1.1. Ethernet Twisted Pair PHY's: A Brief History

Historically, Ethernet twisted pair physical layer (PHY) technology has evolved from simple bi-level digital pulse signaling (e.g. 10BaseT) to more complex, higher density, wider bandwidth encoding schemes that depend upon sophisticated digital processing technologies to both generate and recover digital traffic (e.g. 1000BaseT). The "Base" terminology refers to the fact that Ethernet twisted pair PHY's are baseband signals rather than modulated carriers that are typically found in broadband network connections.

Under IEEE 802.3, strict rules regarding backward compatibility and technology coexistence have assured smooth and seamless transitions to newer and faster networking technologies. The introduction of Auto-Negotiation with 100BaseTx enabled link partners to determine technology compatibility so that links start up with the highest possible performance level. Auto-Negotiation has since evolved to resolve numerous link features between link partners.

On the cabling side, EIA/TIA Category 3 structured cabling created the electrical conduit for 10BaseT. This was a low cost transmission medium with capability to connect link partners up to 100 meters apart. Category 3 cabling offered at least two twisted pairs to support transmission in each direction. Primary RF characteristics included insertion loss (s21), return loss (s11), and crosstalk. These parameters were verified out to 16MHz. 10BaseT's transmission spectrum is tightly spread around 10MHz given the Manchester coding of digital data.

Category 5 structured cabling generally included 4 pairs with parametric verification through 100MHz. This performance enabled a 100BaseTx technology that was largely borrowed from FDDI (TP-PMD). The 100BaseT coding scheme, MLT-3, spread the RF bandwidth of a transmitted signal well beyond 50 MHz, thus requiring the added transmission bandwidth in order to support the 100 meter link length objective. Like 10BaseT, 100BaseTx only required 2 pairs, one for transmission in either direction. This left two "spare pairs" in each cabling path.

While 100BaseTx extended both frequency spectra and information density, 1000BaseT took advantage of the four pair cabling structure and then combined that with a duplex transmission scheme that simultaneously transmits signals in both directions on the same twisted pair. Coupling that with higher density coding schemes, data bandwidth increased an order of magnitude with very little spectral impact. On the cabling side, Category 5e was introduced to slightly improve on attenuation (insertion loss), return loss, and crosstalk specifications. In addition, more crosstalk measurements were defined owing to added sensitivity of 1000BaseT to crosstalk.

10GbaseT was introduced more recently and is largely utilized as a backbone technology at this time. This manual will focus on 10/100/1000BaseT, the predominant Ethernet edge access technologies that are addressed by the PhyView Analyzer.

1.1.1. Ethernet Twisted Pair PHY Specification Summary

The following table briefly summarizes a few of the more important specifications related to the 10BaseT, 100BaseTx, and 1000BaseT physical layers.

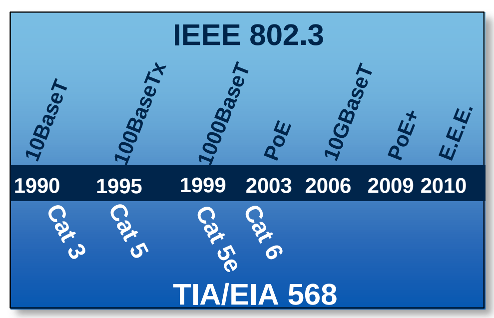

Figure 1.1 IEEE 802.3 & TIA/EIA 568 Timeline

Characteristic	10BaseT IEEE 802.3 Clause 14	100BaseTx IEEE 802.3 Clause 25	1000BaseT IEEE 802.3 Clause 40
Transmission Pairs	1 (unidirectional)	1 (unidirectional)	4 (bi-directional)
Signal Coding	Manchester	MLT-3 (tri-level)	PAM-5 (5-level)
Signal Amplitude	2.5 Vpp	2 Vpp	2 Vpp
Amplitude Band	± 0.3 Vpp	± 0.1 Vpp	± 10.7%
Symbol Rate (Baud)	10 MHz ± .01%	125 MHz ± .005%	125 MHz ± .01%
Symbol Density	1 Bit	1 Bit	8 Bits
Edge Rate Strategy	Manchester, 1 edge / bit	Randomizer	Randomizers
Edge Slew Rate	Mask Fit (typical 5 nsec)	4 ± 1 nsec	Mask Fit (*typical 3.6 nsec*)

Characteristic	10BaseT IEEE 802.3 Clause 14	100BaseTx IEEE 802.3 Clause 25	1000BaseT IEEE 802.3 Clause 40
PHY Layer Error Detection	None	4b/5b Code	Trellis
Forward Error Correction	None	None	Trellis
Differential Impedance	100Ω +~20Ω	100Ω +~16Ω ≤ 30MHz 100Ω +64Ω,-40Ω @ 60MHz	100Ω +~17Ω ≤ 40MHz 100Ω +97Ω,-50Ω @ 100MHz
Interface Return Loss	≤ ~ -20 dB	≤~ -21.5 dB ≤ 30MHz ≤~ -12 dB ≥ 60MHz	≤~ -21.5 dB ≤ 40MHz ≤~ -9.5 dB @ 100MHz
Loss Tolerance	-11.5 dB @ 10MHz	-24 dB @ 100 MHz	-24 dB @ 100 MHz
Mismatch Tolerance	±15 Ω @ < -15 dB R.L.	-15 dB R.L. to 20MHz -8 dB R.L @ 100MHz	-15 dB R.L. to 20MHz -8 dB R.L @ 100MHz
NEXT Crosstalk Tolerance	-26 dB @ 10MHz	-27.1 dB @ 100MHz	-27.1 dB @ 100MHz
Noise Tolerance	264 mVpk (*below 15MHz*)	40mVpp (*below 100MHz*)	40mVpp (*below 100MHz*)
Jitter Tolerance	7 nsec Pk-Pk	1.4 nsec Pk-Pk	1.4 nsec Pk-Pk, total 0.3 nsec Pk-Pk > 5KHz
Pair Delay Skew	Not Applicable	Not Applicable	50 nsec (to 100MHz)
Bit Error Target	10^{-8}	10^{-8}	10^{-10}

1.1.2. Ethernet Symptom Masking

Ethernet links are universally designed to recognize packet transmission problems both in the Physical and in the MAC layer, and when encountering a problem, to drop or discard the erred packet or packet segment. At the physical layer, packet errors appear as defective preambles (or carrier markers) meaning packets that don't properly start or complete. In 100BaseT, coding violations are used as an indication of a flawed transmission. In 1000BaseT, sophisticated digital signal processing can both recognize transmission errors and attempt to correct them before packet data reaches the MAC layer.

In the MAC layer, CRC checking determines if there is a probable data corruption in the packet (or the CRC code). If corruption is discovered, the packet is dropped.

Many Ethernet ports have intelligence to monitor link performance and make decisions to drop link and re-link, and in certain cases, to force a fall-back into a lower rate in order to get better immunity from link impairments. For example, a link that starts out at 1000BaseT may eventually drop down to 100BaseT to overcome those impairments.

On the Internet level, the majority of traffic is governed by a transport layer, most commonly TCP, that has intelligence to know when packets are missing because they were dropped in some link along the network path. TCP can then communicate to the original sender that certain packets need to be re-transmitted, thus giving them all a second chance to traverse the network path.

The bottom line is that it is difficult for a user or observer to know when an Ethernet link gets into trouble unless the problem is so severe that the link is essentially dead. Problems in physical layer communication can be concealed by an array of self-healing behaviors both in the link and in the network levels.

1.1.3. Traditional PHY Conformance Testing

To some degree, IEEE 802.3 specifications have taken ownership for defining the types of tests and test methods that should be applied to verify conformance to each physical layer standard. This is particularly true of 10BaseT and 1000BaseT. 100BaseTx was largely imported from the FDDI TP-PMD (ANSI X3:263 1995) and did not incorporate as much in the way of testing details.

