IBM LAN Architectures and Standards (91770):

Reference Guide - 1st Edition
February 1990

IBM Product Education - Telecommunications
5601 Executive Drive
Irving, Texas 75038
External (214) 550-4000
Internal Tie Line 892-4000

This publication was produced using the IBM Document
Composition Facility (program number 5748-XX9 and
the master was printed on the 3820.

First Edition (February 1990)

Changes are periodically made to the information herein; before using this publication in connection with the operation of IBM systems, consult the latest IBM Program Library reference material.

It is possible that this material may contain reference to, or information about, IBM products (machines and programs), programming, or services that are not available in your country. Such references or information must not be construed to mean that IBM intends to announce such products in your country.

Note: Due to the nature of this course, the visuals used are a mixture of conceptual images and technical drawings. While every effort has been made to create clear, concise, easy-to-read visuals for both the instruction monitor and the student guide, this is not always possible. In those cases where the technical integrity of the visual would be affected by simplification, the instruction staff has decided to use the visuals without modification. This will allow the student to have an accurate representation of the material in the student guide for future reference.

Contents

UNIT 01: Course Introduction and Overview

Session 1 - Introduction

Course Introduction

This course will cover the details of the LAN architectures, standards and protocols currently being implemented by IBM products. LAN media access methods will be discussed along with the common logical link control functions and bridging implementations. The course will also address how three different networking protocols (i.e. NETBIOS, TCP/IP and SNA) are implemented by IBM on LANs. A brief discussion on the proposed FDDI standard will be included.

THIS COURSE IS A PREREQUISITE COURSE FOR ALL DOS AND LAN OS/2 IMPLEMENTATION AND LAN MANAGEMENT EDUCATION (E.G. 91756, 91771).

Audience

This course is for Systems Engineers (SE) and IS Telecommunications personnel who need an understanding of IBM's implementation of LAN architectures and standards or who are planning on attending a LAN implementation, tuning or management course.

Prerequisites

Knowledge of telecommunications (data networking) fundamentals attained either through formal education or equivalent experience.

Although not a prerequisite, familiarity with other architectures such as SNA will be very helpful to the student attending this class.

Objectives

Upon completion of this course, the student should be able to:

- Identify the types of media used for standardized LAN implementations and their general characteristics
- Describe the three LAN media access control (MAC) methods standardized by the IEEE
- Interpret the meaning of all fields as implemented by IBM in the following IEEE standards:
  - 802.5 Token-Ring Media Access Control
  - 802.3 CSMA/CD Media Access Control
  - 802.2 Logical Link Control
- Briefly describe the media access control and physical layer implementations proposed for use with Fiber Distributed Data Interface (FDDI) networks
- State the differences between IEEE LAN implementations and Ethernet LAN implementations

- Describe how NETBIOS provides transparent network support for applications running on the LAN
- Briefly describe the function of each major protocol used in the TCP/IP protocol suite
- Describe how TCP/IP provides support for applications running on a LAN
- Briefly describe how SNA can be implemented in a LAN environment to support the following:
  - Hierarchical networking (i.e. VTAM controlled networking)
  - Advanced Peer to Peer Networking (APPN)
  - Advanced Program to Program Communication (APPC)
- Compare the functions in IBM LAN implementations with the seven layer OSI network model
- Describe bridging techniques used to interconnect LANs and state the difference between source routing and transparent bridging
- Given a specific LAN configuration (hardware and software), state whether the potential for interoperability exists.

Course Outline

The material for this course will be presented in the following units:

- Unit 1 - Course Introduction, Definitions, Media, OSI Network Model
- Unit 2 - Media Access Control and Physical Layer Standards
- Unit 3 - Logical Link Control Standards
- Unit 4 - NETBIOS and Server Message Blocks
- Unit 5 - SNA on Local Area Networks
- Unit 6 - TCP/IP LAN Implementations
- Unit 7 - LAN Interconnection Architectures

References

- IBM Publications
  - *IBM Token-Ring Network Architecture Reference* (SC30-3374)
  - *IBM Local Area Network Technical Reference* (SC30-3383)
  - *TCP/IP Tutorial and Technical Overview* (GG24-3376)
  - *IBM Token-Ring Network Administrator's Guide* (GA27-3748)
  - *IBM Token-Ring Network Telephone Twisted-Pair Media Guide* (GA27-3714)
  - *IBM Token-Ring Network Optical Fiber Cable Options* (GA27-3747)
  - *IBM Cabling System Planning and Installation Guide* (GA27-3361)
  - *A Building Planning Guide for Communication Wiring* (G320-8059)
  - *Internetworking With TCP/IP; Principles, Protocols and Architectures* Douglas Comer; 1988, Prentice Hall, Englewood Cliffs, New Jersey 07632 (SC09-1302)
- Non-IBM Publications
  - *Local Area Networks; architectures and implementations* James Martin with Kathleen Kavanagh Chapman; 1989, Prentice Hall, Englewood Cliffs, New Jersey 07632
  - *Data Communications* magazine (McGraw Hill)
  - *LAN Magazine* (Telecom Library)

Trademarks

- AIX is a registered trademark of IBM Corporation
- DECnet is a registered trademark of Digital Equipment Corporation
- ETHERNET is a registered trademark of Xerox Inc.
- Intel is a registered trademark of Intel Corporation
- NFS is a trademark of SUN Microsystems, Inc.
- UNIX is a registered trademark of AT&T
- VAX is a registered trademark of Digital Equipment Corporation
- AppleTalk and MacIntosh are registered trademarks of Apple Computer Corporation
- Xerox Network Systems (XNS) is a registered trademark of Xerox Corporation
- NetWare, IPX and MINDP are registered trademarks of Novell Inc.
- 3+Share, 3+Mail, 3+3270, 3+Remote, 3+NetConnect, 3+Route and 3+ are registered trademarks of 3COM Corporation
- ProNET is a registered trademark of Proteon Corporation
- StarLan is a registered trademark of AT&T

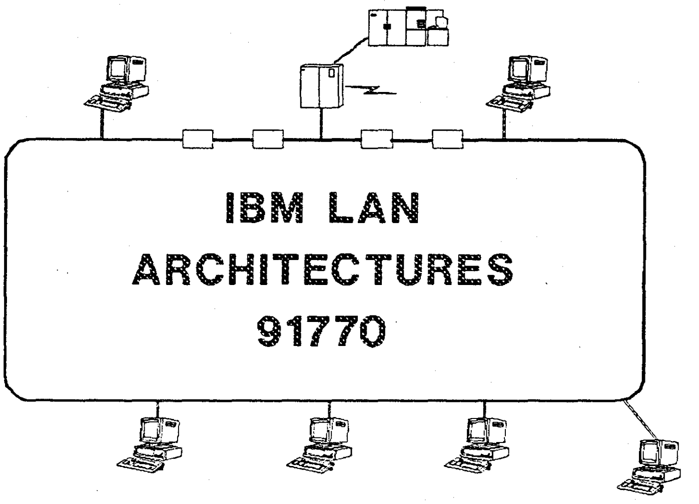

IBM LAN ARCHITECTURES 91770

Who, When, etc.

