Transceiver Solutions

# Gigabit Ethernet Over Copper Performance

by Jason Knickerbocker

June 11, 2003

MARVELL® WHITE PAPER

MARVELL®

MOVING FORWARD FASTER®

**Advanced Information**

This document provides advanced information about the product described. All specifications described herein are based on design goals only. Do not use for final design. Visit the Marvell® web site at www.marvell.com or call 1-866-674-7253 for the latest information on Marvell products.

## Table Of Contents

## 1.0 Executive Summary

Gigabit Ethernet is rapidly becoming the network of choice from the backbone to the desktop. Like its predecessor, Fast Ethernet, Gigabit Ethernet is enjoying quick market adoption due to its ability to offer a significant increase in network bandwidth—ten times more bandwidth—while maintaining backward compatibility to the installed base of Ethernet networking systems. The IEEE 802.3ab Gigabit Ethernet standard represents truly extraordinary technology in that the standard defines 1000 Mbps data transmission over the same copper media defined for Fast Ethernet at 100 Mbps—Category 5 Unshielded Twisted Pair (UTP) cable. This paper reviews some of the challenges in running 1000 Mbps data over Category 5 cable and how the IEEE standard proposed to address them. It is important to note, however, that the actual practice of running 1000 Mbps of data over existing Category 5 cabling contains many challenges, where if not properly addressed can potentially cause significant performance issues, resulting in unreliable networks and costly downtime. At the heart of these issues—and the solution to a robust and reliable Gigabit Ethernet system—is the Physical Layer (PHY) device of Ethernet.

In this paper, some of the theoretical (as addressed by the standard) and practical (real world) challenges of running Gigabit Ethernet over Category 5 cable are explored and the impact of the PHY transceiver examined. It will be shown that the performance challenges and pitfalls of running ten times more data over existing Fast Ethernet cable lines, may be alleviated through the use of robust PHY transceiver technology—devices that greatly exceed the requirements of the IEEE Gigabit Ethernet standard. Finally, it will be shown that the technology that achieves these lofty goals is available through the Marvell Alaska® family of Gigabit Ethernet PHY transceivers.

Sections 2.0 through 4.0 offer the reader an overview of the Gigabit Ethernet market and the technology defined for its implementation, as detailed in the IEEE 802.3ab standard. The reader versed in Gigabit Ethernet may choose to start reading at Section 5.0.

## 2.0 Ethernet Market Dynamics

Over two decades old, Ethernet has become the most widely deployed and well-known networking technology in the world. Throughout its existence, Ethernet has evolved from supporting simple workgroup environments to serving as the primary communications medium for entire corporate campuses. Additionally, through Ethernet's evolution, the line speed (or data rate) has surged from 10 to 100 Mbps, and now to 1000 Mbps. Put into perspective, this represents a 100-fold increase in data bandwidth in less than 10 years.

The advantage of Ethernet's rapid acceleration in available network bandwidth is underscored by the fact that Ethernet technology also enables seamless and easy migration to higher data rates—incrementally, as the network requires it. Each successive higher Ethernet data rate is backward compatible to existing, lower speed, Ethernet. For example, a Gigabit Ethernet Network Interface Card (NIC) will automatically configure to operate at 1000 Mbps if it is connected to a Gigabit Ethernet Switch; however, in the event it is connected to a Fast Ethernet switch, the Gigabit NIC will automatically configure to 100 Mbps. Thus, using Ethernet technology, companies do not have to undergo an expensive, time-consuming infrastructure upgrade. They can migrate portions of the network to higher bandwidth as the network requires it and the budget allows for it.

Fast Ethernet systems currently dominate the network industry as measured by both shipment rate as well as installed base. In this new millennium, with the continued dramatic increases in desktop computing power driven by Intranet-based operations and the increasing demands for time-sensitive delivery between networked users, the network industry is currently driving towards another transition to higher bandwidth—ten times the performance of Fast Ethernet—or 1000 Mbps Gigabit Ethernet.

The rapid conversion of the LAN from 10 Mbps 10BASE-T Ethernet to 100 Mbps 100BASE-TX Fast Ethernet heavily contributed to the mass deployment of Category 5 copper cable, which is now dominant within today's networked buildings. One of the key contributing factors that is driving the successful high volume deployment of Gigabit Ethernet is the fact that Gigabit Ethernet was defined to operate over the existing installed base of Category 5 cable.

## 3.0 Tutorial: How Is Gigabit Ethernet Transmitted Over Category 5 Cable?

To understand how Gigabit Ethernet is transmitted over Category 5 cable, it is necessary to first take a brief look at some of the key characteristics of the Fast Ethernet standard. The Fast Ethernet standard specifies MLT-3 line signaling, which defines a 3-level transmission scheme running at a clock rate of 125 Mbaud. Fast Ethernet, which transmits/receives 100 million data bits per second, uses a 125 MHz clock as the standard also specifies 4B/5B (4 bits to 5 bits) encoding/decoding, adding 25% data overhead (4B/5B encoding/decoding guarantees minimum bit transitions, which is required for the Fast Ethernet over fiber standard, or 100BASE-FX). Another key consideration is the fact that the 100BASE-TX standard defines two Category 5 twisted-pairs, one for transmit and one for receive.