Much of the focus on testing in these documents related to measurements of PHY transmitter characteristics. This is not surprising because those are far easier to measure than are PHY receiver characteristics. Common PHY transmitter parameters that appear throughout these specifications include:

- Signal Amplitude (min & max)
- Transition Slew Rate or Time (min & max)
- Signal Symmetry in Amplitude and/or Time (min)
- Overshoot & Droop (max)
- Edge Jitter (max)
- Return Loss and/or Differential Impedance

Figure 1.3 Specialized Transmitter Test Signals

Receiver testing, when described, is usually limited to Bit Error Rate (BER, or Packet Error Rate, PER) checks. Each 802.3 PHY specification does clearly spell out the types of channel impairments that those physical layers must assume to exist, meaning that in theory, receiver testing and qualification must be done using representative worst case channel impairments. However, the actual testing methods and measurements are loosely defined. Similarly, the 802.3 specifications give wide design discretion to topics such as receiver error detection and indication criteria.

1.1.4. Time Domain Measurements in PHY Conformance Testing

Originally, the baseband nature of 10BaseT signals meant that time domain analysis using oscilloscopes was the natural approach to qualifying transmitted signal integrity to the standard. Most of the above listed transmitter parameters (*see Section 1.1.3*) were a good fit with time domain analysis. With the advent of digital scopes, pulse masks could readily be defined and applied to captured waveforms. Both IEEE 802.3 clause 14 (10BaseT) and clause 40 (1000BaseT) define test procedures that include a detailed pulse mask.

The only inconvenience in testing twisted pair Ethernet signals with an oscilloscope was that typical scope interfaces are either 50Ω or 1MΩ coaxial while each Ethernet twisted pair was a balanced 100Ω channel. This problem could readily be resolved with a matching transformer (or balun) that would convert balanced 100Ω to coaxial 50Ω, the only limitation being that the balun now became a part of the device-under-test.

An alternative interface method was to take each conductor of the twisted pair and route them as center (signal) conductors to a 2-channel oscilloscope whereby the signal analyzed would be the difference between the two channel input. This way, the 50Ω inputs could combine to a total 100Ω termination on the twisted pair, though the pair would no longer be electrically isolated. A third alternative would be to terminate the twisted pair in a nominal 100Ω load, then probe the load with a high impedance, high bandwidth differential (active) probe. This approach keeps the twisted pair isolated but adds cost and possible calibration of the active differential probe.

The good news about interfacing to the original 10BaseT transmitter was that there was only a single pair to ever be concerned with on any given PHY interface. This changed when auto-sensing MDI/MDI-X was introduce whereby a 10/100BaseT PHY would have to be able to transmit on either of two pairs meaning in effect that there were two transmitters that needed to be tested. 1000BaseT then complicated this further since all 4 pairs possess transmitters meaning that more elaborate fixturing was required with more manual connection interventions to perform measurements across the 4 pairs(*see Figure 1.2*). This increased both the equipment and labor costs associated with testing.

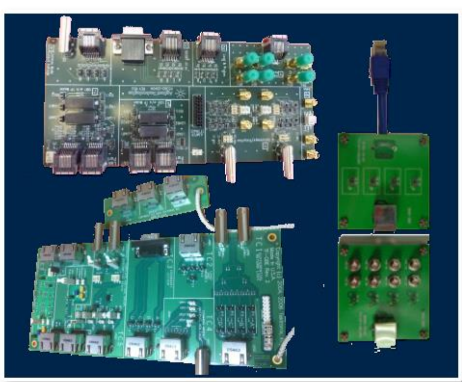

Figure 1.2 Ethernet PHY Test Fixtures

A second dilemma associated with traditional transmitter testing came along with 100BaseTx. Most 100BaseTx PHY's were actually 10/100BaseT and used auto-negotiation with a link partner to activate a 100BaseT transmission. When a transmitter pair is connected to a test instrument, there is no link partner and therefore, there is no 100BaseTx transmitted signaling. This meant that the Ethernet port-under-test either needed some special test mode to activate a 100BaseT transmission, or it needed to start from a linked state with a 100BaseT link partner, then have the transmit pair separately disconnected so it could get routed to the test instrument. IEEE 802.3 did not define any special test mode registers so in many cases, especially when testing finished product, the only option was the latter scheme.

1000BaseT made this even more difficult. When linked, each pair simultaneously transmits baseband signaling in both directions, using local hybrid cancellation on each pair end to sort received signal from transmitted signal. Additionally, PAM5 signaling is digitally filtered, or pre-distorted, in order to shape it to the 100BaseT spectra. So what shows up on each twisted pair is unintelligible. IEEE 802.3 clause 40 addressed this be defining standard test signals, standard registers to activate those test signals, and standard methods for testing those test signals (*see Figure 1.3*). This assured that PHY transmission measurements would be *unavailable* given typical finished product.

1.1.5. Packet Testing: Pros and Cons

Given the complexity, cost, and expertise requirements of traditional PHY conformance testing, it is no surprise that this type testing is done only in design verification, or at a very low sample rate basis, or in a test services setting such as UNH-IOL. Higher volume testing required a much simpler method and in many cases, that method was packet flow testing. In fact, many engineers and technicians equate packet flow testing with physical layer testing even though packet flow testing is purely functional in nature and does not directly measure *any* physical layer parameters.

Packet flow testing is attractive because it is easy and often low cost. Since Ethernet links universally discard defective packets, the core requirement is a device that can transmit, receive and count Ethernet packets. Commercial packet analyzers from manufacturers such as Spirent and Ixia have been widely deployed for years to help with Ethernet interface and protocol testing. Lower cost, software-based solutions and even Ethernet switches with modified firmware can serve the purpose of a packet generator / counter.

As a method for evaluating the physical layer, packet flow testing is in truth quite limited in utility and defect coverage for several reasons:

Limitation #1: It is inherently difficult to resolve packet loss to defects in a transmitter versus a receiver. If two **unknown ports** are linked and packets drop in one or both directions, there is no information to suggest whether this is a transmission problem or a receiver problem. On the other hand, if a **known good port** is connected to an **unknown port** and packets are pinged (meaning the port–under-test has a full protocol stack including ICMP available), any packet loss is still ambiguous as to whether the **unknown port** failed to receive them properly or corrupted the transmission back. This limitation can be partially overcome by accepting Limitation #2.

Limitation #2: The device under test must be a hub, switch, router, or repeater with 2 or more ports, one of which is a **known good port**. Many of the twisted pair Ethernet ports ever produced are in fact multi-port devices that can bridge traffic between ports. This allows **known good** (or *instrument grade*) ports to be connected to the **known good port** and to any other **unknown port** so that the direction of packet loss on the **unknown port** should point to a transmission versus receiver problem.

Limitation #3: The degree of bit errors in a link *may* provide *hints* to the severity of a problem but does not shed any light on the nature of the problem. *Figure 1.4* depicts the types of relationships found when comparing Bit Error Rate (BER) measurements to various physical impairments or defects. From this figure, it is apparent that bit errors typically do not appear until the degree of Physical Impairment crosses some threshold level that will depend on the type of defect or impairment. At that point, the rate of errors will increase at some slope that is also dependent upon the type of impairment involved. As the graph suggests, a transmit jitter problem may cause a very steep slope from zero to high Bit Error Rate while reduction in signal power (or higher channel loss) may have a much more gradual impact. Summarizing, knowledge of Bit Error Rate does not point to the nature of a defect nor does it really describe the magnitude of an impairment.