- Instructor(s): ______________________________
- PROFS (VM) ID & Node: ______________________________
- Telephone Number: ______________________________
- Class dates: ____________________

What This Course Is All About ...

This course will cover the common building blocks for most IBM LAN product implementations. These common building blocks are based on a combination of international standards and industry accepted de facto standards (e.g., NETBIOS, SNA, XNS, etc.).

The course will address the following areas:

- Media
- Media access control
- Logical link control
- Networking protocols

Although the emphasis will be on IBM′s LAN implementations, other vendor implementations will be acknowledged.

This course is not intended to address specific products. Specific product information is contained in the LAN Concepts and Products course (91739). It is the intent of this course to provide you with the generic LAN implementations on which the products rely on to participate in this environment.

Premise For This Class

This class was designed with one major premise in mind:

Knowledge of the underlying architectures will allow you, the student, to make statements relative to product interoperability in a LAN environment.

In order words, if you connect product ′A′ to product ′B′, will it work? Let′s take a look at a few possible scenarios (these scenarios are not in the student guide and will be presented by the instructors).

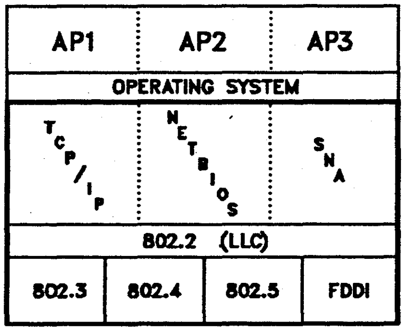

LANA235

Definitions and LAN Concepts

Definitions

Medium

- In this course, the term relates to 'a physical entity used to transport an information signal'
- In other words twisted pair copper wire, coaxial cable and fiber optic cable
- 'Media' is the plural of medium; represents the combination of all entities used for information transport

Protocol

- "A set of semantic and syntactic rules that determines the behavior of functional units in achieving communications"[1]
- A set standard on how the electronic signals will be interpreted

Server

- A data station that provides services to other data stations on the LAN
- Examples include file servers, communications servers and print servers

Baseband and Broadband Transmission

- Baseband
  - Transmission can use the entire bandwidth of medium

[1] Definition from *Dictionary of Computing*, IBM, March 1987.

- Only one user at a time transmits on the medium
- Usually digital signalling is used although analog signalling is an alternative (e.g., modems)
- Medium is shared by multiple users through time sharing or time division multiplexing (TDM)

- Broadband
  - Bandwidth of medium is divided into multiple channels, each using a different range of frequencies
  - Multiple users use medium at the same time
  - Only users operating within the same range of frequencies can communicate with each other
  - Sharing technique is called frequency division multiplexing (FDM)
  - Signalling must be analog

Bridges, Routers and Gateways

There are several ways to interconnect networks together; the specific function used to make this connection will depend on the types of network being interconnected and how many fields need to be altered when going from one network to the next. Bridges, routers and gateways all interconnect networks; however, each does so differently.

We will discuss network layering later in this unit. For now, it will only be necessary to point out that most networks are implemented in a layered fashion, and each of these three functions are implemented at a different network level.

- Bridges
  - Bridges connect networks, or more correctly segments of the same network, at a very low functional level
  - Portions of the link header and trailer may need to be changed, but most of the data in the frame is left unaltered
  - Bridges are the most efficient of the three since they usually require the least amount of processing
- Routers
  - Routers are used to route frames (or packets as they are also called) from one network to another
  - Routing requires changing one or more addresses, making decisions as to the best route or path to take through the interconnected networks and possibly changing link headers and trailers
  - Routers are usually less efficient than bridges but more efficient than gateways
- Gateways
  - Gateways are used to change protocols from one set to another
  - This usually means removing the information from all protocol envelopes and repackaging it in another set of protocol envelopes
  - Gateways are the least efficient way to interconnect networks, but allow the interconnection of networks with different protocol implementations (e.g., SNA and NETBIOS)

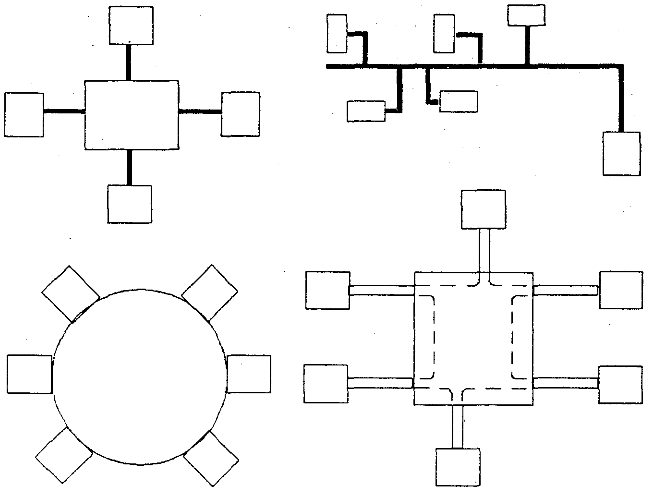

Topology

- TOPOLOGY refers to the physical layout of the LAN's medium used to connect the attached devices
- There are three predominant topologies used for LANs:
  - Star topology
  - Bus or tree topology
  - Ring topology
- Star: All devices connect to a central control hub which allows devices to communicate directly with each other (through the hub).
- Bus: A main cable provides a common communications path; devices are attached to the common cable (called cable taps) and can potentially communicate with any other device on the cable.
  - Tree: A variation of the bus where the physical cable distances are extended using amplifiers and signal splitters; operates the same as a bus.
- Ring: Devices are connected together so that each device is attached to adjacent devices only; data transmission is unidirectional.

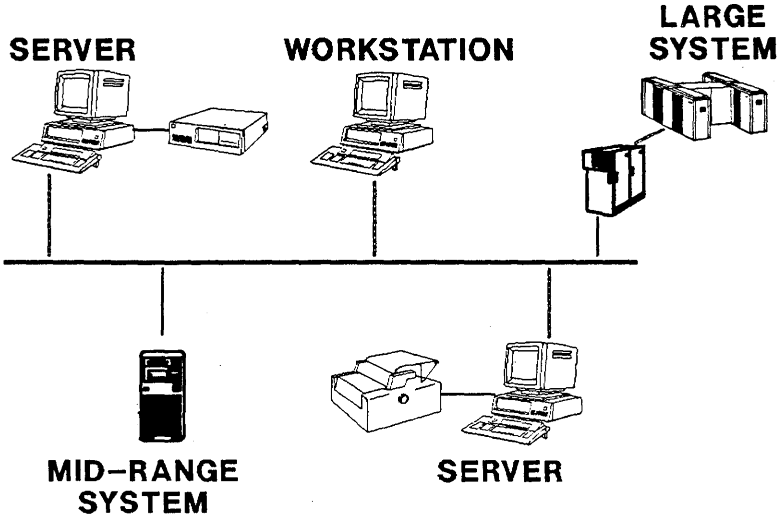

LAN Overview

The following overview provides a description of local area networks, LAN uses, benefits, and layouts.