Similar to Fast Ethernet, the Gigabit Ethernet standard defines a 125 MHz clock resulting in many advantages. First, using the same clock rate simplified the design of dual data rate 100/1000 Mbps transceivers. An additional advantage is that the frequency spectrum of Fast and Gigabit Ethernet are similar—both with a null at 125 MHz. The 1000BASE-T standard achieves 1000 Mbps of data transmission by using four Category 5 twisted-pairs (as opposed to two pairs used in Fast Ethernet) as well as a more sophisticated five-level coding scheme called PAM-5. Another important difference relative to Fast Ethernet is that 1000BASE-T sends and receives data simultaneously on each of the four pairs.

Of the five levels provided in PAM-5 coding (-2, -1, 0, + 1, +2), four levels are used for data transmission and represents 2 bits of information. The fifth level is used for "Forward Error Correction" and will be discussed later in this paper. The symbol rate of PAM-5 coding is only half that of 2-level signaling, and thus a 250 Mbps data signal can be transmitted at a rate of 125 Msymbols/sec. Therefore, 5-level PAM coding provides better bandwidth utilization than binary signaling, in which each transmitted symbol represents just one bit—0 or 1.

## 4.0 Tutorial: Theoretical Challenges of Running Gigabit Ethernet Over Category 5 Cable

As the Gigabit Ethernet standard specifies a ten-fold increase in data bandwidth while maintaining the same cabling infrastructure deployed for Fast Ethernet, clearly it stretches the limits of technology. This section reviews some of the issues in transmitting 1000 Mbps of data over Category 5 and how the standard addresses them.

### 4.1 Signal-To-Noise Ratio (SNR) Margin

SNR margin is a measure of a communications system's immunity to noise. SNR margin is expressed in dB and represents the level of additional noise that the system can tolerate before violating the required Bit Error Rate (BER). For example, an SNR margin of 3 dB means that if the noise level were increased by 3 dB, the system would be subject to excessive errors. Simply stated, SNR margin is a measure of how much additional noise a system can tolerate or how far the system is from not working properly.

Figure 1 below demonstrates that increasing the number of signal levels while maintaining the same transmit voltage results in a degradation of the SNR margin. The reason for this is that as the vertical opening of the eye gets smaller, the system can tolerate less noise before bit errors begin to occur. For example, increasing the number of logic levels from 3 (MLT-3 used in Fast Ethernet) to 5 (PAM-5 used in Gigabit Ethernet), while keeping the overall voltage constant at 2 Volts, reduces the vertical eye opening of each logic state by 50%. The noise voltage required to cause a symbol error on a 3-level signal is half (or 6 dB lower) than the voltage required to cause a symbol error on a binary signal. So a 3-level signal has 6 dB less SNR margin than a binary signal.

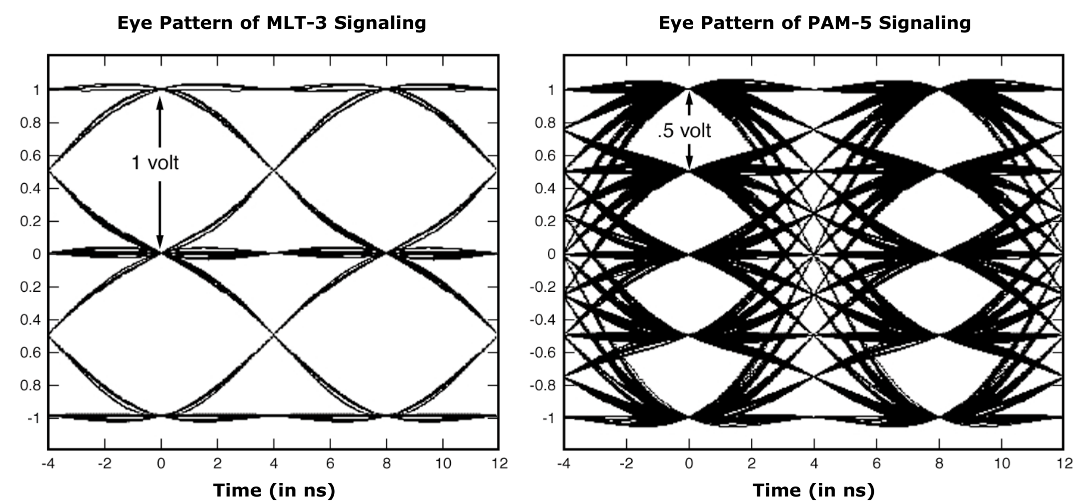


**Figure 1: MLT-3 and PAM-5 Signal Levels**

#### 4.1.1 Forward Error Correction (FEC)

The Gigabit Ethernet standard specifies the use of 5-level PAM-5 coding, although only 4-levels are required to represent 2-bits of information.The extra level provides redundancy for use in FEC. FEC provides a second level of coding that helps to recover the transmitted symbols in the presence of high noise and crosstalk. The error correction method consists of Trellis coding (4-Dimensional 8-State) in combination with Viterbi decoding. It is estimated that the FEC error correction logic enhances the system's SNR margin by up to 6 dB. Based on this assumption, the extra 6 dB of SNR margin gives the 5 level encoded signal the noise immunity of a 3 level signal (i.e., similar to the Fast Ethernet's MLT-3 three-level encoding).

## 4.2 Total Noise in a Gigabit Ethernet Channel

Each twisted-pair cable in a 1000BASE-T Gigabit Ethernet channel is subject to three major sources of noise; crosstalk noise, transmit echo noise and ambient noise. Crosstalk noise consists of two components, near-end crosstalk (NEXT) and far-end crosstalk (FEXT). Each noise component is discussed in the following sections.