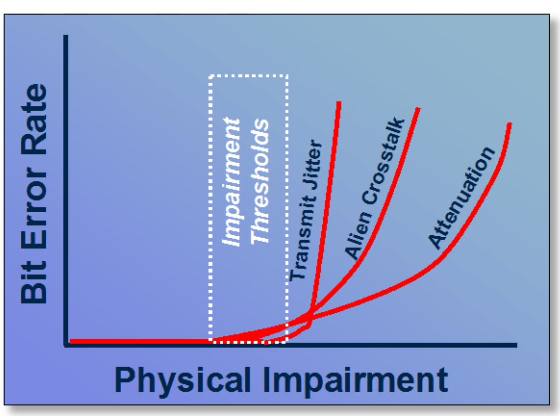

Figure 1.4 BER versus Physical Impairment

Limitation #4: Packet flow does not directly measure Bit Error Rate. There is a statistical relationship between packet loss and BER and the effect of that relationship, as described in *Figure 1.5*, is to dramatically sharpen, or increase the slope when Packet Loss is related to Physical Impairment level. Smaller packets help soften this relationship a bit, however Ethernet restricts the minimum packet size to 64 bytes including packet overhead. This translates to 512 bits for 10BaseT and to 640 bits for 100BaseTx and 1000BaseT owing to coding overheads. The BER resolution problem is worsened by 1000BaseT where a degree of forward error correction built into the physical layer will hide lower bit error levels by reconstructing partially damaged packets. This in effect increases the impairment threshold (see Figure 1.4) and sharpens the slope relating each Physical Impairment to Packet Loss.

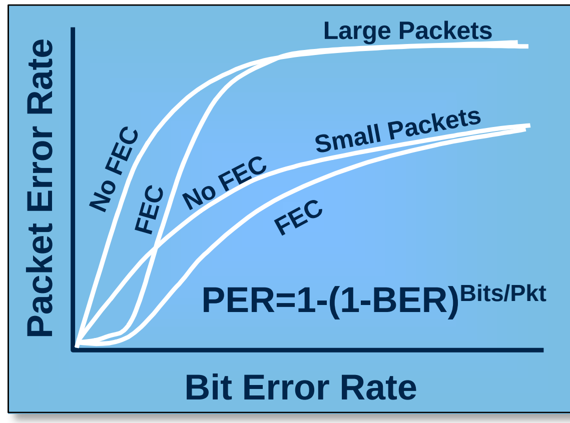

Figure 1.5 Packet Error Rate vs Bit Error Rate

Limitation #5: When performing receiver testing, unless the degree of Physical Impairment is carefully manipulated, it is hard to know whether the Ethernet port-under-test is capable of tolerating IEEE 802.3 specified maximum channel impairment levels. Measuring zero Packet Loss at zero impairment provides no insight to this capability and therefore represents minimal defect coverage. On the other hand, if a receiver is exposed to specified maximum channel impairments of various types and Packet Loss is zero, then it can be determined that the receiver is indeed immune to the worst-case types of channel impairments it can expect to encounter while in service.

Summarizing, packet flow testing has thrived because it is convenient, low cost, and because the traditional PHY compliance tests are both expensive and laborious. The PhyView Analyzer was developed to provide much of the direct physical layer test coverage associated with traditional PHY compliance tests while providing all of the convenience of a fully automated packet flow test.

1.2. The PhyView Analyzer Metrology

The PhyView Analyzer is designed to make "pure" physical layer measurements. These measurements are made while linked to an 10/100/1000BaseT port-under-test. Connections are made using high grade RJ-45 cables and plugs so there is no need of any special fixturing or probing as is common with traditional physical layer test methods.

PhyView measurements and test capabilities can be classified into four different topics:

1. **Link Configurations & Assessments**: Basic analysis of port-under-test link capabilities and link status
2. **Transmission Measurements**: Direct measurements of the port-under-test transmitter function
3. **Interface Measurements**: Evaluation of electrical interface parameters
4. **Receiver Tests**: Assessments of port-under-test receiver performance characteristics

Each of these areas will be addressed in the following paragraphs.

1.2.1. Link Configurations and Assessments

Link configuration and assessment is the basic ability to determine link state and link configuration, and to control link configuration. Since link-up and link configuration is essential to all other PhyView Analyzer measurements, the instrument must have these capabilities. Link configuration features are presented in the following table.

Link Configuration	Options	Description
Link Rate	Auto, 10, 100, 1000, 10/100	Forces specified or at auto-negotiated rate
Connection Mode	Auto, MDI, MDI-X	Forces specified or link determined connection mode
Duplex	Auto, Half, Full	Forces specified or auto-negotiated duplex mode
Gigabit Mode	Auto, Master, Slave	Forces specified or auto-negotiated 1000BaseT gigabit timing mode

Link assessment includes two specific capabilities: **Link Monitor** and **Link Partner** assessment.

The **Link Monitor** reports current link status and link-up parameters such as link rate, duplex, connection mode, and gigabit mode (1000BaseT only). This can be sampled at any time. The Link Monitor can also be configured to perform automatic sampling and link status counts over a user-specified period of time. This resource is a key component to the Receiver Tests topic presented in Section 1.2.12 below.

The **Link Partner** measurement captures and reports the auto-negotiation and auto-MDI characteristics of a 10/100/1000BaseT port-under-test. This resource is important in qualify which automated tests can be run from the PhyView PHY Performance Test Suite that is the topic of Section **Error! Reference source not found.**.

1.2.2. Transmitter Verification Overview

The PhyView Analyzer will make transmitter measurements while linked using either 100BaseTx or 1000BaseT. While there are no direct transmitter measurements available for 10BaseT in the PhyView Analyzer, given the simplicity of a 10BaseT transmission, there are not many possible fault modes that would adversely affect 10BaseT and not affect 100BaseTx or 1000BaseT. The following PhyView Analyzer measurements assess transmitter performance:

- Power Spectral Distortion
- Signal-Noise Ratio
- Transmission Level
- Pair Skew

1.2.3. Power Spectral Distortion

Power Spectral Distortion (PSD) is a measurement of the spectral power characteristics of a 100/1000BaseT transmitter normalized to a nominal, undistorted 100BaseTx or 1000BaseT transmission spectrum. Figure 1.6 graphically depicts both 100BaseTx and 1000BaseT transmission spectra with and without packet flow. These spectra have been adjusted to align peak power levels to 0 dB so that relative power loss as a function of frequency is readily visible.

Looking at these spectra, it can readily be observed that the distribution of power over frequency is very much tilted to the lower frequencies. This is critical to how these technologies overcome the insertion loss of 100M cable channels. Secondly, it is clear that these spectra are largely not impacted by packet flow. This is critical to assuring that IDLE links remain linked with no loss of receiver synchronization.

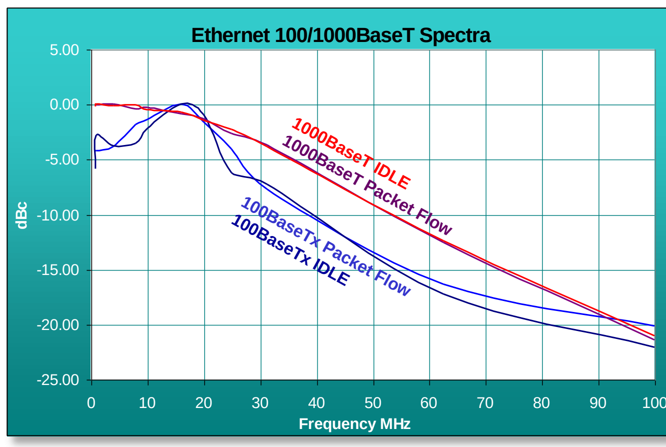

Figure 1.6 100BaseTx and 1000BaseT Transmission Spectra

The 100BaseTx spectrum shows a distinct peak in the 16 MHz range, then rolls off to –20 dB by 90 MHz. This characteristic peak relates to a statistical behavior of the MLT-3 encoding used in 100BaseT. MLT-3 is designed to reduce spectral bandwidth to below 30 MHz.

