f in MHz) is recommended to be at least:

$$ICR(f) \geq 12.5 - 20\log\left(\frac{f}{5GHz}\right) \quad (69\text{–}20)$$

For 1000BASE-KX, the range of f is 100 MHz to 1250 MHz. For 10GBASE-KX4, the range of f is 100 MHz to 3125 MHz. For 10GBASE-KR, the range of f is 100 MHz to 5000 MHz.

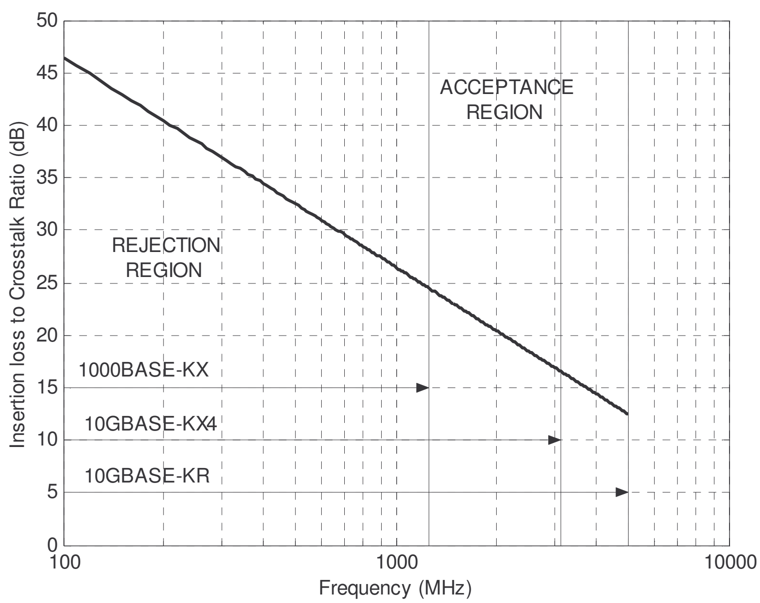

Figure 69–7—Insertion loss to crosstalk ratio limit

69.4 Delay constraints

Predictable operation of the MAC Control PAUSE operation (Clause 31, Annex 31B) demands that there be an upper bound on the propagation delays through the network. This implies that MAC, MAC Control sublayer, and PHY implementers must conform to certain delay maxima, and that network planners and administrators conform to constraints regarding the cable topology and concatenation of devices.

Table 69–1 contains the values of maximum sublayer round-trip (sum of transmit and receive) delay for the 1000BASE-KX port types in bit time as specified in 1.4.

Table 69–1—Round-Trip Delay Constraints for 1000BASE-KX (informative)

Sublayer	Maximum (bit time)	Notes
MAC Control, MAC, and RS	696	
1000BASE-X PCS and PMA	320	See 36.5.1 and 38.1.1
1000BASE-KX PMD	32	See 70.3

Table 69–2 contains the values of maximum sublayer round-trip (sum of transmit and receive) delay for the 10GBASE-KX4 and 10GBASE-KR port types in bit time as specified in 1.4 and pause_quanta as specified in 31B.2.

Table 69–2—Round-Trip Delay Constraints for 10GBASE-KX4 and 10GBASE-KR (informative)

Sublayer	Maximum (bit time)	Maximum (pause_quanta)	Notes
MAC Control, MAC, and RS	8192	16	See 46.1.4
XGXS and XAUI	4096	8	Round-trip of 2 XGXS and trace for both directions, see 47.2.2
10GBASE-X PCS and PMA	2048	4	See 48.5
10GBASE-R PCS	3584	7	See 49.2.15
10GBASE-KX4 PMD	512	1	See 71.3
10GBASE-KR PMA and PMD	512	1	See 72.3

69.5 State diagrams

State machine diagrams take precedence over text.

The conventions of 1.2 are adopted, along with the extensions listed in 21.5.

69.6 Protocol Implementation Conformance Statement (PICS) proforma

The supplier of a protocol implementation that is claimed to conform to any part of IEEE 802.3, Clauses 70 through 73, demonstrates compliance by completing a Protocol Implementation Conformance Statement (PICS) proforma.

A completed PICS proforma is the PICS for the implementation in question. The PICS is a statement of which capabilities and options of the protocol have been implemented. A PICS is included at the end of each clause as appropriate. Each of the Backplane Ethernet PICS conforms to the same notation and conventions used in 100BASE-T (see 21.6).