The Status of the 10-gigabit Ethernet Standard

White Paper

By David G. Cunningham

Agilent Technologies
Network Solutions Division
Semiconductor Products Group

Agilent Technologies

Abstract

This paper discusses the status of the 10-Gigabit Ethernet standard along with its objectives and architecture. Ten-gigabit Ethernet is designed for use in two major applications: firstly, local area networks (LAN) and, secondly, connection to SONET/SDH-based wide area networks (WAN). The use of 10-gigabit Ethernet in these applications is introduced. The four port types being specified for LAN applications and the three for WAN applications are reviewed.

Introduction

In common with previous IEEE 802.3 standards[1], 10-Gigabit Ethernet (IEEE 802.3ae) will define a standard that guarantees interoperation between different vendors' implementations. Essentially, the standard will specify physical layers (PHY); only a very slight change will be made to the medium access control (MAC). A major theme of earlier versions of Ethernet has been the pragmatic adoption of cost effective but robust technologies. In large part, this enabled Ethernet to dominate the LAN market. One of the major challenges addressed by the standards effort has been the development of specifications that are friendly to directly modulated lasers—it is believed this will facilitate very cost effective implementations. It is important to note that 10-Gigabit Ethernet represents the coming together of both data communications and telecommunications.

Some of the important features adopted by 10-Gigabit Ethernet are:

- Wide range of cost/reach options
- much longer maximum reach than previous Ethernets
- a four bit wide electrical bus extender (XAUI)
- a very low overhead, scrambler-based, 64B66B code
- an option for transport in SONET/SDH like frames
- two serial physical layer types
- a coarse or wide wavelength division multiplexed (WWDM) physical layer
- line rates of 10.3125 (LAN), 9.95328 (WAN, OC-192 rate) and 3.125 (LAN, 4 wavelengths) GBd

The standardization of 10-Gigabit Ethernet began during March of 1999. Currently, the draft standard is in its final review process and it will become a standard during 2002. For 10-Gigabit Ethernet, this paper will summarize the following: its objectives, its applications, its key features and, some implementation examples.

Objectives of 10-Gigabit Ethernet

The objectives for 10-Gigabit Ethernet can be summarized as follows:

- Full duplex operation (continuous transmission, point-to-point, bi-directional, dual optical fiber links)
- Support for a MAC data rate of 10 Gb/s
- The specification of two families of physical layer:
  - LAN PHY, operating at a MAC data rate of 10 Gb/s
  - WAN PHY, operating at a MAC data rate compatible with OC-192c/SDH VC-4-64c
- Definition of a mechanism to adapt the MAC data rate to the data rate of the WAN PHY
- The support of both single-mode (SMF) and multimode (MMF) fiber types
- Specification of a 10-Gigabit Ethernet media independent electrical interface (XGMII)
- Backward compatibility with earlier versions of Ethernet: same frame format, same frame size etc.

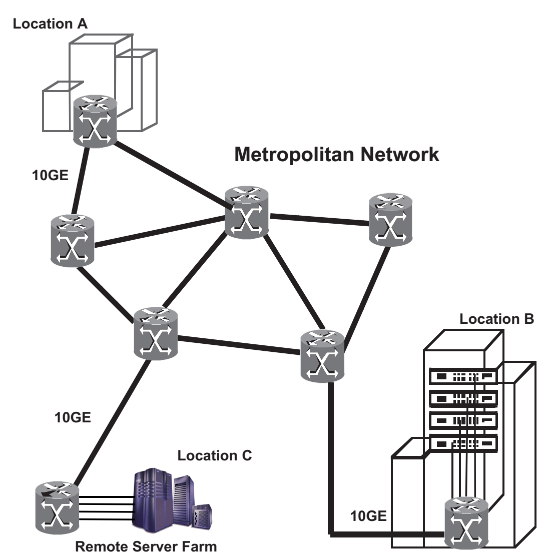

Figure 1: A native 10-Gigabit Ethernet MAN.