The PAM5 encoding of 1000BaseT would not naturally produce the band limiting behavior of 100BaseT. So 1000BaseT applies a digital averaging, or low pass filter to roll off the spectrum to roughly match that of 100BaseT. This protects against high RF emissions and crosstalk that would otherwise result given unshielded twisted pair transmission media. The pre-distortion filter can then be corrected by the 1000BaseT receiver that knows this exact filter function. High speed digital signal processing (DSP) is essential in 1000BaseT PHY's.

The PSD measurement recovers the spectral shape of a transmitted signal, then compares that shape to a nominal, or idealized spectral shape that depends on the link rate, 100BaseTx or 1000BaseT. This means that the ideal PSD measurement would be a flat 0dB from low frequency to high frequency. Looking at Figure 1.7, the NOMINAL PSD trace represents the ideal PSD measurement result.

Many impairments that distort a transmitter signal in the time domain will produce a different type of distortion in the frequency domain. Figure 1.7 diagrams a number of important distortion types and the expected PSD response to those impairments. These impairments include:

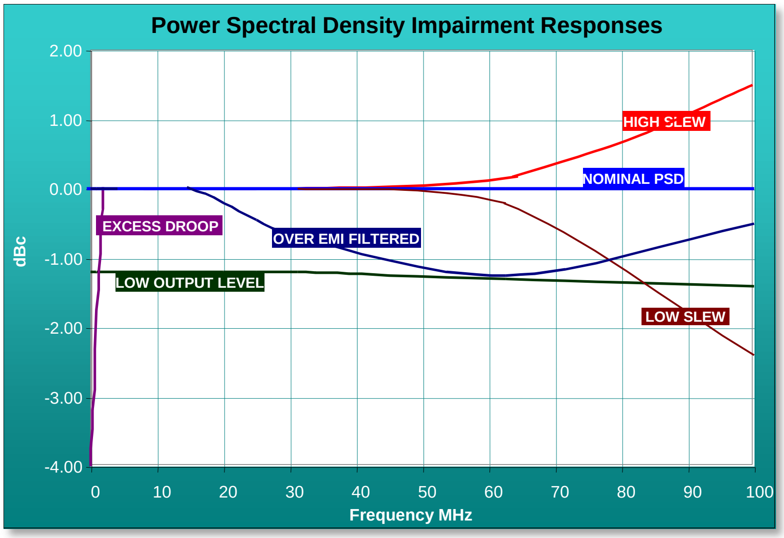

Figure 1.7 PSD Responses to Various Impairments

Low (or High) Output Level: A low output level as shown *Figure 1.8a* would equate to a LOW OUTPUT LEVEL trace in a PSD measurement (*see Figure 1.7*). PSD will be below 0 dB, especially at the lower frequencies where most of the spectral energy exists. Similarly, high output level would correspond to a higher PSD in the low frequency bands. Lower frequency PSD values are approximately related to voltage using the relationship 20*LOG(**Vp-p/Vp-p_nom**) where **Vp-p** is actual peak-peak voltage measurement and **Vp-p_nom** is the ideal peak-peak voltage, e.g. 2Vpp for 100BaseTx. Assuming for the moment that PSD at 2 MHz is linearly related to Vp-p, a 2 MHz 1000BaseT PSD that measures –1.4 dB at 2 MHz would equate to 2Vpp * $10^{(-1.4/20)}$ = 1.7Vpp. (*Note: The 802.3 standard would expect a minimum of 1.79Vpp for 1000BaseT*.)

Low (or High) Slew Rate: Slew rate, or Rise/Fall time, tends to affect 100BaseTx and 1000BaseT signals in the frequency range above 65 MHz with most of the impact observable above 75 MHz. This means a nominal power signal with low slew rate (*see Figure 1.8c*) will approximate 0 dB until the upper frequencies where it will decline. Both a LOW SLEW rate and a HIGH SLEW rate are depicted in Figure 1.7. Relating the low slew rate to the Ethernet transmission spectra in Figure 1.6, the higher frequencies would roll off even faster. Low slew rate can be especially harmful in a 100M link where attenuation of higher frequencies is already a challenge to the link integrity. Conversely, high slew rate will be detrimental in very short links and will lead to RF emissions problems.

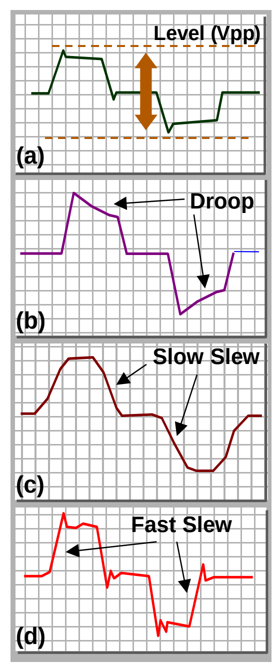

Figure 1.8 Impairments in the Time Domain

Excess Droop: Droop in any twisted pair baseband transmission (*see Figure 1.8b*) typically comes about because a transformer is removing very low frequencies from the transmitted signal. Transformers are band limited devices that work over a min-max frequency range. Because PSD can assess frequency response as low as 20KHz, it is able to discern EXCESS DROOP (*see Figure 1.7*) that might typically affect frequencies below 100KHz. In 100BaseT, droop can be directly compared to baseline wander, that is, the inability of the transmitted signal to hold a steady voltage for a longer period of time given an all zero's encoded data pattern. For example, 100KHz would correspond to 10µsec, 1,250 symbol periods, or a data pattern of 125 bytes. Power-over-Ethernet can worsen droop if the delivery of DC current is not evenly split across both conductors of a twisted pair. This impairment, referred to as DC Unbalance, causes saturation in the transformer which in turns clips off low frequency response.

Over Filtering: Most 10/100/1000BaseT interfaces use transformer magnetics not only for isolation but also for EMI and common mode suppression. This requires a low-pass filtering effect that should nominally roll off above 100MHz. If this filtering cuts off at lower frequencies like 30 – 50 MHz, then PSD will naturally pick this effect up as a spectral distortion. This is shown in Figure 1.7 as OVER EMI FILTERED.

The PHY Performance Test Suite (see Section **Error! Reference source not found.**) takes advantage of relationships between PSD measurements and certain time-domain measurements in order to predict parameters such as **Vpp** (100/1000BaseT), **Rise/Fall Time** (100BaseT), and **Template Fit** (1000BaseT). Those relationships start with theoretical underpinnings including the topics of these paragraphs, and then add in a body of empirical correlation experience in order to tighten the predictive accuracy of those critical 802.3 conformance parameters.

PSD is a calibrated measurement so that effects of cabling, connectors, and even test receivers are properly compensated. Fully automated calibrations requiring no external calibration standards simplify the task of periodic calibration. PSD can be readily used to measure insertion loss across link components including cabling, connectors, and patch panels. PSD is also useful for assessing signal integrity at any service point outlet.

1.2.4. Signal-to-Noise Ratio

Signal-Noise Ratio (SNR) is a measure of residual distortion on a transmitted signal. Residual distortion refers to non-correctable distortions that occur either because of imperfections in a transmitter or because of problems in a link. Examples of non-correctable distortion include **amplitude noise** (or alien crosstalk), **phase noise** (or severe jitter), **clipping**, **signal compression** (or non-linear gain/loss), and severe inter-symbol interference (**ISI**) that would occur given non-linear or notch-like filtering.