Characteristics of LANs

- Local Area Networks are a way of connecting devices in order to:
  - Allow more efficient computing
  - Share computing resources
  - Create a new communications network implementation which allows new applications of computers
- LANs connect intelligent (programmable) devices together by cable so that they can communicate with each other
- Almost any electronic device has the potential to be part of a LAN using either a direct attachment or through some LAN attached device:
  - Personal Computers
  - Mini-computers
  - Mainframe computers
  - Facsimile machines
  - Disk drives
  - Printers
  - Word processors
- Devices on a LAN must use compatible data stream formats and communication protocols in order to communicate with each other
  - LANs provide physical connectivity for devices, but do not guarantee communications among the attached devices at the application level

  - For example, for two SNA applications to communicate, both applications must use the same SNA implementation; the SNA layers are above the common LAN protocols

- Devices using different data stream formats can share the same LAN, but will be unable to communicate with other devices on the LAN

  - For example, an SNA application can not communicate with a TCP/IP application but both can use the same common LAN protocols

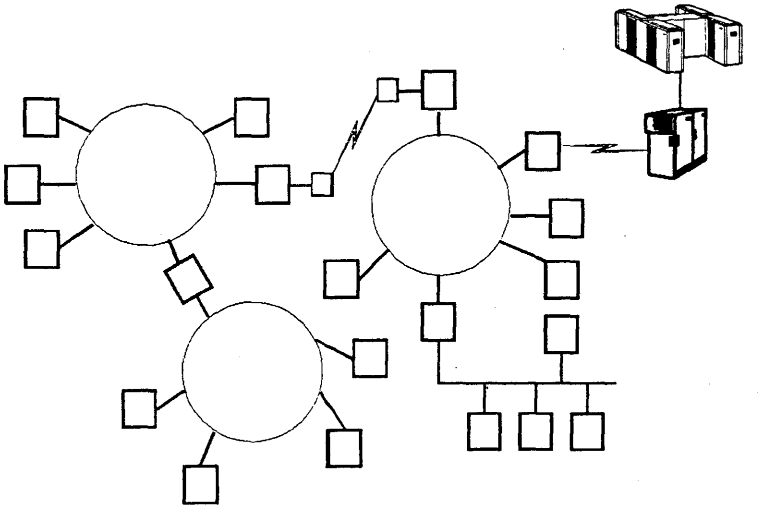

Interconnecting LANs

Note: Details of LAN interconnections will be covered in a later unit of this course.

- Similar LANs can be linked to each other by means of a BRIDGE
  - Similar LANs may use different IEEE MAC sublayer and physical layer implementations (e.g., IEEE 802.5 or 802.3), but the same LLC sublayer protocols (e.g., IEEE 802.2)
  - The bridge usually connects the LANs at the 802.2 LLC sublayer of the Data Link Control layer
- Multiple networks are interconnected using a ROUTER
  - Network interconnection is accomplished using OSI layer 3 functions
- Dissimilar networks (i.e., a LAN and a WAN, or different types of LANs) can be connected through a GATEWAY
  - Gateways provide speed, protocol, and/or code conversion
  - A gateway acts as a translator, allowing messages to pass between different types of networks
  - A gateway may be required if the LAN data is transmitted over a WAN to a host system

Session 2 - LAN Media

Objectives

Upon completion of this session you should be able to:

- Identify the types of media that can be used in a LAN
- Describe the general characteristics of each type of medium
- Identify the advantages and disadvantages of each type of medium as they relate to LAN installations.

Introduction

In this session, the characteristics of the different types of wire and fiber will be reviewed, along with the common terminology used when discussing these mediums.

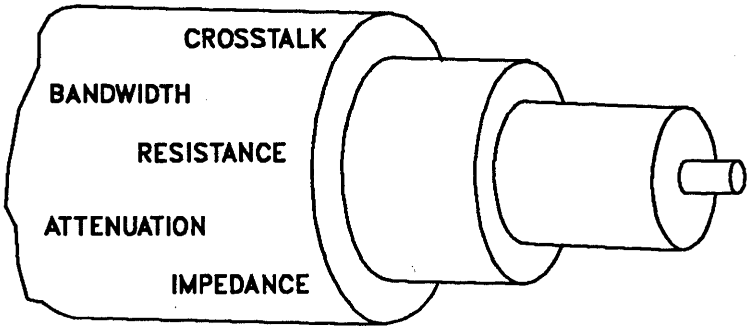

LANA101

Media (Wire) Characteristics

DC Resistance

- The opposition in an electrical circuit to DC current
- Stated in 'ohms'
- Determined by the characteristics of the wire (primarily the size and conductor type)
- Resistance causes gradual signal loss of a DC or low frequency AC signal as the signal travels along the wire

Characteristic Impedance

- The opposition in an electrical circuit to AC current flow
- Stated as 'ohms at a particular frequency'
- Inversely related to frequency (characteristic impedance goes down as frequency goes up)
- Determined by the characteristics of the wire (wire size, construction, insulation size and type, etc.)
- The cable must be terminated with a load which matches the characteristic impedance of the cable

  Note: Specific types of cable will sometimes be referred to by their characteristic impedance, such as 50 ohm coax cable.

Attenuation

- A decrease in the magnitude of current, voltage or power of a signal as it travels along a wire
- Stated in 'dB/unit length at a specific frequency'
- Attenuation characteristics are related to the resistance and impedance of the cable

Crosstalk

- The unwanted energy induced into one wire pair from another wire pair in close proximity
- An example would be crosstalk caused by telephone ring signals
- Crosstalk can be reduced by keeping signals on wires low, twisting wire pairs and/or shielding the cable

Continuity

- DC continuity
  - Continuous (end to end) low resistance DC electrical path
- AC continuity
  - Continuous (end to end) low loss AC electrical path
  - Measures the ability of the wire to successfully transport an AC signal at a given frequency
  - Related to the characteristic impedance, bandwidth and integrity of the wires and shield in the cable
- DC continuity does not equate to AC continuity

Bandwidth

- A range of frequencies which can be transmitted over a medium without significant loss of power (attenuation)
- Determined by the physical properties and construction of the media
  - Unshielded twisted pair cable (usually called telephone cable) - low to medium bandwidth (up to several MHZ)
  - Shielded twisted pair cable - medium bandwidth (up to about 20 MHZ)
  - coax cable - medium to high bandwidth (up to 300 MHZ)
  - fiber optic cable - high to very high bandwidth (up to several GHZ)
- Bandwidth is primarily related to the attenuation characteristics of the cable
- As the bandwidth of the cable increases, so does the price (usually)