### 4.2.1 Crosstalk Noise: NEXT & FEXT

Crosstalk is defined as unwanted or undesirable signals caused by the electric or magnetic fields of one signal affecting another adjacent signal. As there are a bundle of four adjacent signals in Gigabit Ethernet transmission, crosstalk becomes a complicated problem to analyze (see figure below). Crosstalk is characterized in reference to the transmitter. Near-end crosstalk (NEXT) is crosstalk that appears at the output of a wire pair at the transmitter end of the cable and far-end crosstalk (FEXT) is crosstalk that appears at the output of a wire pair at the far end of the cable from the transmitter (see Figure 2).

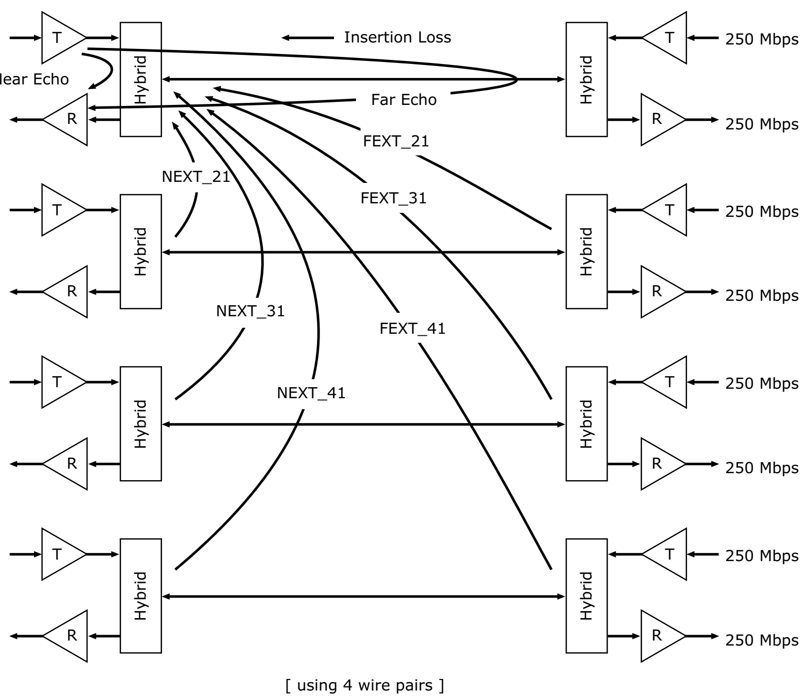


**Figure 2: Sources of Noise in a Gigabit Ethernet System**

### 4.2.2 Transmit Echo Noise

As previously stated, Gigabit Ethernet achieves 1000 Mbps operation, in part, by transmitting and receiving simultaneously over each of the four twisted-pair cables (250 Mbps equivalent data rate per line). As a direct result of this full duplex transmission, transmit echo noise is introduced to the system. In simplified terms, the outbound transmit signal reflects (or echos) off the channel (or cable) and creates an unwanted signal at the receiver of the same channel. This echo noise signal interferes with the actual receive signal. The IEEE Gigabit Ethernet standard defines echo cancellers which attempt to separate the transmit echo from the received signal. This concept of echo also applies to a telephone connection over a single twisted pair. Your phone has built-in echo canceling (much more simplified than Gigabit Ethernet) so you don't hear your voice reflected back.

To implement bi-directional full duplex transmission over each of the four Category 5 pairs, the transmitter and receiver are connected to each pair through a directional coupler circuit, known as a hybrid, which separates the outbound transmit signal from the inbound receive signal. Hybrid networks with good trans-hybrid loss minimize the amount of transmitter signal that is coupled into the receiver but cannot remove all of the transmit echo signal. The residual transmit signal due to the trans-hybrid loss combined with the cable return loss produce the unwanted echo noise. The magnitude of the reflection, or echo, is proportional to the return loss of the channel. Return loss, usually expressed in dB, may be defined as a measure of the dissimilarity between impedances in metallic transmission lines and loads. The return loss is a measure of the amount of power reflected due to cabling impedance mismatches—the higher the return loss, the greater the cable impedance mismatch.

### 4.2.3 Ambient Noise

Some examples of typical sources of ambient noise include background white noise and impulse noise generated by power lines and telephone voltages. Ambient noise can also include interfering wireless signals and alien crosstalk. Due to its random nature, the ambient noise cannot be cancelled in the receiver and so directly detracts from the SNR margin of the system. Thus, the best defense against ambient noise is a robust system with extra SNR margin. The higher the SNR margin, the less prone the system will be to errors caused by ambient noise.

## 4.3 Attenuation

Attenuation is defined as the reduction of signal strength during transmission across the cable (it is the opposite of amplification). Attenuation, measured in dB, increases (worsens) with increased frequency and distance. If the signal attenuates too much, it becomes unintelligible, and will result in network errors. Gigabit Ethernet (as well as 10 and 100 Mbps Ethernet) define a maximum cable distance of 100 meters. Thus, given known attenuation models for Category 5 cable, it may be determined what the minimum transmit amplitude should be to guarantee a certain minimum distinguishable voltage level at the downstream receiver—up to 100 meters away. Based on this logic, the IEEE Gigabit Ethernet standard defines these minimum voltage levels. It is important to note, however, that as the receiver sensitivity is increased, extra cable distance margin is achieved (i.e., in the presence of higher attenuation, the receiver can still distinguish voltage logic levels and accurately recover data).