The application of 10-Gigabit Ethernet

Initially, 10-Gigabit Ethernet will be a switch-to-switch interconnection for statistically multiplexing packet traffic from lower data rate (10/100/1000 Mb/s) Ethernets. Therefore, 10-Gigabit Ethernet is primarily a backbone technology that is targeted at the enterprise LAN or the telecom WAN. This use of 10-Gigabit Ethernet is illustrated in Fig. 1, which shows various locations interconnected by a 10-Gigabit Ethernet MAN. At each location the 10-Gigabit Ethernet capable switch/router multiplexes the traffic from the local 10/100/1000 Ethernet LAN onto the 10-Gigabit Ethernet MAN—no end station (computer) has a 10 Gb/s connection.

Overview of the 10-Gigabit Ethernet Standard

The draft standard specifies seven port types as listed in Table 1. Six of the port types use bit serial optical transmission whilst the remaining port type multiplexes MAC data across four wavelengths. The WWDM physical layer (see Table 1) can support both multimode and single-mode fiber.

10-Gigabit Ethernet Port Types
As can be seen from Table 1, two categories of port types are defined:

- LAN PHY for native Ethernet applications
- WAN PHY for connection to the installed base of SDH/SONET 10 Gb/s networks

Layered model for 10-Gigabit Ethernet

The layered model for 10-Gigabit Ethernet is shown in Fig. 2. Sublayers for the two families of PHY (LAN and WAN) are included in the diagram. Also shown are the specified interfaces as follows:

- XGMII, 10-Gigabit Media Independent Interface
- XAUI, 10-Gigabit Attachment Unit Interface
- XSBI, 10-Gigabit 16-bit Interface
- MDI, Media Dependent Interface

For clarity the XGMII extender sublayer (XGXS) is not shown in Fig. 2. However, if XAUI is used to extend the XGMII then the XGXS is required too.

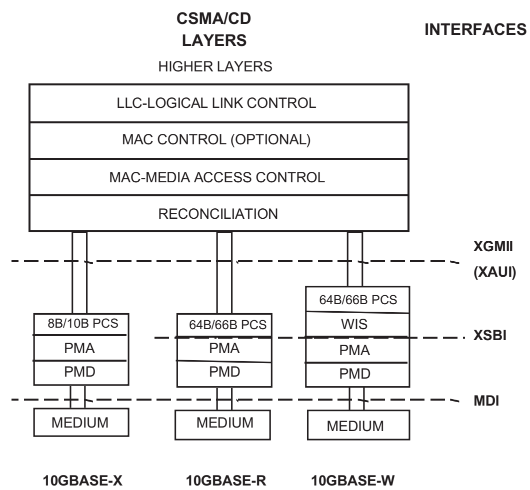

Figure 2: Layered model for 10-Gigabit Ethernet.

10-Gigabit Ethernet MAC
Obviously, the normal MAC data rate (the rate at which the MAC transfers its information to the PHY) for 10-Gigabit Ethernet is 10 Gb/s. However, for WAN PHY operation the MAC data rate must be throttled down to the slightly lower data rate (9.58464 Gb/s) of SDH/SONET equipment by a very slight change to the MAC. This is achieved by dynamically adapting the interframe spacing. Even in past versions of Ethernet the interframe spacing could be lengthened whilst remaining compliant with the Ethernet standard—at the cost of lower throughput.

Table 1. 10-Gigabit Ethernet Port Types.

Description	Name	Comments
850 nm Serial LAN PHY	10GBASE-SR	Directly modulated VCSEL, MMF, 2-300 m
1310 nm Serial LAN PHY	10GBASE-LR	Directly modulated DFB laser, SMF, 2-10 km
1550 nm Serial LAN PHY	10GBASE-ER	Modulator, DFB laser, SMF, 2-40 km
1310 nm WWDM LAN PHY	10GBASE-LX4	Directly modulated VCSEL or DFB, MMF (300 m) or SMF (2-10 km)
850 nm Serial WAN PHY	10GBASE-SW	Directly modulated VCSEL, MMF, 2-300 m
1310 nm Serial WAN PHY	10GBASE-LW	Directly modulated DFB laser, SMF, 2-10 km
1550 nm Serial WAN PHY	10GBASE-EW	Modulator, DFB laser, SMF, 2-40 km