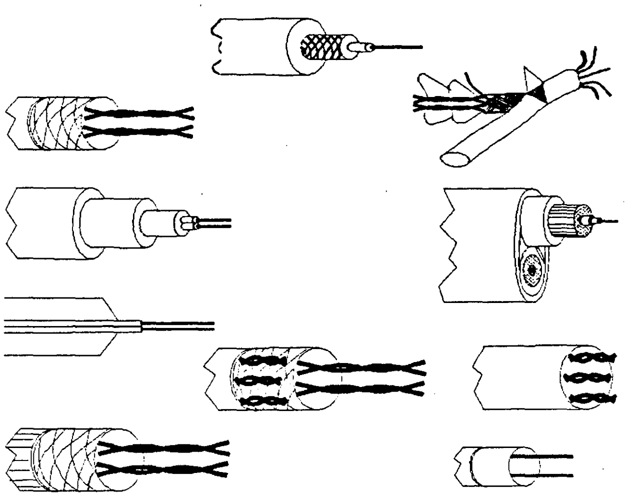

Specific Types of Media

75 Ohm Coaxial Cable

- Often referred to as 'broadband coaxial cable'
- Construction:
  - Single conductor with shield
  - Signal return path is cable shield
- 75 ohm characteristic impedance
- Type RG 59/U, RG 11/U, RG 6/U, etc.
- Large bandwidth (usually 300 MHZ)
- Commonly used for broadband LAN implementations

50 Ohm Coaxial Cable

- Primarily used for Ethernet baseband implementations
- Construction:
  - Single conductor with shield
  - Signal return path is cable shield
- 50 ohm characteristic impedance
- Type RG 58/U and others
- Large bandwidth (usually 300 MHZ)
- Comes in several different diameters

- 0.395″ diameter, sometimes called 'thick Ethernet', used for the main cable path (backbone)
- RG 58/U, sometimes called 'thin Ethernet' or 'cheapernet', used for connection of network nodes to the backbone cable through the use of taps

  Note: More details of the Ethernet connections will be covered in a later Unit

- 0.405″ and 0.375″ dia. as specified by IEEE 802.3 standard

93 Ohm Coaxial Cable

- Often referred to as '3270 coaxial cable'
- Construction:
  - Single conductor with shield
  - Signal return path is cable shield
- 93 ohm characteristic impedance
- Type RG 62/U
- Medium bandwidth (supports 3270 2.3 Mbps data rate)
- Conductor size, insulation and shield is different from the 75 ohm coax cable; they are NOT interchangeable
- This cable is not used on any IBM LAN implementations but is used for a few vendor LAN implementations (e.g., ARCNET)

Twinaxial Cable

- Construction:
  - Two parallel conductors with shield
  - Shield is not used as part of signal path
- 110 ohm characteristic impedance
- Medium bandwidth
- This cable is not used for any current LAN implementation

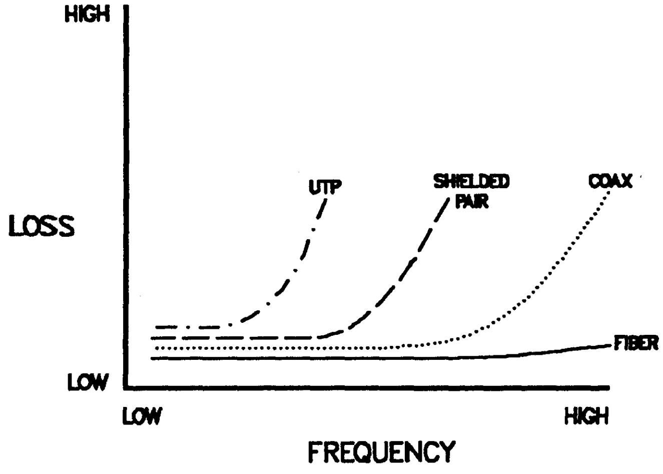

LANA104

Unshielded Twisted Pair Cable

- Used for voice applications and some data (LAN) applications
- General construction:
  - One or more pairs of conductors in same sheath
  - Each wire pair has a certain number of twists per foot
- Characteristic impedance will range from 600 ohms at low frequencies (i.e. voice) to about 100 ohms for higher frequencies (> 100KHZ)
- Attenuation characteristics at higher frequencies impose limits on the length of cable that can be used in any given configuration
  - Length limitations can sometimes be addressed through the use of repeaters at various points in the network (e.g., at both ends of a cable segment)
- Additional limitations on the use of this cable result from the fact that it is not shielded, and will therefore emit radiated energy
- If the signal level on the cable is too high, can violate FCC rules governing limits of radiate energy
- Cable is highly susceptible to induced noise
- Referred to as UTP or TTP (telephone twisted pair) cable

  Note: The cable discussed here is packaged in a variety of ways (2 pair, 25 pair, 50 pair, etc.).

  Note: The small, flat telephone extension cable, commonly found in 25 and 50 foot lengths, should NEVER be used for LAN node attachments. The cable has a much higher loss characteristic than UTP, and is NOT twisted.

Shielded Twisted Pair Cable

- Used for LAN applications
- General construction:
  - Usually two pair of wires
  - Shield may be common for both pair, or each pair may be individually shielded
  - Each wire pair has a certain number of twists per foot
- Characteristic impedance will be around 100 ohms at higher frequencies
- Although not as severe as UTP, attenuation characteristics at higher frequencies will impose limits on total cable lengths
  - Length limitations can sometimes be addressed through the use of repeaters at various points in the network
- Shielding provides the cable with high immunity to noise and almost eliminates the effects of radiated energy

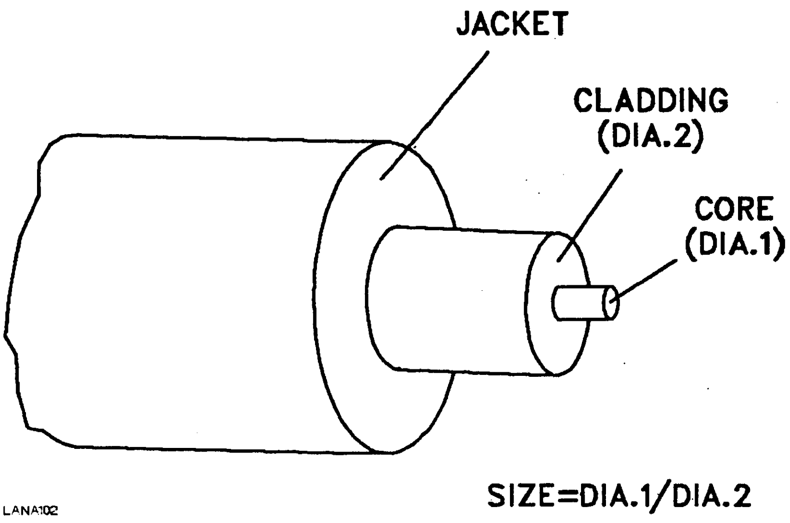

Fiber Optic Cable

Fiber optic cable provides many advantages over other types of media (i.e. wire cables). Some of these advantages include:

- Security
- Immunity from the effects of noise and radiated energy
- Very high bandwidth

Before we take a look at the characteristics of fiber, let's take a brief look at how information transmission across fiber works.