Because residual distortion is not correctable, it will *for certain* adversely affect bit error rate and packet flow in a receiver. SNR is reported in units of dB, that is, the *power* of the signal in ratio to the *power* of non-ideal or distortion components. 36 dB SNR would mean that the ratio between true signal power and distortion components is $10^{(36/10)}$ = 3981. (*Note*: SNR = 10*LOG[$\mathbf{P}_{\mathbf{SIGNAL}}$ / $\mathbf{P}_{\mathbf{NOISE}}$]).

Modern, DSP-based twisted pair PHY receivers are designed to correct for predictable distortions including the spectral shaping done in 1000BaseT and the effects of insertion loss and linear phase shifting that are characteristic of transmission media, namely cabling. The PSD measurement discussed earlier in Section 1.2.3 is designed to be sensitive these types of impairments or distortions and will be affected by both "correctable" and certain residual distortions. To the extent that a receiver successfully compensates the *predictable* distortions, the bit error rate, and ultimately packet flow, is unaffected by the presence of those distortions.

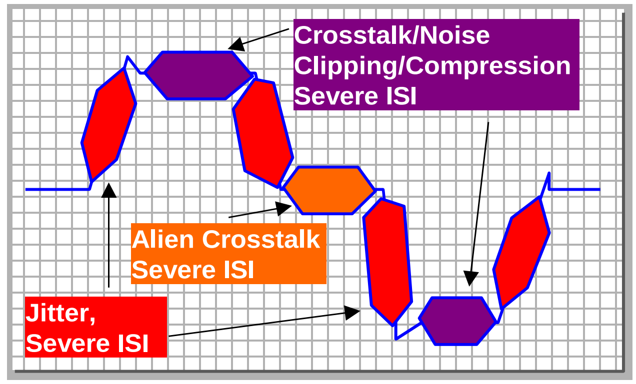

Figure 1.9 Sources of Residual Distortion

One special case of residual impairment is jitter. While jitter is random, it is to a certain extent, correctable. Receivers are designed to acquire signal timing from the received signal and thus have some ability to track out, or nullify, randomly varying signal timing. That ability is very much receiver implementation dependent and will have limitations both in jitter magnitude and jitter rate. Lower magnitude and lower rate jitter is easily tracked out while high rate (or high frequency) jitter may easily flow through the receiver and turn up as residual noise. This is why the 1000BaseT standard places a 5KHz bandwidth on the specification of peak-peak jitter.

In 1000BaseT, receivers have to deal with a special form of quasi-random noise in the form of echoed transmitted signals and local crosstalk from adjacent transmitting pairs. 1000BaseT receivers are designed to use matched filters to remove these quasi-random interferers, however, to the extent that those filters are not perfect or that the hybrid echo cancellation is not perfect, those interfering signals will show up as residual noise and affect the bit error rate. While the PhyView Analyzer as a measuring instrument compensates for this internal error, this effect also places an upper bound on the maximum SNR (36 dB) that the meter can report.

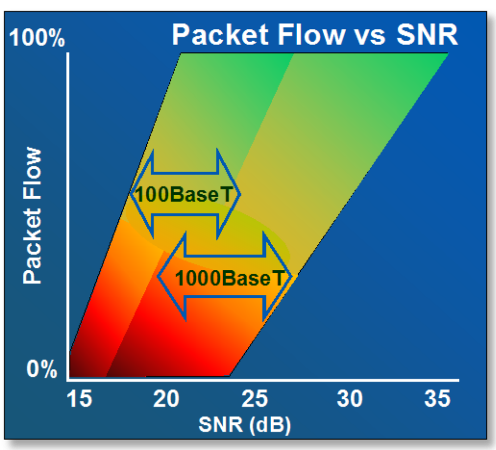

Figure 1.10 SNR vs Packet Flow

The relationship between SNR, Bit Error Rate, and Packet Flow will generally be design implementation dependent. Figure 1.10 diagrams this using numbers that are realistic but show the spread of potential performance. From this diagram, it is evident that 100BaseTx is more tolerant of lower SNR than 1000BaseT in general. It is also evident that SNR in the 30's will seldom result in any packet loss at any receiver while SNR in the low 20's is a problem in many implementations for 1000BaseT and in the mid-teens will likely be a problem for 100BaseTx on most implementations.

Finally, it should be noted that low transmit signal amplitude will adversely impact measured SNR. This relationship is essentially in the definition of SNR since lower signal power means that the ratio of signal power to ambient noise power *must* decline. This fact should be considered when creating impaired receiver tests as will be discussed elsewhere in this manual. It also means that SNR measurements can be artificially reduced simply by adding fixed loss in line with a transmitter-under-test.

1.2.5. Transmit Level Measurement

Transmit Level (or transmit power) is a bulk measurement of the *power* of a transmitted signal. This measurement is a factory calibrated and normalized such that a 0 dB reading for 100BaseTx corresponds to a nominal, 2Vpp, 4 nsec Rise/Fall time transmitted signal and a 0 dB reading for 1000BaseT corresponds to a nominal peak-to-peak voltage of 1.5V from Point A to Point B on the gigabit Test Signal #1 (*see Figure 1.11*).

Transmit Level (also referred to as **Tx Level**) should be thought of as the integrated power within the full signal spectra as compared with that of a nominally powered signal. It bears an approximate relationship with the peak-peak level of a transmitted signal, though as stated in Section 1.2.3, the peak-peak voltage will relate more closely to lower frequency power (e.g. 2 MHz PSD) than to the wide-band power measurement made with this meter.

The Tx Level measurement is used inside of PSD calibrations to assure a more accurate zero dB baseline to PSD measurements. When evaluating Ethernet ports, Transmit Level is a quick indication of total signal power relative to a nominal 100BaseTx or 1000BaseT signal. It will respond dB-for-dB if a fixed, flat frequency power loss (or pad) is placed in line with an Ethernet transmitter.

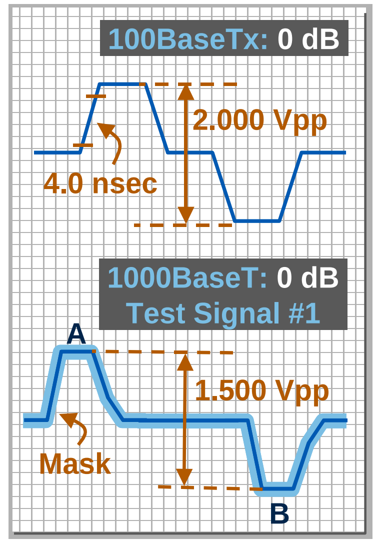

Figure 1.11 Zero dB Transmit Level

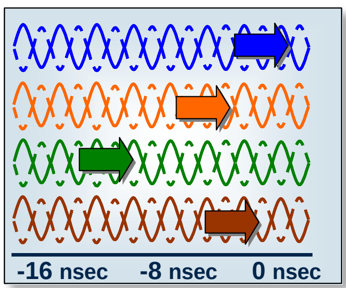

Figure 1.12 Pair Timing Skew

1.2.6. Pair Skew Measurement

The Pair Timing Skew measurement is only applicable in 1000BaseT links. This meter very simply reports the timing differences between the 4 incoming pairs (*see Figure 1.12*). Measurements are reported in nanoseconds and are normalized to one of the four pairs meaning one pair will always read zero nsec. Meter granularity is 8 nsec, or one symbol period.

The IEEE standard requires that delays between the four pairs should never exceed 50 nsec. This equates to just over 6 symbol periods. This requirement exists to assure receivers can properly reassemble data streams split amongst the four pairs.

1.2.7. Interface Verification – Overview

The PhyView Analyzer will verify two types of interface characteristics while linked at 1000BaseT.