## 5.0 Practical (or Real World) Challenges of Running Gigabit Ethernet Over Category 5 Cable

Section 4.0 of this paper reviewed some of the key issues/challenges in running 1000 Mbps of data over existing Category 5 cable, and what the IEEE standard for Gigabit Ethernet specifies as a solution. It is important to note that in order to generate an industry standard such as IEEE 802.3ab 1000BASE-T, several assumptions dictating the content of the standard had to be made. Below, a subset of the assumptions are highlighted:

- Standards-compliant Category 5 cabling is used as specified in ISO/IEC 11801:1995 and ANSI/EIA/TIA-568-A (1995) and tested for additional performance parameters specified in Section 40.7 of the IEEE 802.3ab standard.
- Four (4) twisted-pairs of the above-described standards compliant Category 5 cable are available.
- The network supports a BER of less than or equal to $10^{-10}$.
- The *horizontal cabling subsystem* is no greater than 100 meters as specified in ANSI/EIA/TIA-568-A and ISO/IEC 11801.
- The *horizontal cabling subsystem* consists of standards-compliant Category 5 patch cords, cables and connecting hardware as specified in ANSI/EIA/TIA-568-A and ISO/IEC 11801.
- The impedance of the Category 5 cable is 100 ohms +/- 15%.
- The Gigabit systems meet or exceed FCC Class A/CISPR operation.

The environment required by the 1000BASE-T standard for successful transmission of Gigabit data exists for the majority of the current installed base of network infrastructure. However, due to the extreme requirements and challenges of the 1000BASE-T standard, as well as the fact that cabling subsystems have been evolving over time, there are many circumstances in which a 1000BASE-T installation will fail. In the remainder of this section, we will explore some of these cases.

### 5.1 The Quality and Existence of Category 5 Cabling Is Not a Given

Like networking standards, cabling standards are constantly changing and evolving to meet the changing and growing needs of the industry. Cabling standards are primarily defined by the ANSI/EIA/TIA-568 specification. The initial ANSI/EIA/TIA standard was published as recent as 1991 and, at that time, was targeted for 10BASE-T electricals or Category 3 cable with a maximum frequency of 16 MHz. Prior to this specification, cabling was not standardized, and as a result, cabling existed that was sub-Category 3.

Additionally, a standard-defined Category 4 grade of cable was created with an operating frequency up to 20 MHz, and was intended for use in Token Ring networks. In 1995, Category 5 cable was standardized (100 MHz frequency), with the initial driving force being a copper version of the fiber-based FDDI networking protocol, and later, Fast Ethernet. With respect to Gigabit Ethernet performance, it's important to realize that not all Category 5 cable was/is created equal. Category 5 cable existed prior to the standard being ratified and as a result could result in poorer performance. Additionally, the performance of Category 5 cable is, in part, a function of the chemical compounds used in the manufacture of the cable. It is estimated that from 1994 on, there are 105 different electrical designs of Category 5 cable.

In summary, the majority of the installed base of cabling is Category 5 rated compliant to the ANSI/EIA/TIA-568-A standard; however, as cabling technology and standards have been evolving over the past decade or so, there exists sub-Category 5 cable in today's installed base. When constructing Gigabit Ethernet networks using system equipment that merely meets the IEEE 802.3ab 1000BASE-T standard in these less than ideal cabling installations, there exists the possibility for significant network performance problems.

### 5.2 The Length or Attenuation of the Cable Is Not Always to Specification

As stated previously, the ANSI/EIA/TIA-568-A and ISO/IEC 11801 standards specify that the horizontal cabling subsystem is to be no greater than 100 meters. The horizontal cabling subsystem is defined as the connection between a workstation and the local telecommunications closet and consists of cables, patch cords and connecting hardware (see Figure 3). With proper, certified installation it is very likely that the 100 meter maximum distance limit will be guaranteed—in fact, this is the case in the vast majority of new cabling installations. However, as with most things, there are exceptions. There are numerous cases where the 100 meter maximum distance limit is exceeded due to uncertified or careless installations, as well as intentional violations.  For example, consider a case where a small number of workstations slightly exceed the 100 meter distance limit to the wiring closet. As opposed to incurring the expense of installing an additional intermediary wiring closet, the distance limit may be stretched.

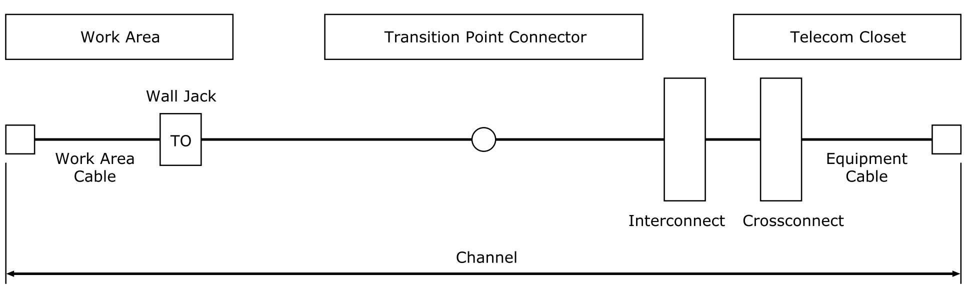


**Figure 3: ANSI/EIA/TIA Horizontal Cabling Subsystem**

As discussed, a key variable in determining cable attenuation is the distance of the cable—attenuation is proportional to cable distance. However, there are additional real-world variables that will affect/increase the attenuation of the cable. Some cabling installations will exhibit higher attenuation due to poorer grades of cable or due to imperfections in cabling installations resulting in kinked or distorted cables. The additional attenuation due to these imperfections will decrease the budget for cable length, or conversely, increase the effective cable length (i.e., a 100 meter cable will have the attenuation characteristics of a 120 meter cable).