- Transmission medium is fiber optic cable which is designed to transport light energy as opposed to electrical energy
- Data presented at the input of the transmitter is used to modulate a light source (on and off)
- At the receiver, a light sensitive device looks at the arriving light signals, and converts them back to data
- Signal on the fiber is unidirectional digital baseband

Characteristics of Fiber Optic Cable

- Light-transporting material (core) is either glass or plastic
- Cladding material encloses glass/plastic core
- Strength members and protection layers are built around core and cladding to give the cable the strength required to withstand handling
- The core is classified by 'mode', which represents how the light is conducted through the core

  - multi-mode graded index
  - multi-mode stepped index
  - single mode

- Size of fiber:
  - Fiber is identified by two numbers; one representing the diameter of the core and the other representing the diameter of the cladding in micrometers
    - A micrometer is equal to 39.37 millionths of an inch
  - Although there are a number of different sizes, the most common are:
    - 62.5/125 (used as the AT&T Technologies standard)
    - 50/125 and 85/125 (defined by ANSI along with 62.5/125)
    - 100/140 (IBM Cabling System Type 5 cable)

- Light sources and wavelengths
  - Two most common light sources used are Laser and light emitting diodes (LEDs)
    - A Laser is a device that produces a very narrow, intense beam of coherent light
    - Lasers emit light that has a narrower angle of emission and less pulse spreading than LED's; this means that laser light can be transmitted for longer distances before regeneration of the signal is required
    - An LED is a device consisting of layers of positive and negative material which is electrically stimulated to produce a coherent light output
    - LEDs are more stable, more reliable, and less expensive than laser sources, but LEDs exhibit a wider angle of emission and have more pulse spreading, reducing the bandwidth capacity (limits there use primarily to short (< 2KM) distances)
  - Light sources are also identified by their wavelength; this wavelength must be matched to the wavelength of the fiber (usually 850 or 1300 nanometers)

- Wavelength
  - coherent light sources (i.e. lasers and LEDs) produce light at a specific frequency; the distance traveled by one cycle of light at the specific frequency is called the wavelength
  - the core material of the fiber is designed to conduct light most efficiently at a specific wavelength (analogous to the characteristic impedance of wire)

- Attenuation
  - A light signal suffers a loss (attenuation) as it travels through the fiber the same as an electrical signal traveling in a wire
    - loss is measured in dBs per length (usually in Kilometers)
  - Loss also occurs at each connecter (usually 0.1 to 2 dB)
  - Overall loss should be limited to 16 dB when using fiber
    - To keep the overall loss within tolerance, fiber optic repeaters can be used when transmitting over longer distances.

Session 3 - Standards and the ISO OSI Model

Objectives

Upon completion of this session you should be able to:

- Identify the standards organizations involved with local area networks
- Briefly describe the purpose of the ISO Open Systems Interconnect (OSI) model
- Briefly describe the functions of each OSI layer.

Standards

- Communication standards, including those for local area networks (LAN), identify common procedures and protocols that allow products to communicate with each other
- LAN standards are needed for two reasons:
  1. To establish a hardware-independent framework for network design
  2. To establish standards for software protocols and hardware interfaces to allow different vendor products or offerings to communicate with each other
- Although a number of organizations are involved in standards work either directly or indirectly related to local area networks, several are predominant:
  - ISO - International Standards Organization
  - IEEE - Institute of Electrical and Electronic Engineers
  - ANSI - American National Standards Institute
  - CCITT - International Telegraph and Telephone Consultative Committee

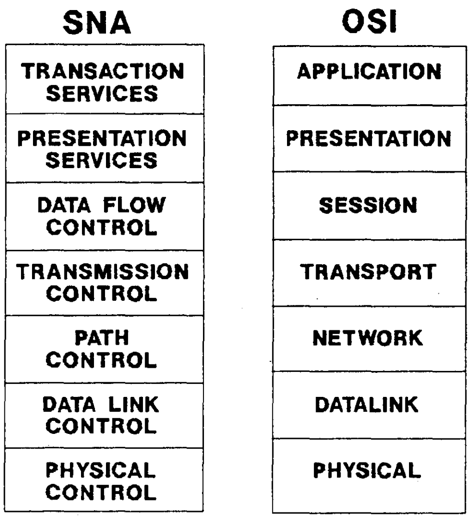

SNA	OSI
TRANSACTION SERVICES	APPLICATION
PRESENTATION SERVICES	PRESENTATION
DATA FLOW CONTROL	SESSION
TRANSMISSION CONTROL	TRANSPORT
PATH CONTROL	NETWORK
DATA LINK CONTROL	DATALINK
PHYSICAL CONTROL	PHYSICAL

ISO-OSI Reference Model

- Developed by the ISO, the Open System Interconnect (OSI) reference model specifies a standard set of functions for exchanging information among different vendor network products
- The OSI model establishes terminology, identifies the primary components of a system and component functions, and defines the relationship of the functions to each other
- The OSI model is a layered architecture that allows interconnection of different vendor products
- The OSI model was designed primarily for wide area networks (WANs) but can be (and is) used for local area networks
- The OSI model organizes all system networking functions into seven (7) separate and independent layers

Overview of OSI Layers

- Physical
  - Physical transmission of the signal including coding and electrical and mechanical specifications, and activation/deactivation of the physical link
- Data Link
  - Protocols and functions for transferring information across a link
  - Services can be reliable (connection-oriented) or best effort basis (connectionless)
  - Includes link level flow control and error recovery as appropriate

- Network
  - Support for internetworking and routing functions between networks
  - Networking services transparent to the upper layer functions
- Transport
  - Reliable end-to-end data transfer, including flow control and error recovery
  - End-to-end transport transparent to upper layers
- Session
  - Establishes, maintains and terminates connections between applications
  - Controls data flow between applications
  - May provide name to address resolution
- Presentation
  - Standardized application interface
  - Data presentation services, including code translation and datastream conversion, encryption, and compression
- Application
  - User interface
  - Services provided can include file transfer protocols, transaction server functions and network management

Feedback Questions

Directions:

First select the correct answer for each of the following questions by circling the appropriate letter. In a few minutes, you will be asked to input your answers, one at a time, using the Student Response Unit.

1. (True/False) CATV cable is '75 ohm' coaxial cable; this means the characteristic impedance of the cable is about 75 ohm.