- Bulk Return Loss (or Echo)
- Bulk Pair-Pair Crosstalk

These characteristics are separated from Transmitter and Receiver testing since they are not strictly related to the PHY transmitter or receiver. The measurements made required that the port-under-test support gigabit Ethernet. The values obtained however are not likely related to link rate and are generally applicable at any of the link rates 10, 100, or 1000BaseT. Of course, 10BaseT is much less sensitive to return loss or crosstalk because of it's relative narrow frequency band.

1.2.8. Bulk Return Loss (Echo) Measurements

Bulk Return Loss (or Echo Response) is a measure of the power ratio between transmitted power and incident reflected power. Transmitted power is reflected whenever there are discontinuities or mismatches in the transmission path of a signal. 10/100/1000BaseT interfaces are designed for 100Ω transmission paths meaning that source and termination impedances, as well as characteristic impedance, are nominally 100Ω over the applicable frequency range.

The term "bulk" indicates that the measurement is not reported as a function of frequency. It is simply a wideband bulk power ratio. It is available only in 1000BaseT because that technology was designed to enable signal power measurements that are correlated only to a specific transmission direction on a specific pair. In other words, the Bulk Return Loss measurement takes advantage of features that are unique to 1000BaseT.

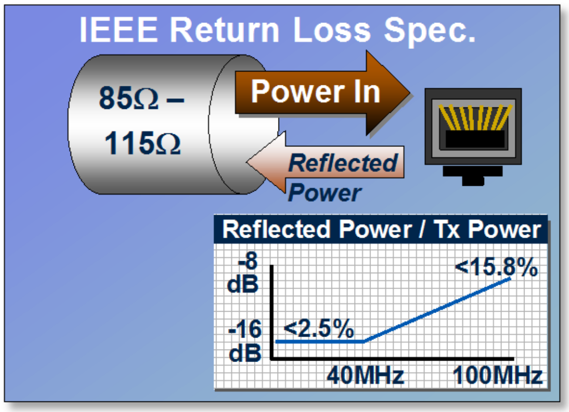

Figure 1.13 IEEE 802.3 Return Loss Specification

The IEEE 802.3 specification for Return Loss was written with the constraint that Category 5 cabling may have a characteristic impedance between 85 Ω and 115Ω (*see Figure 1.13*). Normally, a return loss measurement would be specified assuming a single, fixed channel impedance such as 50Ω. However in 802.3, it is essentially specified with two reference impedances, 85Ω and 115Ω. It is then expressed as a formula that changes over frequency. The limit line formula from 802.3 clause 40.8.3.1 is:

Return Loss (≤40MHz) ≤ - 16 dB
Return Loss (>40MHz) ≤ - (10–20*LOG(MHz / 80)) dB

To make this specification relevant to a test apparatus that has fixed 100Ω impedance, some transformations are necessary. The following table shows what the port-under-test must look like to a 100Ω source in order to meet the return loss requirements of -16dB (40MHz) and –8dB (100MHz). It also computes a figure for –15.5 dB since, given the known 1000BaseT signal spectra (*see Figure 1.6*), -15.5 dB approximately represents the wideband Echo power of a hypothetical port-under-test that tracks the 802.3 limit line function over frequencies from 1 to 100MHz.

Source Impedance	Return Loss	Port Under Test Impedance	Observation
85 Ω	-16 dB	61.8 Ω to **117** Ω	83.5 Ω to 117 Ω will produce ≤ -16 dB given a reference anywhere between 85 Ω and 115 Ω
	-8 dB	36.6 Ω to **197** Ω	
	-15.5 dB	60.6 Ω to **119** Ω	49.5 Ω to 197 Ω will produce ≤ -8 dB given a reference anywhere between 85 Ω and 115 Ω
115 Ω	-16 dB	**83.5** Ω to 158 Ω	
	-8 dB	**49.5** Ω to 267 Ω	82.1 Ω to 119 Ω will produce ≤ -15.5 dB given a reference anywhere between 85 Ω and 115 Ω
	-15.5 dB	**82.1** Ω to 161 Ω	
PVA Impedance	**Port Under Test Impedance**	**Expected Return Loss**	**Applicable Frequencies**
100 Ω Frequency Selective	83.5 Ω	-20.9 dB	Frequency Range 0 – 40 MHz
	117 Ω	-22.1 dB	
	49.5 Ω	-9.4 dB	@ 100MHz
	197 Ω	-9.7 dB	
100 Ω Wideband	82.1 Ω	-20.1 dB	Bulk Return Loss Limits that assure MDI Return Loss specs across 85 Ω to 115 Ω
	119 Ω	-21.2 dB	

Using this table, the *wideband criteria* for a compliant interface would be in the –20 to –21 dB Return Loss region. The Return Loss metrology is restricted to a floor of –26 dB and an error magnitude of 1 to 1.5 dB. So a fair criteria to apply to a typical port-under-test would be in the –19 to –20 dB range.

Bulk Return Loss is a calibrated measurement so that effects of cabling, connectors, and even test receivers are properly compensated. Fully automated calibrations that require no external calibration standards simplify the task of periodic calibration. Bulk Return Loss can be readily used to measure return loss across link components including cabling, connectors, and patch panels.

1.2.9. Bulk Pair-Pair Crosstalk

Crosstalk is a ratio between the power transmitted on a single cable pair and the portion of that transmitted power that leaks across into other cable pairs. It is generally thought of as a property of multi-pair cabling and connector components. Typically, crosstalk increases with frequency owing to capacitive coupling between components. IEEE 802.3 does not explicitly specify crosstalk performance of a 100/1000BaseT PHY, however it does specify the use of TIA/EIA 568 (ISO/IEC 11801) qualified connectors that are in turn subject to various crosstalk specifications. IEEE 802.3 also specifies expected near-end crosstalk (*see Figure 1.14*) performance of physical links that connect Ethernet ports with the formula:

NEXT ≤ 27.1 – 16.8*LOG(MHz / 100) dB

Ethernet ports can contribute to crosstalk in several ways. Defective components and circuit trace layout issues may lead to excessive crosstalk either as a design feature or on an occasional port basis. Many Ethernet ports utilize magnetics packages that combine all 4 pairs into a single compact package with tight spacing between coils. Mechanical defects in sockets may also add to crosstalk.

The PhyView Analyzer offers a calibrated Crosstalk measurement that assesses near-end crosstalk at the port-under-test interface. The measurement reports crosstalk by pair combination, that is, crosstalk between Pair #1 and Pair #2 or between Pair #2 and Pair #4. (*See section 1.2.10 for discussion of pair numbering convention*.)

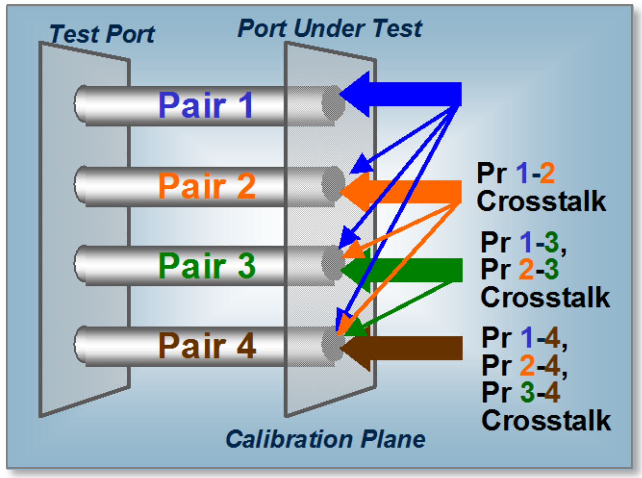

Figure 1.14 PVA Crosstalk Measurement

There are six unique pair combinations (*see Figure 1.14*):

- Pair 1 to 2
- Pair 1 to 3
- Pair 1 to 4
- Pair 2 to 3
- Pair 2 to 4
- Pair 3 to 4

The PVA Crosstalk measurement assumes that all crosstalk is bi-directional, that is, that crosstalk is really a measurement of electrical isolation between two pairs. Well isolated pairs will report low crosstalk and poorly isolated pairs will report higher crosstalk.

