

Gigabit Ethernet

Raj Jain
Professor of Computer and Information Sciences
The Ohio State University
Columbus, OH 43210
<http://www.cis.ohio-state.edu/~jain/>

The Ohio State University

Raj Jain

9-1

MBone Instructions

- ❑ Handouts for the class are available on-line:
<http://www.cis.ohio-state.edu/~jain/cis788-97/index.html> or
<http://www.netlab.ohio-state.edu/~jain/cis788-97/index.html> or
<ftp://netlab.ohio-state.edu/pub/jain/cis788-97/>
- ❑ The schedule keeps changing. Please always check current schedule at:
<http://www.cis.ohio-state.edu/~jain/cis788-97/schedule.html>

The Ohio State University

Raj Jain

9-2

Instructions (Cont)

- ❑ Please email your positive and negative feedback about the quality of the reception as well as the content with a subject field of “**Feedback**” to mbone@netlab.ohio-state.edu
- ❑ If you are not able to receive the program due to some technical difficulties, please email “**Feedback**” to mbone@netlab.ohio-state.edu
- ❑ Please email technical questions with the subject field “**Question**” to mbone@netlab.ohio-state.edu. We will try to answer selected questions live.



- ❑ Current design goals
- ❑ PHY Issues
- ❑ MAC Issues
- ❑ Schedule

Since the Gigabit Ethernet standard is still in a draft stage, some statements made here may change.

Gigabit Ethernet

- ❑ Being standardized by 802.3z
- ❑ Project approved by IEEE in June 1996
- ❑ 802.3 meets every three months \Rightarrow Too slow
 \Rightarrow Gigabit Ethernet Alliance (GEA) formed.
It meets every two weeks.
- ❑ Decisions made at GEA are formalized at 802.3 High-Speed Study Group (HSSG)
- ❑ Based on Fiber Channel PHY
- ❑ Shared (half-duplex) and full-duplex version
- ❑ Gigabit 802.12 and 802.3 to have the same PHY

The Ohio State University

Raj Jain

9-5

How Much is a Gbps?

- ❑ 622,000,000 bps = OC-12
- ❑ 800,000,000 bps (100 MBps Fiber Channel)
- ❑ 1,000,000,000 bps
- ❑ 1,073,741,800 bps = 2^{30} bps ($2^{10} = 1024 = 1k$)
- ❑ 1,244,000,000 bps = OC-24
- ❑ 800 Mbps \Rightarrow Fiber Channel PHY
 \Rightarrow Shorter time to market
- ❑ Decision: 1,000,000,000 bps \Rightarrow 1.25 GBaud PHY
- ❑ Not multiple speed \Rightarrow Sub-gigabit Ethernet rejected
- ❑ 1000Base-X

The Ohio State University

Raj Jain

9-6

Physical Media

- ❑ Unshielded Twisted Pair (UTP-5): 4-pairs
- ❑ Shielded Twisted Pair (STP)
- ❑ Multimode Fiber: 50 μm and 62.5 μm
 - Use CD lasers
- ❑ Single-Mode Fiber
- ❑ Bit Error Rate better than 10^{-12}

How Far Should It Go?

- ❑ Full-Duplex:
 - Fiber Channel: 300 m on 62.5 μm
at 800 Mbps \Rightarrow 230 m at 1000 Mbps
 - Decision: 500 m at 1000 Mbps
 \Rightarrow Minor changes to FC PHY
- ❑ Shared:
 - CSMA/CD without any changes
 \Rightarrow 20 m at 1 Gb/s (Too small)
 - Decision: 200 m shared
 \Rightarrow Minor changes to 802.3 MAC


Gigabit Ethernet Objectives

- ❑ 1000 Mb/s MAC
- ❑ 802.3 Ethernet Frame format
- ❑ Meet all 802 requirements except possibly Hamming distance
- ❑ Preserve min and max frame size of 802.3
- ❑ Full and half-duplex operation
- ❑ Support star-wired topologies
- ❑ Use CSMA/CD with at least 1 repeater

- ❑ Support Fiber and, if possible, copper
- ❑ At least 500 m over multimode fiber
- ❑ At least 25 m over copper
 - ⇒ Wiring-closet or data center backbone
 - 100 m desirable
- ❑ At least 2 km on single mode fiber
- ❑ Collision domain diameter of 200 m
- ❑ Accommodate 802.3x flow control
- ❑ Price: 2 or 3 times fast Ethernet ports and NICs

PHY Issues

- ❑ Fiber Channel PHY:
 - 100 MBps = 800 Mbps
 - ⇒ 1.063 GBaud using 8b10b
- ❑ Changes to get 500 m on 62.5- μ m multimode fiber
 - Modest decrease in rise and fall times of the transceivers



The Ohio State University

Raj Jain

9-11

- ❑ Symbol Codes for Specific Signals: Jam, End-of-packet, beginning of packet
- ❑ PHY-based flow Control: No. Use the XOFF flow control of 802.3x

The Ohio State University

Raj Jain

9-12

850 nm vs 1300 nm lasers

- ❑ 850 nm used in 10Base-F
 - Cannot go full distance with 62.5- μ m fiber
 - 500 m with 50- μ m fiber
 - 250 m with 62.5- μ m fiber
- ❑ 1300 nm used in FDDI but more expensive
 - Higher eye safety limits
 - Better Reliability
 - Start with 550 m on 62.5- μ m fiber
 - Could be improved to 2 km on 62.5- μ m fiber
 - ⇒ Needed for campus backbone