WWDM LAN PHY
The WWDM LAN PHY uses a physical coding sublayer (PCS) which is based on four lanes of 8B10B code (the same code as used by Gigabit Ethernet). Each lane operates at a data rate of 2.5 Gb/s and a line rate of 3.125 GBd. On reception the 8B10B PCS can tolerate interlane skews much greater than that which could be generated by the transmission of the four different wavelengths through 10 km of single-mode fiber. The physical media attachment (PMA) conditions the signal levels for transmission by the WWDM physical medium dependent (PMD) sublayer. In some circumstances additional eye opening circuits may be required between the PMA and the WWDM PMD. The WWDM PMD includes four laser transmitters, laser drivers, an optical multiplexer, an optical demultiplexer, four photoreceivers and post amplifiers. Most likely the Management Data Input/Output (MDIO) electronics and interface will also be present.

Serial LAN PHY
Initially it was thought that the serial PHY would reuse the 8B10B code (from Gigabit Ethernet). From a standards viewpoint this was very attractive since, if this were true, all that would have been required to specify 10-Gigabit Ethernet would be to copy Gigabit Ethernet with a scaled bit time. However, it was quickly realized that directly modulated lasers have difficulty transmitting 10 Gb/s signals. It was generally viewed that 12.5 GBd, as required for 8B10B coding, was unreasonable for cost effective serial implementations. Therefore, the committee adopted a PCS that provides the codec function for a new more efficient 64B66B code. For serial LAN PHY operation this reduced the serial baud rate to 10.3125 GBd.

The PMA for the serial PHY is essentially a serialize-deserialize function. In the standard, for specification convenience, the interface between the serial PCS and the serial PMA is assumed to be sixteen bits wide. The sixteen wide data path also matches the data path of current generation 10 Gb/s optical transponder products. The serial PMD includes a laser, laser driver, and optical receiver. As shown in Table 1 three wavelengths of operation are specified, 850 nm, 1310 nm, and 1550 nm respectively. The worst case, maximum link lengths for these PMDs are at 300 m (new high bandwidth multimode fiber), 10 km (single-mode fiber) and 40 km (single-mode fiber) respectively.

Serial WAN PHY
For the serial WAN PHY an additional sublayer, the WAN Interface Sublayer (WIS) is required between the serial PCS and the serial PMA. The basic functions of this sublayer are: to map the output of the serial PCS into a SDH/SONET like frame, to map the contents of received SDH/SONET like frames into the PCS format and to process the SDH/SONET like frame overhead including pointers and parity checks.

Reconciliation sublayer and the 10-Gigabit Ethernet Interfaces
The reconciliation sublayer (RS) adapts the protocol of the Ethernet MAC into the parallel encoding of the 10 Gb/s PCS. Although the physical implementation of the XGMII is optional, for the purposes of specifying 10-Gigabit Ethernet, the XGMII is assumed to be the interface between the RS and PCS sublayers. The XGMII uses 32 bit data paths that are partitioned into four transmit and four receive lanes, 8 bits per lane. Also, each lane has a control bit associated with it. The RS maps MAC data octets to (from) the lanes of the XGMII in round-robin order. At the request of the MAC or PHY the RS also maps MAC control signals to (from) the XGMII.

Optionally, the transmission distance of the XGMII can be extended using the XGXS and XAUI. Both XGXS and XAUI use the 10GBASE-X (see Fig. 2), PCS and PMA. XAUI associates one of its serial 8B10B lanes, operating at a data rate of 3.125 Gb/s, to each XGMII lane. Essentially, the XGXS and the XAUI interface provide a narrow 4 bit wide, self timed, full duplex, data bus. Repetitive XAUI control signals (for example, idle) are scrambled to prevent excessive electromagnetic interference.

The optional sixteen bit wide interface between the serial PCS or WIS and the serial PMA is called the 10-Gigabit sixteen-bit interface (XSBI). The XSBI is a fully differential, LVDS, clocked interface that is very similar to the SFI-4 interface of the Optical Internetworking Forum (OIF). For specification convenience the standard is written in terms of the XSBI.