### 5.3 Is the Horizontal Cabling Subsystem Completely Category 5 Rated?

As defined in Section 5.2, the horizontal cabling subsystem is defined as the connection between a workstation and the local telecommunications closet and consists of cables, patch cords and connecting hardware (see Figure 3). In addition to specifying the cabling, the ANSI/EIA/TIA-568 standard defines that the remaining components of the horizontal cabling subsystem—patch cords, connecting hardware, wall outlets, and punch-down blocks—be Category 5 or better. In real world installations, however, occasionally incorrect hardware is used. For example, a Category 3 connector may inadvertently be used. Incorrect or faulty connectors result in impedance mismatches in the link and contribute significantly to increased echo noise (as discussed in Section 4.2), which could exceed the SNR margin of the system and result in excessive BER.

### 5.4 Cable Installation Is a Factor

Cable installation is also a possible source of network errors, even when the proper standards-defined cable is used. Cable installation methods and procedures are also standardized by the EIA/TIA-568-A specification. Category 5 cable requires careful and precise installation to maximize efficiency and system performance. Anything that kinks the cable or disturbs the precise alignment of the pairs inside the cable has the potential to create future performance problems. Some sources of failure are as follows:

- Excessive force used when "pulling" the cables during installation. The EIA/TIA-568-A specification limits the pulling tension to 25 pounds maximum.
- Crushed or pinched cable bundles from sharp edges during installation or from applied weight.
- Crushed cable bundles from support hardware or staple guns.
- Improper termination to connecting hardware. To meet Category 5 cabling standards, the untwisting that takes place where cabling is terminated to connecting hardware should not exceed 13 millimeters, or a 1/2 inch.
- Loose, worn or abused wiring can cause the twisted pair to untwist more than necessary. The greater the distance that wires run parallel to one another, the greater likelihood that magnetic fields might cause impedance mismatch or induce current charges.
- Splicing of wires between the telecom closet and outlet locations.

## 6.0 Impact of Real World Gigabit Networking Considered

As highlighted in the above sections, there are several examples of real world networking environments that exceed the requirements and assumptions of the 1000BASE-T standard. For the successful market adoption and mass deployment of 1000BASE-T systems, it is critical that Gigabit Ethernet systems not only be compliant to the 1000BASE-T standard, but that they exceed the requirements of the standard enabling the deployment of Gigabit systems in the vast majority of networking environments.  Because in the end, when an end-user or IT personnel purchase and install 1000BASE-T Gigabit Ethernet equipment, it must work properly and transparently. This after all, is the legacy of Ethernet technology—it is simple and works with little or no intervention.

Networking systems manufacturers benefit greatly from ensuring that their equipment performs to specification, even in the event that the equipment is installed in an environment that may be in violation of the requirements of the 1000BASE-T standard. Systems manufacturers do not necessarily have control over the end-user's cabling plant/infrastructure, and thus by building extra performance margin into their products, they are able to cover a wider range of installations. When networking systems fail to function properly upon purchase and installation, the systems manufacturer is exposed to the following disadvantages:

- Increased cost for customer support and thus decreased product profit margins. Support includes phone calls into the support line as well as the potential for field visits.
- Negative impact to the systems manufacturer's quality and customer satisfaction level.
- Potential for lost business/revenue. The end-user or IT personnel may decide to replace the non-working systems with equipment from a competing supplier.
- For the customer, network reliability is a top priority. If the network equipment is unstable or unreliable in any given installation, there is higher probability of network down time. For the end-user, network downtime directly translates to lost productivity/revenue.

## 7.0 The Marvell Alaska® Gigabit Ethernet PHY Transceivers Reach Well Beyond the 1000BASE-T Standard

As highlighted in this section, the Marvell Alaska® family of Gigabit Ethernet PHY devices greatly exceed the IEEE 802.3ab 1000BASE-T standard. By utilizing Alaska technology, Gigabit Ethernet systems maximize the number of installations a given 1000BASE-T network may deploy. Some of the key advantages of the Marvell Alaska technology are:

- Full duplex Gigabit transmission up to 180 meters of Category 5 cable while maintaining a BER of $10^{-10}$ or better. This represents a 80% margin in cable distance relative to the 1000BASE-T standard.
- Tri-speed (10, 100, 1000 Mbps) operation up to 100 meters of **Category 3** cable.
- Superior echo cancellation circuitry permits full duplex Gigabit transmission up to 125 meters of Category 5 cable under worst case impedance mismatch scenario. The test involved 20% impedance mismatch (100 ohm cable to 120 ohm cable).
- Verified operation even with the use of non-Category 5 connectors and cords.
- Superior alien crosstalk noise rejection and better common mode noise rejection.
- Superior transmit wave shaping techniques resulting in FCC Class B/CISPR qualification.

Some of the advantages of the Alaska Gigabit Ethernet PHY transceiver technology are discussed in more detail in the following sections.

### 7.1 Industry-Leading DSP Technology

Simply stated, a Gigabit Ethernet PHY transceiver is only as good as the DSP technology used in its implementation. Marvell has a long sustained track record of delivering to production the industry's most advanced Digital Signal Processing (DSP) technology. Since 1995, Marvell has introduced five generations of industry-leading read channel ICs for the storage market. The read channel may be viewed as the PHY transceiver of the disk drive and is highly DSP centric. Additionally, there are several similar functional blocks between the read channel and 1000BASE-T physical layers. By leveraging the advanced DSP technology developed for the read channel, Marvell was able to quickly come to market with a highly advanced 1000BASE-T PHY device, greatly exceeding the performance requirements of the IEEE standard as well as competing solutions.