   [✓ a.] True

   b. False

2. Which of the following statements is true relative to the use of unshielded twisted pair cable for LAN communications:

   a. Can be used as a direct replacement for shielded twisted pair cable

   b. Can not be used for LAN communications

   c. Is an excellent substitute for coaxial cable because of its price

   [✓ d.] Can be used as an alternative to shielded cable, but with distance limitations

3. The OSI Transport layer (4) provides:

   [✓ a.] Reliable end-to-end data transfer

   b. A standardized application interface

   c. Reliable link level data transfer

   d. Name to address resolution

UNIT 02: MAC and Physical Layer Standards

Session 1 - Overview of the IEEE Standards

Introduction

This unit is divided into 7 sessions. The first session will discuss the IEEE standards in general terms. Each specific IEEE LAN standard will then be covered in a separate session. The IEEE standards define the common protocols and media access methods used on most vendor LAN implementations. Although not an IEEE standard, Ethernet version 2.0 will also be covered because of its wide acceptance in the marketplace. This unit will end with a brief discussion of a new standard still under development by ANSI: Fiber Distributed Data Interface, or FDDI for short.

Objectives

Upon completion of this session, you should be able to:

- State the area of responsibility for each IEEE 802 committee as it relates to LANs.

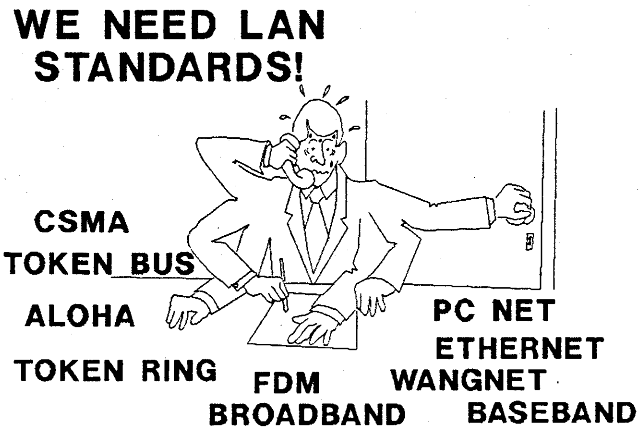

IEEE Project 802

- IEEE Project 802 was established in February, 1980
  - Purpose: to standardize Local Area Networks
- December, 1980, tentatively accepted two protocols for standardization and named two topologies for implementation. This gives us three LAN standard configuration possibilities:
  - CSMA/CD using Bus Topology
  - Token Passing using Bus Topology
  - Token Passing using Ring Topology
- Project 802 has been divided into five (5) separate standard committees:
  - IEEE 802.1 - Documentation and overview
  - IEEE 802.2 - Logical Link Control
  - IEEE 802.3 - CSMA/CD
  - IEEE 802.4 - Token Bus
  - IEEE 802.5 - Token-Ring

IEEE 802.1 - Documentation and Overview of LAN standards

- Coordination of the committees is done by this group. Also, LAN standards must reflect other network standards. Members of the following groups provide input to LAN committees:
  - ISO - International Standards Organization
  - ECMA - European Computers Manufacturing Association
  - ANSI - American National Standards Institute
  - CCITT - International Telegraph & Telephone Consultative Committee

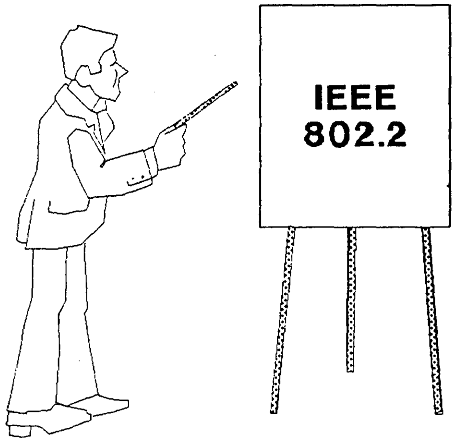

IEEE 802.2 - Logical Link Control

- One of two sublayers (the other being the MAC sublayer) within the OSI Data Link Control layer
- Responsible for defining a common logical link protocol
  - LLC protocol defines the "packet" used to transport data within a LAN independent of the type of LAN being used
  - LLC "packet" includes the data which will be visible to the higher OSI functional levels (e.g. SNA, TCP/IP, NETBIOS)
  - The LLC "packet" is called a LLC Protocol Data Unit (L-PDU)

  Note: Details of the L-PDU will be covered in the next Unit.
- The MAC frame has two (2) formats and can be implemented on 2 topologies resulting in three IEEE MAC standards:
  - CSMA/CD, bus topology (802.3)
  - Token passing, bus topology (802.4)
  - Token passing, ring topology (802.5)
- The L-PDU is contained within all of the three Media Access Control (MAC) frames defined by the IEEE
- The MAC frame is used to pass data across the physical network

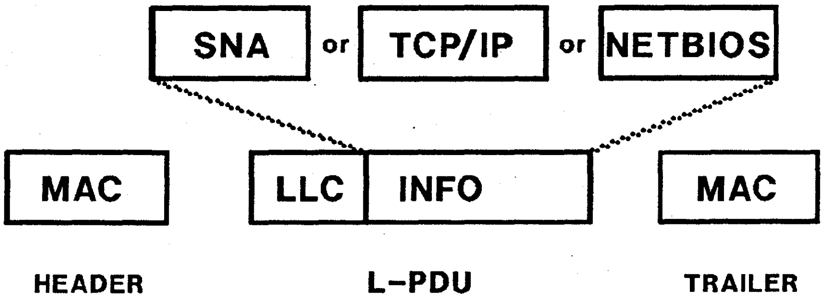

LANA230

MAC and LLC Transparency

- When data is passed from the upper layers to the LLC sublayer, control information will be included to indicate the destination of the data
- LLC sublayer does not look at the data, only the control information
- LLC assembles the appropriate header and passes the LLC frame containing the imbedded data to the MAC sublayer, again including control information
- MAC sublayer does not look at the LLC frame, only the control information (e.g., destination and source addresses)
- MAC sublayer assembles the MAC frame and sends the information using the appropriate physical layer protocols
- Because of the way that the LLC and MAC frames are assembled, they are independent of the data content, thereby providing transparent network services to the upper layers of the machine

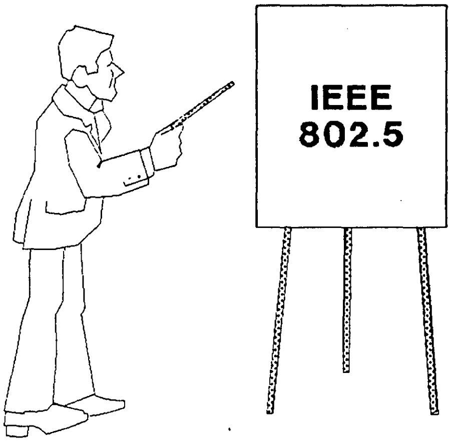

Session 2 - IEEE 802.5 Token-Ring Standard

Objectives

Upon completion of this session, you should be able to:

- Briefly describe the purpose and operation of an MAU
- Describe the operation of the IBM Token-Ring at 4Mbps and 16Mbps
- Explain the types of physical addressing used on an IBM Token-Ring
- Interpret the meaning of the bits in the IEEE 802.5 MAC frame as implemented by IBM
- Identify the functions of IEEE 802.5 MAC frames
- Briefly explain the Token-Ring adapter ring insertion process
- List the IBM products that attach to the Token-Ring LAN.