The Ohio State University

Raj Jain

9-13

Media Access Control Issues


- ❑ Carrier Extension
- ❑ Frame Bursting
- ❑ Buffered Distributor

The Ohio State University

Raj Jain

9-14

Carrier Extension




- ❑ 10 Mbps at 2.5 km \Rightarrow Slot time = 64 bytes
- ❑ 1 Gbps at 200 m \Rightarrow Slot time = 512 bytes
- ❑ Continue transmitting control symbols.
Collision window includes the control symbols
- ❑ Control symbols are discarded at the destination
- ❑ Net throughput for small frames is only marginally better than 100 Mbps

The Ohio State University

Raj Jain

9-15

Frame Bursting




- ❑ Don't give up the channel after every frame
- ❑ After the slot time, continue transmitting additional frames (with minimum inter-frame gap)
- ❑ Interframe gaps are filled with extension bits
- ❑ No no new frame transmissions after 8192 bytes
- ❑ Three times more throughput for small frames

The Ohio State University

Raj Jain

9-16

Buffered Distributor




- ❑ All incoming frames are buffered in FIFOs
- ❑ CSMA/CD arbitration inside the box to transfer frames from an incoming FIFO to all outgoing FIFOs
- ❑ Previous slides were half-duplex. With buffered distributor all links are full-duplex with frame-based flow control
- ❑ Link length limited by physical considerations only

The Ohio State University

Raj Jain

9-17

GMII



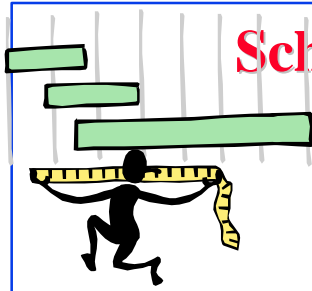
- ❑ Gigabit Media Independent Interface
- ❑ MII between MAC and PHY
- ❑ Allows any PHY to be used with a given MAC
⇒ Allows fiber channel phy to be used with CSMA/CD MAC and TX
- ❑ Allows both full-duplex and half-duplex modes

The Ohio State University

Raj Jain

9-18

Schedule



- November 1996: Proposal cutoff
- January 1997: First draft
- March 1997: Second draft
- July 1997: Working Group Ballot
- March 1998: Approval

The Ohio State University Raj Jain

9-19

Status

- On Schedule
- First draft reviewed in January 97
- Third draft was issued in July'97
- 1000Base-X: Gigabit Ethernet based on Fiber Channel Phy
- Phy modified for 1000 Mbps operation
- Phy modified for ISO 11801 standard for premises cabling ⇒ 550 m intra-building backbone runs
⇒ 1300-nm lasers on 62.5- μ m multimode fiber
850-nm lasers on 62.5- μ m fiber ok for 300 m

The Ohio State University Raj Jain

9-20


1000Base-X

- ❑ 1000Base-LX: 1300-nm laser transceivers
 - 2 to 550 m on 62.5- μm or 50- μm multimode, 2 to 3000 m on 10- μm single-mode
- ❑ 1000Base-SX: 850-nm laser transceivers
 - 2 to 300 m on 62.5- μm , 2 to 550 m on 50- μm . Both multimode.
- ❑ 1000Base-CX: Short-haul copper jumpers
 - 25 m 2-pair shielded twinax cable in a single room or rack.
 - Uses 8b/10b coding \Rightarrow 1.25 Gbps line rate

1000Base-T

- ❑ 100 m on 4-pair Cat-5 UTP
 - Network diameter of 200 m
 - Requires new coding schemes
 - Under development.
 - New PAR approved in March 1997
 - 802.3ab task force

Summary



- Ethernet will run at 1000 Mbps
- Will compete with ATM for campus backbone and desktop
- Both shared and full-duplex links
- Fully compatible with current Ethernet

The Ohio State University

Raj Jain

9-23

References

- For a detailed list of references, see http://www.cis.ohio-state.edu/~jain/refs/gbe_refs.htm
- "Media Access Control (MAC) Parameters, Physical Layer Repeater and Management Parameters for 1000 Mb/s Operation," IEEE Draft P802.3z/D3, June 20, 1997.

The Ohio State University

Raj Jain

9-24

References (Cont)

- Email Reflector:
stds-802-3-hssg@mail.ieee.org
 - To join send email to majordomo@mail.ieee.org
 - subscribe stds-802-3-hssg@mail.ieee.org <your email address>
- FTP Site:
ftp://stdsbbs.ieee.org/pub/802_main/802.3/gigabit
- Gigabit Ethernet Consortium
<http://www.gigabit-ethernet.org>

Gigabit Ethernet Products

- **Shipping:** Acclaim Communications, Cabletron, NBase
- **Announced:** 3Com, Alteon Networks, Foundry Networks, Packet Engines, XLNT Designs, Xylan, HP
- **Planning:** Bay Networks, Cisco, Compaq, Digital, FORE, Extreme Networks, IBM, Intel, Ipsilon, Madge Networks, Neo Networks, Plaintree, Prominet, Rapid City, Sun, YAGO Systems
- Ref: Network World, March 17, 1997

Current Schedule

7/17/97 Priority and Multicasting on LANs

7/22/97 **No Class**

7/24/97 Virtual LANs

7/29/97 Gigabit Ethernet

7/31/97 Quiz 2 (No MBone transmission)

8/5/97 Residential broadband: Cable Modems, xDSL

8/7/97 Multimedia: Compression Standards

8/12/97 Multimedia over IP: RSVP, RTP

8/14/97 Wireless LANs and WANs

8/19/97 Quiz 3 (No MBone transmission)

The Ohio State University

Raj Jain

9-27

Credits

This MBone transmission was made possible by:

- Mark Fullmer, OSU/UTS
- Mike Iverson, OSU/UTS
- Mike Douglas, OSU/UTS
- Jayaraman Iyer, OSU/CIS
- Sohail Munir, OSU/CIS

The Ohio State University

Raj Jain

9-28