As discussed in previous sections, the clock rate of the 1000BASE-T PHY transceiver is 125 MHz. Thus, the DSP engines implementing the PHY device also run at 125 MHz. As shown in the Figure 4 below, during the development of the initial Alaska Gigabit PHY product, Marvell was in production with a read channel that operated at frequencies up to 750 MHz; therefore, the DSP engines were running at 750 MHz (1.3 GHz read channels are now currently available from Marvell). Thus, from a DSP technology standpoint, the requirement of 125 MHz DSP for the Gigabit PHY transceiver did not represent a significant technology barrier for Marvell. On the other hand, other semiconductor suppliers started with DSP technology operating at much lower speeds, and thus were forced to "stretch" the technology up to 125 MHz. When technology must be stretched, the result is less efficient designs from both a power and die size perspective, as well as an impact on the performance capability of the devices.

| PHY | System | Media | Clock Rate |
|---|---|---|---|
| Read Channel | Disk Drive | | 750 MHz |
| Transceiver | Fast and Gigabit Ethernet | | 125 MHz |
| Cable Modem Transceiver | Cable Modem | | 25 MHz |
| Baseband Processor | Wireless Ethernet | | 22 MHz |

Marvell

Competition

**Figure 4: The Marvell Industry-Leading DSP Technology**

## 7.2 Alaska PHY Exhibits Superior Cable Distance Performance

The Alaska Gigabit Ethernet PHY transceivers feature the industry's greatest cable distance margin, operating up to 180 meters of Category 5 cable in full duplex transmission mode. The "eye diagram" of the Alaska device after transmission over 140 meters of standard Category 5 Unshielded Twisted Pair (UTP) cable is shown in Figure 5. The key issue to highlight is that the "eyes" at 140 meters are wide open. The eyes are the black openings between the signal traces. Generally, a larger eye opening results in better BER performance. To determine the proper logic level, the waveform is sampled in the middle of the data "eye" in between data transitions as this is the location where the waveform is most settled, and the height of the eye opening is at its maximum. Other than the Marvell Alaska PHY devices, no other Gigabit PHY products in the industry operate at 140 meters—the link cannot be established, the eye does not exist (it is closed).

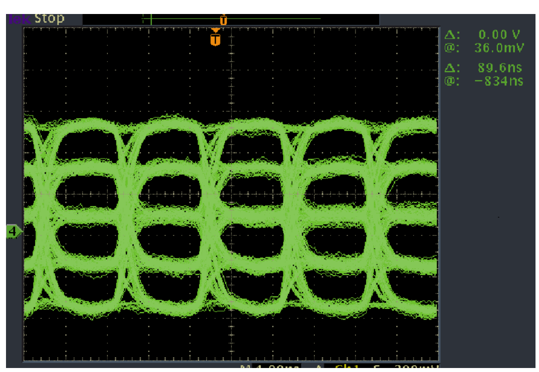


**Figure 5: The Eye Diagram at 140 Meters Category 5 Cable Using the Marvell Alaska Gigabit PHY Transceiver**

### 7.3 Alaska PHY Exhibits Superior Echo Cancellation Performance

As previously discussed, Gigabit Ethernet achieves 1000 Mbps operation, in part, by transmitting and receiving simultaneously over each of the four twisted-pair cables (250 Mbps equivalent data rate per line). As a direct result of this full duplex transmission, transmit echo noise is introduced to the system. The outbound transmit signal reflects (or echos) off the channel and creates an unwanted signal at the receiver of the same channel. Although the ANSI/EIA/TIA-568 standard defines that the components of the horizontal cabling subsystem (cable, patch cords, connecting hardware, wall outlets, and punch-down blocks) be Category 5 or better, occasionally incorrect hardware is used. The use of cabling not exhibiting 100 ohm impedance, or the use of incorrect or faulty connectors result in impedance mismatches in the link and contribute significantly to increased echo noise. The Marvell Alaska family of Gigabit Ethernet transceivers feature extremely robust echo cancellation circuitry resulting in added SNR margin and proper functionality meeting or exceeding the specified BER— even in the presence of these non-standard cabling environments.

Figure 6 shows a test setup defined in the IEEE standard and the UNH Interoperability Lab. To create impedance mismatch, 1 meter segments of 120 ohm cables are used at each end of the wiring subsystem. The 100 ohm impedance Category 5 cable will vary in length for this test. In this test scenario, a 20% change in impedance causes a 9% reflection (or echo) of the signal back to the transmitting node. In the above test setup, systems containing the Marvell Alaska PHY transceivers function properly with the 100 ohm Category 5 cable segment up to 125 meters in length. Running the same test using Gigabit Ethernet systems based on non-Marvell devices, result in errors (down link) beyond 41 meters of 100 ohm Category 5 cable.