Implementation examples

The discussion so far has been rather abstract. Therefore, in this section we will show how two common types of optical transceivers for 10-Gigabit Ethernet are likely to be implemented. Of course, the Ethernet networking equipment that uses these transceivers must include the implementation of the higher layers of the standard as illustrated in Figures 3 and 4.

Front panel, connectorized (no fiber pigtails), electrically hot pluggable, optical transceivers have become very popular for 100 Mb/s and 1000 Mb/s Ethernet equipment. Many equipment manufacturers would like to continue this methodology for 10-Gigabit Ethernet equipment. The XAUI interface provides a narrow electrical interface for such through-panel, pluggable, optical transceivers. A simplified block diagram of an XAUI-based, hot pluggable optical transceiver is shown in Fig. 3. A multisource agreement (MSA) called XENPAK[2] for XAUI-based transceivers has been published.

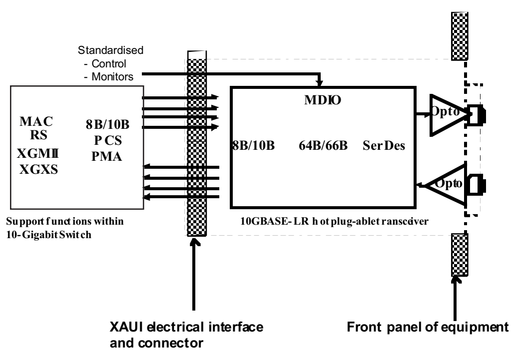

Figure 3: XAUI-based, hot pluggable optical transceiver.

Front panel, electrically hot pluggable, optical transceivers are less common in telecommunications equipment. For telecommunications applications it is more usual for the optical transceiver to be pluggable but mounted toward the rear of a linecard. Additionally, telecommunications optical transceivers usually have a connectorized pigtail that is connected to the front panel of the linecard. A simplified block diagram of an XSBI-based, pluggable, optical transceiver is shown in Fig. 4.

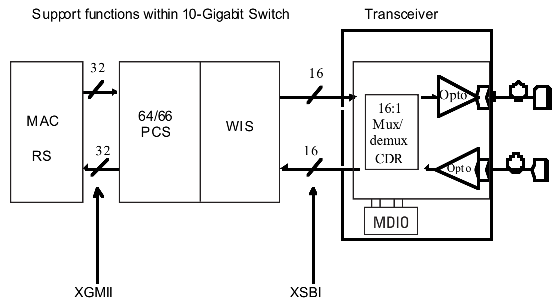

Figure 4: XSBI-based, pluggable, optical transceiver.

Concluding remarks

This paper has reviewed the status of the 10-Gigabit Ethernet draft standard, its target applications, its objectives, its layers and two example optical transceiver types. The draft standard is now reasonably stable: the MAC, RS, XGMII, XGXS, XAUI, PCS, WIS, PMA sublayers are essentially finished and generate few comments during ballots. Only some fine details regarding the PMDs are still in a slight state of flux. Many component and network equipment manufacturers will soon begin interoperation testing. The results of these tests are expected to validate the draft standard and quicken its completion. Therefore, during 2002, expect 10-Gigabit Ethernet to come to a network near you.

References

1. http://standards.ieee.org/getieee802/
2. http://www.xenpak.org/

www.agilent.com/semiconductors

For product information and a complete list of distributors, please go to our web site.

For technical assistance call:
Americas/Canada: +1 (800) 235-0312 or (408) 654-8675

Europe: +49 (0) 6441 92460

China: 10800 650 0017

Hong Kong: (+65) 271 2451

India, Australia, New Zealand: (+65) 271 2394

Japan: (+81 3) 3335-8152(Domestic/International), or 0120-61-1280(Domestic Only)

Korea: (+65) 271 2194

Malaysia, Singapore: (+65) 271 2054

Taiwan: (+65) 271 2654

Data subject to change.

December 12, 2001
5988-4762EN

Agilent Technologies