Much like the Return Loss measurement, Crosstalk is a wideband or bulk measurement across all frequencies, can only be performed while linked at 1000BaseT, and is reported in dB. It is available only in 1000BaseT because that technology was designed to enable signal power measurements that are correlated only to a specific transmission direction on a specific pair.

Since there are no explicit limits from IEEE 802.3, hypothetical performance limits need to be discerned from other specifications. In a 10/100BaseT channel, the only crosstalk that matters is Pair 2-3 since pairs 1 and 4 are not utilized for data communication. In 1000BaseT, all four pairs are utilized in a full duplex fashion and any crosstalk that can be measured on any pair will create an uncorrectable distortion at the receiving end of the link. If the link between two ports were ideal, then the full budget of the NEXT formula limit line (*see abo*ve) could be applied to the port-under-test. Crosstalk would then range from –60dB at 1 MHz up to –27dB at 100MHz.

Given the spectral characteristics of a 1000BaseT transmission, the wideband equivalent to the NEXT formula limit line approximates –37.8 dB. That is, the allowable crosstalk between any two pairs is about 0.017% of transmitted power. The Crosstalk metrology is restricted to a floor of –39 dB and an error magnitude of 1 to 1.5 dB. This means that the Crosstalk measurement is particularly useful for exposing design or manufacturing defects that cause performance to degrade by worse than 3 to 6 dB relative to what the 802.3 standard hypothetically allows.

1.2.10. Pair Designators in the PhyView Analyzer

The PhyView Analyzer adheres to a pair numbering convention used in TIA/EIA-568-B. This convention is commonly recognized throughout the cabling and connector industry. In this convention, pair numbers are related to colors and RJ-45 connector pin numbers as follows:

Pair Number	Pair Colors	RJ-45 Pins	Description
1	Blue & Blue/White	4 & 5	Spare Pair for 10/100BaseT
2	Orange & Orange/White	2 & 1	MDI Transmit / MDI-X Receive Pair for 10/100BaseT
3	Green & Green/White	6 & 3	MDI-X Transmit / MDI Receive Pair for 10/100BaseT
4	Brown & Brown/White	8 & 7	Spare Pair for 10/100BaseT

1.2.11. Receiver Verification Overview

Unlike transmitter testing, 10/100/1000BaseT receiver testing is not generally addressed in 802.3 standards. 802.3 standards do define the worst-case link (or cabling systems) environment for Ethernet ports with specifications for channel insertion loss, return loss, and crosstalk. These specifications are ultimately derived from TIA/EIA 568 specifications related to Category 5/5e cabling systems. Transmitted signal characteristics such as signal amplitude range, clocking frequency, distortion parameters, and timing jitter are also thoroughly specified in the standards and by inference create additional constraints on required receiver performance. Finally, allowances for added noise, or alien crosstalk also affect receiver requirements. This collection of receiver requirements is depicted in *Figure 1.15*.

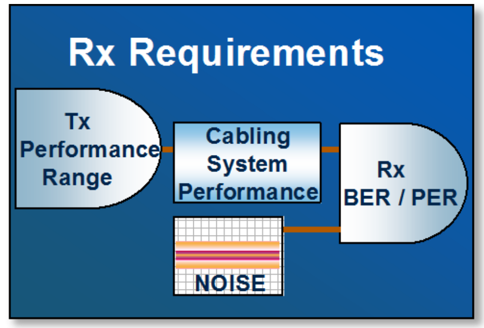

Figure 1.15 Ethernet Port Rx Requirements

The 10BaseT and 1000BaseT specifications also define explicit goals for bit error rate (BER) performance. The 100BaseTx specification, and its subsidiary TP-PMD specification make no mention of bit error rate targets. Section 1.1.5 above raised the additional issue that BER is not directly measurable and must be inferred from packet flow measurements.

The table below summarizes specifications that are relevant to 10/100/1000BaseT receiver testing.

Link Rate	Target BER	Cabling System	Rx Noise Insertion	Target Packet Error Rate
1000BaseT	10^{-10}	Cat5e -24dB @100MHz	25mVp-p 100BaseTx Spectrum	10^{-7} (125 octet frames, *no account for coding ovhd.*)
100BaseTx	Unspecified (UNH: 10^{-8})	Cat5 -24dB @100MHz	40mVpp Band Limited 100MHz	Unspecified (UNH: $1.5*10^{-5}$, 64 octet frames)
10BaseT	10^{-8}	Cat3 -11.5dB @10MHz	300mVpp Band Limited 15MHz	Unspecified ($5.12*10^{-6}$, 64 octet frames)

Summarizing, receiver test methodology is largely the responsibility of Ethernet port designers and implementers. The task of modeling all of the possible signal and channel impairments while assessing bit error rate performance can be daunting and the exact criteria for success (pass or fail) can be ambiguous.

1.2.12. Receiver Testing Metrology

The PhyView Analyzer offers two distinct measurement techniques for receiver testing:

1. **Link Monitor**: Receiver Measurements on stand-alone Ethernet ports
5. **Packet Flow**: Receiver Measurements on multi-port (bridging) devices such as switches, hubs, and repeaters

The Link Monitor was introduced in Section 1.2.1. This measurement resource assesses both instantaneous and time-sampled **Link Status**, as well as other link configuration parameters. In receiver testing, the key parameter is sampled Link Status. Link Status can be one of three physical layer indicators:

1. **Link State**, that is link **UP** (linked) or link **DOWN** (unlinked) – available for 10/100/1000BaseT
6. **Remote Rx Status**, a physical layer indicator from the link partner indicating Rx “OK” or Rx “Not OK” – *available for 1000BaseT only*.
7. **Local Rx Status**, a physical layer indicator from the test port receiver indicating Rx “OK” or Rx “Not OK” – *available for 1000BaseT only*.

Link Status can be a single, instantaneous sample or can be configured to report a count of up to “UP” or “OK” samples evenly spaced with sampling intervals of 20, 50, or 100msec, whereupon it becomes a **Link Stability** measurement.

Generally, the criteria for link “UP” in 10BaseT is the continuous receipt of link test pulses (*see Figure 1.16*). A highly impaired 10BaseT receiver-under-test may or may not choose to continue transmission of link pulses meaning there is no sure way for one link partner (or Test Port) to discern the receive condition of the other link partner. This means that **Link Stability** generally will *not* correlate linearly to a hypothetical port-under-test bit error rate or packet error rate.

When testing 100BaseTx receivers, the ability of the **Link State** to predict or estimate effective packet loss in a port-under-test receiver is very much implementation dependent since the 100BaseTx standard is nebulous about the action of receivers that are experiencing any degree of packet loss. Some PHY implementations may recognize error conditions and attempt to force a re-link to overcome it. Any re-link can be detected as a link drop by the Link Monitor.

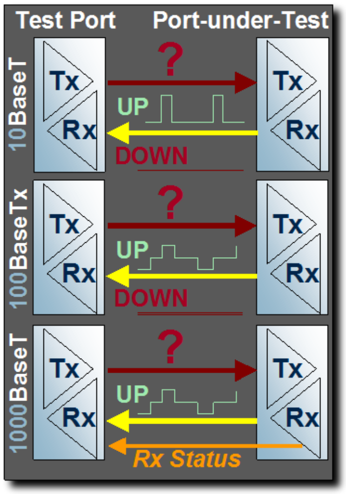

Figure 1.16 Link Status Monitor

In re-linking, some 100BaseTx PHY implementations may naturally drop back to 10BaseT in an attempt to overcome serious channel impairments. The loss thresholds at which any given PHY implementation makes decisions to drop link and re-link is unspecified. Conversely, some PHY implementations, when auto-negotiated to 100BaseT, will attempt to maintain 100BaseT link by transmitting continuous IDLE MLT-3 signaling (*see Figure 1*.16) even when *incoming* packet traffic and IDLE signaling is 100% impacted by bit errors. Summarizing, for 100BaseTx, the relationship between Link Stability and packet loss can range from very weak to very strong depending the PHY implementation in the port-under-test.