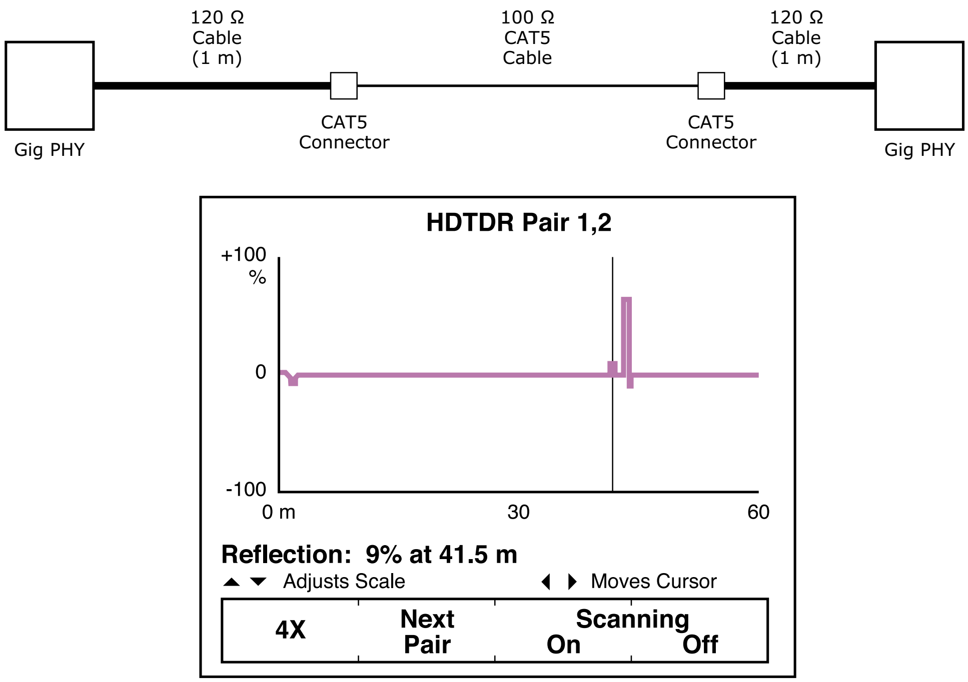


**Figure 6: Gigabit Ethernet Impedance Mismatch Test Setup**

The Alaska Gigabit PHY devices feature extremely robust noise cancellation circuitry, greatly exceeding the requirements of the standard. The DSP-based NEXT and echo cancellers implemented in the Alaska PHY transceivers are the largest functional block of IC (Integrated Circuit), consuming nearly 40% of die real estate. There are 160 echo taps and 20x3 NEXT taps for each of the four channels contained in the Alaska device. Competing Gigabit PHY transceiver solutions typically implement only as many as 60 taps.

### 7.4 Alaska PHY Exhibits Superior EMI Performance

As reviewed in Section 5.0, the 1000BASE-T standard requires only that Class A/CISPR FCC requirements be met. Using superior pulse shaping techniques, the Alaska family of devices meets the more stringent Class B/CISPR FCC requirements. Pulse shaping is used to minimize the transmitted signal energy at frequencies where distortion and disturbances are significant. Rather than typical wave shaping filters, the Alaska devices use a complex SIN(X)/X function for superior performance.

### 7.5 The 2-Pair Problem of Gigabit Ethernet

As reviewed in prior sections, the Gigabit Ethernet standard requires four Category 5 pairs for Gigabit transmission. However, in actual cabling systems, there are cabling installations that contain only two Category 5 cables. Typically cabling bundles include four twisted-pair conductors; however, in many cases IT personnel have routed only two twisted-pair conductors to an end-user. For example, consider a case where a 4-pair bundle is split between two users in order to reduce installation costs. This would be a common practice particularly in environments of employee expansion (expansion beyond the existing cable infrastructure), or in environments where end-users are consolidated in a single office/cubicle.

Prior to Gigabit Ethernet technology, splitting a four-pair bundle into two sets of two twisted-pairs could be viewed as acceptable, as the Ethernet technology at the time (both 10BASE-T and 10/100BASE-T) required only two twisted-pairs for transmission. Today, deploying Gigabit Ethernet systems in a two-pair cabling environment would result in disaster—a link or connection would not occur. In this environment, the Auto-Negotiation process would be carried out successfully as this scheme involves transmission on only 2 twisted-pairs. Further, the Auto-Negotiation process would inform the Gigabit systems at each end of the link that the other is Gigabit ready. Data transmission would proceed, however, this will result in failure as Gigabit data transmission requires four pairs. The data link would then go back into Auto-Negotiation mode and this process will continue infinitely.

Marvell recognized this issue and as a result designed a feature into the Alaska PHY devices that will automatically detect this 2-pair environment and downshift operation to 100 Mbps. Downshifting to 100 Mbps enables the network to properly link. Under similar conditions, competing devices cannot link and the network would go down.

## 8.0 The Marvell Alaska® Gigabit Ethernet Physical Layer Family: Product Overview

The Marvell Alaska® family of Gigabit Ethernet over copper transceivers, currently consisting of sixteen different devices, has offered the networking industry the enabling technology required to ease and accelerate the deployment of Gigabit Ethernet to the desktop and throughout the network. Marvell continues to be the undisputed technology and market leader in the Gigabit Ethernet PHY transceiver marketplace by setting multiple industry and technology milestones with its single, dual and quad-port Alaska family of Gigabit Ethernet transceivers.

With its first generation single-port Alaska and Alaska+ PHY transceivers, which are currently in volume production, Marvell set a new technology standard by offering the lowest power dissipation, smallest form factor packaging and highest performance devices. Leveraging this Gigabit Ethernet technology expertise, Marvell again led the market with the introduction of its Alaska II and Alaska II+ PHY transceivers—the world's first dual-port Gigabit Ethernet PHY transceivers offering the networking industry even higher integration and lower system cost.