Multi-Station Access Units (MSAU or MAU)

- Device which converts the star wiring topology into a logical ring topology
- Active or passive wiring concentrator which provides for inserting or bypassing stations on the ring
  - Number of stations supported will vary by manufacturer
  - IBM 8228 is a passive device supporting 8 station connections
- Larger rings are built by connecting MAUs serially (RO on one MAU to RI in next MAU)

MAU Operation

1. Empty positions with or without lobe connections plugged in are automatically bypassed
2. Positions with cables installed will be bypassed until the adapter becomes active
3. When an adapter becomes active, ring diagnostics are run between the station and the MAU
4. If the diagnostics are successful, a + 5VDC signal is induced across the 2 wire pairs (referred to as a 'phantom' signal)
5. This signal is used by the MAU to activate the relay that allows the lobe to become part of the ring
6. When the adapter becomes inactive or is unplugged from the MAU, the + 5VDC signal is removed and the MAU drops the relay, bypassing the lobe
7. The MAU also provides a backup path, using either the self-shorting 'RO' and 'RI' plugs (if so equipped) or electronic switching technology, allowing recovery from single open ring segments

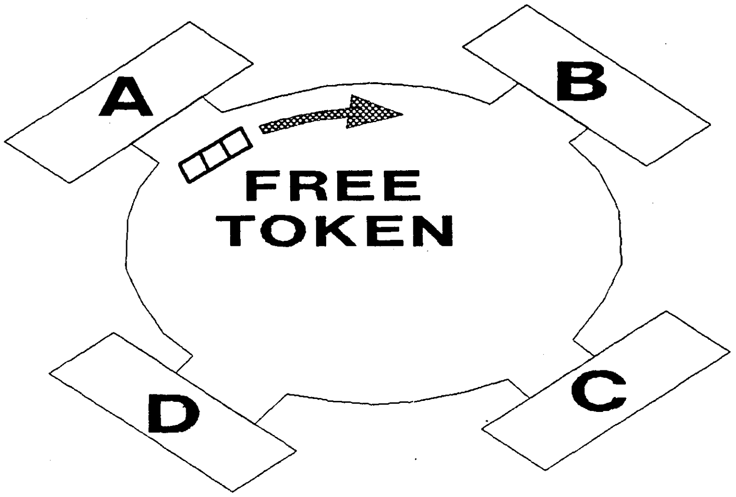

Token Passing Ring Operation

- IEEE 802.5 defines the MAC standard for data transmission using a token passing protocol on a physical ring topology
- Network has a sequential arrangement, with data moving in a single direction, station to station
- Standardized for baseband twisted-pair wire, coax cable or fiber optic cable[2]
- Standard defines the protocol (token passing), media, physical connections, speed and coding scheme

Token-Ring Operation at 4M bps

1. Device ready to send looks for a free token
2. When the free token arrives at the device, the device 'captures' the token by changing the token status to busy and by appending the data (called a data frame, busy token frame or MAC frame)
3. Data frame passes around the ring through each active device
4. When the destination (receiving) device recognizes its address, it copies the data (including the source address and any routing information contained in the frame)

[2] IBM does not support the use of coax cable in a Token-Ring LAN environment.

- Receiving device indicates the data has been received and copied by changing several bits in the frame status field at the end of the frame

5. Data frame continues around the ring until it returns to the sender

6. Sender MUST remove the data from the ring and, in IBM's implementation, MUST issue a free token

7. Basic rules:

   - Sender removes the data; the receiver only copies it
   - There is only one token on the ring at any given time
   - A "monitor" will continuously check the ring for token faults (missing or extra tokens, circulating busy token, etc.)

8. The physical layer (i.e. 802.5) is responsible for delivery of the information to the receiver; any and all error recovery is the responsibility of some higher level function (i.e. LLC sublayer or Transport layer)

16 Mbps Early Token Release

IBM recently announced support for Token-Ring operation at both 4 Mbps and 16 Mbps (on separate rings). In order to improve performance at a ring speed of 16 Mbps, the IEEE 802.5 standard has been enhanced to provide support for early release of the token.

- First device looks for free token, captures it and makes it busy, and appends the data (same as 4 Mbps)

- At the end of the data frame, the sending device appends a free token

- Second device on the ring wanting to send data will pass the data from the first device through the adapter, and capture (make busy) the free token at the end of the data frame

- Second device appends its data frame to the first frame, and appends a free token to the end of its data frame

- Implementation can result in multiple data frames on the ring at any given time, but never more than one free token

- Rule that sending device must remove its own data still applies

- Early token release can be implemented on all IBM '16/4' TRN adapters

FREE TOKEN FORMAT

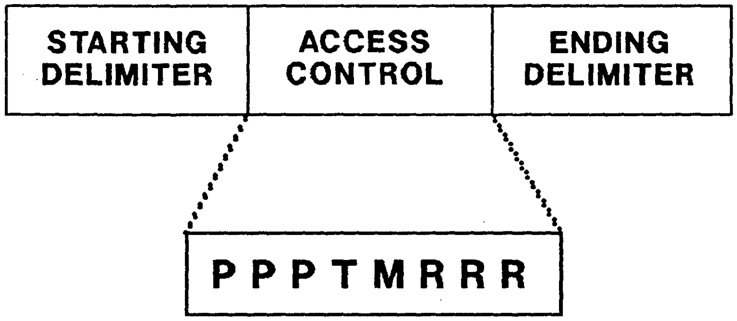

STARTING DELIMITER | ACCESS CONTROL | ENDING DELIMITER

P P P T M R R R

Token-Passing MAC frame formats

There are three different MAC frame formats defined for the token-passing ring implementation:

- Token format - used to pass the free token from one node to another
- Frame format - used for MAC and LLC frames (frames are also called busy tokens)
- Abort format - used during ring recovery

Token Format (Free Token)

Ref: *Token-Ring Network Architecture Reference* SC30-3374

The free token frame consists of three bytes.