1000BaseT introduced the concept of a receive status indicator combined with a purely physical layer signaling protocol to communicate receiver status to the link partner while linked (*see Figure 1.16*). Because of DSP processes including forward error correction that are inherent in receiving 1000BaseT signals, PHY implementations could more readily discern and communicate bit error problems even when link performance was reasonably good. The Link Monitor takes advantage of this feature in 1000BaseT by allowing **Link State** ("UP" or "DOWN") to be replaced by **Remote Rx Status** "OK" or "Not OK". **Remote Rx Status** can be sampled periodically from 1 to 100 samples at periodicity of 20, 50, or 100msec. **Remote Rx Status** is a reasonably good predictor of bit error and packet loss performance.

The second option for receiver measurements in the PhyView analyzer is a simple **Packet Flow** measurement. This involves the transmission of user-specified packets (size, gap, payload) into the port-under-test. The packets generated by each PhyView Analyzer test port are restricted to MAC frames with programmable address, size, packet gap, and repeating 4-byte payload. This technique therefore requires that the port-under-test must forward any successfully received packets to another Ethernet port so that they may be returned to another PhyView Analyzer test port. Hence, the device being tested must be a multi-port device with layer 2 bridging capability such as an Ethernet switch, hub, or repeater.

The obvious advantage of packet flow testing is that it will always provide insight to the bit error performance of a receiver regardless of link rate, that is 10BaseT, 100BaseTx, or 1000BaseT.

A disadvantage of Packet Flow is that it cannot be used on single-port devices or devices that don't bridge layer 2 packet traffic. Another disadvantage is that it consumes two PhyView Analyzer test ports and is generally a bit more complicated to configure and initialize than the Link Monitor. Often, layer 2 bridging must be "primed" in an Ethernet switch as it executes spanning tree protocols in order to isolate connection paths. Finally, the PhyView Analyzer will not filter received packets so that any rogue, unsolicited packets (e.g. LLDP or CDP) will add to received packet counts, though usually at very low levels.

1.2.13. Receiver Testing with Impairments

A major contribution of the PhyView Analyzer toward 10/100/1000BaseT receiver testing is the ability to apply a range of controlled and relevant impairments. Impairments are vital for moving the BER performance of any receiver-under-test to a point where it can be measured much more quickly and can also be statistically compared to the BER performance of other receivers.

Figure 1.17 diagrams this phenomenon and provides a quantitative perspective. For example, if the BER performance of an unimpaired receiver is 10^{-11}, that is one error in 10^{11} bits, testing at 100BaseTx would yield a test time of over 13 minutes to discover the first packet error even when testing with the maximum packet size of 1518 octets. In 1000BaseT, this time is reduced by a factor of 10 and becomes just under 1½ minutes. Following up on the relationships between packet errors and bit error rate (*see Section 1.1.5*), statistically, it takes many packet errors to assess BER.

Impairing the receiver-under-test will drive up the BER, increase packet drops and packet error, and assess that receiver under typical and worst case stresses.

The PhyView Analyzer offers a powerful combination of impairments that place realistic stresses on 10/100/1000BaseT receivers while measurements such as Link Monitor or Packet Flow are being performed. Many of these impairments can be applied asymmetrically so that the port-under-test experiences impairment that the PhyView Analyzer test port does not, thus assuring that link instability is fully attributable to the port-under-test.

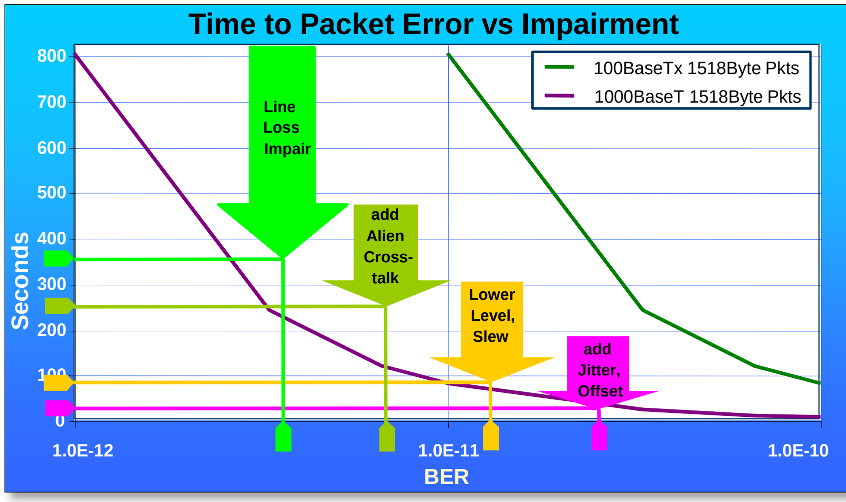

Figure 1.17 Packet Error Yield with Impairments

Physical impairments offered within each PVA test port include:

- **Line Loss**: Simulated worst case IEEE 802.3 cabling system insertion loss
- **Mismatch**: Fixed –12 dB Return Loss
- **Alien Crosstalk**: Frequency-contoured ingress noise with programmable level
- **Tx Clock Offset**: Offset transmit signal timing with 6 programmable settings
- **Tx Clock Jitter**: Phase-noise vs frequency contoured transmit jitter with programmable level
- **Tx Level**: Transmit signal amplitude with programmable level
- **Tx Slew**: Transmit signal slew rate with programmable rise time

These impairments are each discussed with greater detail in Section **Error! Reference source not found.**.

Most of the impairments are independent in that they may be combined. For example, an outgoing 100BaseT transmit pair may have the Line Loss impairment coupled with Alien Crosstalk, Tx Clock Jitter, low Tx Level, and low Tx Slew. Combinations of impairments such as these are very effective at driving bit error rate levels into regions where they can be efficiently characterized and accurate comparisons between different ports can be made in a short period of time.

The fully automated PHY Performance Test Suite addressed in Section **Error! Reference source not found.** has a set of receiver tests that take full advantage of the impairments and receiver testing metrology to perform comprehensive analysis of 10/100/1000BaseT receivers in a several minutes and without any user intervention.

1.2.14. Reference Manual Overview

The remainder of this manual will familiarize users with the various resources available within each test port and the various means of accessing and using those resources for link configurations, impairment configurations, measurements, calibrations, and fully automated test suites.

Section Error! Reference source not found. will introduce basic test port architecture, describe each testing resource in detail, discuss PVA software components, and present technical specifications for the PhyView Analyzer.

Section Error! Reference source not found. will present the PVA Interactive graphical user interface application available for intuitive and interactive control of the PhyView Analyzer. It will include some sample exercises to help familiarize users with common tasks.

Section Error! Reference source not found. will review PowerShell PSA, the live command line control and script automation environment for the PhyView Analyzer and for all PowerSync Analyzer instruments from Sifos Technologies.

Section Error! Reference source not found. will cover the fully automated PHY Performance Test Suite for 10/100/1000BaseT Ethernet port testing. This optionally configured test suite makes multi-port PHY testing and analysis a one-button-click task with colorful, graphical reporting including test limits for excellent versus marginal versus non-compliant parameter performance. Transmitter, interface, and receiver tests are included in this test suite.

Finally, **Section** Error! Reference source not found. will provide some additional information for users that must automate testing with the PhyView Analyzer and need tips for integrating with other software environments and/or making tests more efficient.