Targeted at high-density Gigabit Ethernet switches, the Marvell Alaska Quad Gigabit Ethernet transceiver family represents a significant technology breakthrough for the Gigabit Ethernet market by offering the world's most complex and highly integrated mixed-signal DSP-based quad-port communications devices ever developed. The Alaska Quad family offers the industry's lowest power and highest integration Gigabit Ethernet over copper PHY transceivers. When used with stacked RJ45s, the Alaska Quad device enables 48-port Gigabit systems.

All of the Marvell Alaska Gigabit Ethernet transceivers, the Alaska+, Alaska II+ and Alaska Quad+ devices, provide system manufacturers with the ability to bridge fiber-optic and copper networks with the built-in serializer/deserializer (SERDES) function.

### 8.1 The Marvell Alaska Ultra Family of Single-Channel 1000BASE-T PHY Transceivers

The Marvell Alaska Ultra device is the industry’s first Gigabit Ethernet over copper transceiver optimized for the mobile computing market. The Alaska Ultra transceiver delivers the lowest power consumption, smallest form factor package and advanced power management modes. The device supports advanced power management modes including support for Wake-on-LAN, reducing power consumption to near-zero. These low power management modes are critical to client applications such as Gigabit-connected laptop computers, where the use of the Alaska Ultra device can substantially prolong battery life.

The Marvell Alaska Ultra+ device, the 88E1011S, is the industry's first and only Gigabit solution to provide complete flexibility in media selection—either copper or fiber media. The Alaska Ultra+ transceiver offers simultaneous support of both copper and fiber interfaces, whereby the copper port is connected to a transformer and RJ45 connector, and the fiber port is connected to a Gigabit fiber-optic module or GBIC device. The device's Media Detect feature automatically senses whether the end-user has inserted a Category 5 copper or fiber-optic cable, performs Auto-Negotiation and provides seamless plug-and-play configuration.

| Part Number | Number of Ports | Package | PHY/MAC Interface | SERDES Media Interface | Media Detect |
|---|---|---|---|---|---|
| **88E1011** | 1 | 117-pin TFBGA or 128-pin PQFP | GMII, RGMII, SGMII, TBI, RTBI | No | No |
| **88E1011S** | 1 | 117-pin TFBGA or 128-pin PQFP | GMII, RGMII, TBI, RTBI | Yes | Yes |

**Table 1: The Marvell Alaska Ultra Device Seletion Table**

## 8.2 The Marvell Alaska Quad Family of 1000BASE-T Gigabit Ethernet PHY Transceivers

The Marvell Alaska Quad Gigabit Ethernet transceiver family represents a significant technology breakthrough for the Gigabit Ethernet market by offering the world's most complex and highly integrated mixed-signal DSP-based quad-port communications devices ever developed. The Alaska Quad family offers the industry's lowest power and highest integration Gigabit Ethernet over copper (10/100/1000BASE-T) PHY transceivers. When used with stacked RJ45's, the Alaska Quad enables 48-port Gigabit systems. The Alaska Quad+ (88E1040S/88E1041S/88E1042S) transceivers provide flexibility in the selection of copper or fiber-optic interfaces on a per-port basis using a built-in SERDES function.

| Part Number | Number of Ports | Package | Package Dimension | Package Ball Pitch | PHY/MAC Interface | SERDES Media Interface |
|---|---|---|---|---|---|---|
| **88E1040/ 88E1040S** | 4 | 352-pin TBGA | 35mm x 35mm | 1.27mm | GMII, RGMII, SGMII, TBI, RTBI | 88E1040S Only |
| **88E1041/ 88E1041S** | 4 | 352-pin TBGA | 27mm x 27mm | 1.0mm | GMII, RGMII, SGMII, TBI, RTBI | 88E1041S Only |
| **88E1042/ 88E1042S** | 4 | 256-pin TBGA | 27mm x 27mm | 1.27mm | RGMII, SGMII, RTBI | 88E1042S Only |

**Table 2: The Marvell Alaska Quad Device Seletion Table**

## 9.0 References


[1] "Physical Layer Parameters and Specifications for 1000 Mbps Operation Over 4-Pair of Category 5 Balanced Copper Cabling, Type 1000BASE-T," IEEE Std. 802.3ab, 1999.

[2] He, R., et al., "A DSP Based Receiver for Gigabit Ethernet over CAT-5 Cables," *ISSCC Digest of Technical Papers*, February 2001.

[3] P. Roo, et al., "A CMOS transceiver analog front-end for Gigabit Ethernet over CAT-5 cables." Submitted to *ISSCC 2001, Subject Area: Wireline communications*.

[4] "Gigabit Ethernet 1000BASE-T Whitepaper." Gigabit Ethernet Alliance, 1997.

[5] "Running 1000BASE-T: Gigabit Ethernet Over Copper Whitepaper." Gigabit Ethernet Alliance, September 1999.

[6] "Gigabit Ethernet Over Category 5." Agilent Corporation.

[7] "802.3ab, A Tutorial Presentation."

[8] "Category 5: How Did We Get Here and Where Do We Go Next?" Anixter Inc, 1997.

[9] "Do You Really Know Your Cable Configuration?" Danpex Corporation Web Site. http://www.danpex.com/faqs/cable.htm/.

**Marvell Semiconductor, Inc.**

700 First Avenue
Sunnyvale, CA 94089

Phone 408.222.2500

www.marvell.com