- SD is a 1 byte starting delimiter
  - Purpose of the SD byte is to provide synchronization pulses to each adapter to allow the adapter to properly interpret the information following the SD byte
  - SD byte consists of several 0 bits and Manchester code violations (J and K) in a specific pattern
  - Appendix B of the reference manual contains information on the differential Manchester coding used by Token-Ring adapter cards
- AC is a 1 byte access control byte which contains the token indicators and priority bits

- T = token bit
  - 0 indicates a free token
  - 1 indicates a busy token (MAC or data frame)
- M = monitor bit
  - This bit is set to 0 when a source device begins sending a data frame
  - It is changed to a 1 as it passes through the active monitor
  - If the bit is in the '1' state when it comes to the active monitor, the active monitor assumes a failure (e.g. source failed to remove the frame) and will purge the ring; a new free token will then be issued
- P = priority bits (3)
  - Stations can be assigned priorities from 0 to 7 (0 being the lowest) based on their relative importance on the ring
  - For example, a bridge function is usually assigned a priority of 3
  - P bits in the token determine the priority of the circulating token and can be used to control access to the ring; only stations with a priority equal to or higher than the priority in the AC field can capture the token
  - This implementation is used to provide faster ring access to devices requiring more immediate service
  - A station issuing a free token will set the P bits based on the value of the R bits
- R = reservation bits (3)
  - A station with information pending for the ring will attempt to capture a free token
  - If the station sees a data frame (busy token), it will examine the P bits; if the P bits are less than the priority of this station, the station will set its priority in the R bits
  - When the sending device removes its data, it will look at the R bits, and if set to a level higher then the existing token priority, the P bits will be changed to reflect the new priority
  - Only stations on the ring with a priority equal to or greater then the token priority can capture the token
  - When the free prioritized token returns to the station that issued it, the station will release a normal (priority = 0) token

  Note: Priorities are only implemented at 4M bps.

- ED is a 1 byte ending delimiter
  - Purpose of the ED byte is to indicate to the receiving device that this is the end of the frame
  - ED byte consists of several 1 bits, Manchester code violations (J and K) in a specific pattern and an E bit for errors (discussed below)

- **UNIQUE**
  - **UNIVERSAL**
  - **LOCAL**
- **MULTI-STATION**
  - **FUNCTIONAL**
  - **GROUP**
  - **BROADCAST**

LANA225

Data Frame

There are two types of information frames (also called busy token or data frames):

- LLC frames containing an LLC header and may contain user information (data from the upper layers)
- MAC frames containing special information (e.g. network management, configuration changes, etc.) passed between stations on the ring

Ring station addressing

Prior to discussing the data frame, we need to look at the addressing conventions used for token-ring stations.

- IEEE standards allow addresses to be 2 bytes long (local networks only) or 6 bytes long (local or global networks); IBM implements the 6 byte format only
- Individual station addresses can be in one of two formats:
  - Addresses can be ″universally administered″ by the IEEE
  - Addresses can be locally administered by someone other then the IEEE
  - The DA and SA fields, 1st byte, 2nd bit identifies the format being used
    - If the bit is off, the address is universal (also called ″burned in″) and is permanently assigned to that adapter at time of manufacture; if assigned by the IEEE, it will always be unique[3]
    - If the bit is on, the address is local and will start with ′4000........′; local addresses are assigned when the adapter is ′opened′ on the ring (usually through a software command)

[3] Cards currently manufactured by IBM are assigned addresses starting with ′10005A......′.

- In addition to recognizing its unique address, every adapter will recognize broadcast addresses, and can be conditioned to recognize one or more group or functional addresses[4] .
  - Group addresses are used to send messages to multiple stations within the same group (e.g., a server domain)
  - Functional addresses are predefined addresses for specific functions in a LAN (see below)
  - A broadcast address is an all stations address and all active adapters on the ring will copy the frame at the MAC level

Functional Addresses

- Functional addresses are architected to implement network functions on the Token-Ring LAN
- They are required since adapters need to communicate to these functions without knowing the specific (unique) address of the machine in which the function is located
- Functional addresses implemented by IBM include:
  - Active Monitor
  - Ring Parameter Server
  - Ring Error Monitor (REM)
  - Configuration Report Server
  - NETBIOS
  - LAN Manager
  - Bridge
  - User defined
- Functional addresses are turned on by higher level software functions (i.e. layers 3 through 7)

[4] Specific information on the three types of addressing will be found in Chapter 2 of the Architecture Reference manual.

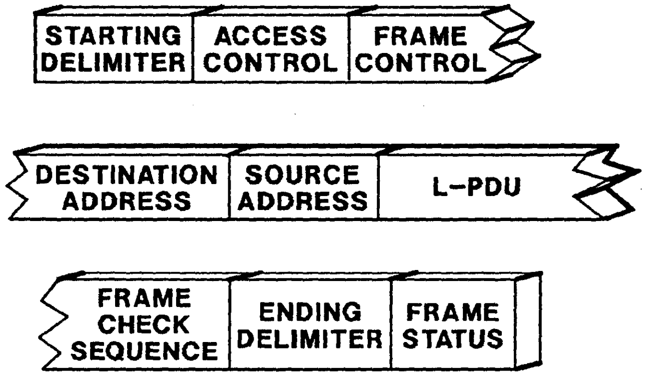

BUSY TOKEN WITH DATA

LANA266

Data Frame

The data frame (MAC or LLC information frame) consists of a minimum of 21 bytes plus the contents of the L-PDU (LLC and user data).

- SD and AC bytes are identical to the free token frame except that the T bit = 1 indicating a busy token
- FC = frame control byte
  - The purpose of this field is to identify this frame as an LLC type frame or a MAC type frame
  - LLC frames will be passed up to the LLC layer after removal of the MAC envelope and will be processed by the LLC function
  - MAC type frames are decoded and handled by the MAC layer functions on the adapter card; they are used for MAC level information transfer between adapter cards

    Note: MAC frames will be discussed after the data frame
- DA is a 2 or 6 byte destination address (target or receiver)

  Note: Two (2) byte addressing will not be discussed here since it is vendor specific and is not supported by IBM.

  - DA can be the target station's unique address, or a group or functional address (multiple receivers)
  - Frames addressed to a unique address can only be copied by the target station or bridges
  - Frames addressed to a group or functional address will be copied by all stations recognizing the address
  - Most of the DA field bits are used for the address; a few are used to indicate address format:

    - Byte 0, bit 0 (I/G bit)
        - Indicates whether the address is an individual or group address
    - Byte 0, bit 1 (U/L bit)
        - Indicates whether the address is universally administered or locally administered
    - Byte 2, bit 0 (FAI bit)
        - Indicates if the address is a functional address (FA)

- SA is a 2 or 6 byte source address (sender)
    - SA must be a unique address; group, functional and broadcast addresses can not appear in the SA field
    - SA will always be the address of the station that originated the frame, not an intermediate station such as a bridge
    - Most of the SA field bits are used for the address; a few are used to indicate address format:
        - Byte 0, bit 0 - Routing Information Indicator (RII)
            - Indicates whether routing information immediately follows the SA field

              Note: This will be discussed in more detail in a later unit.
        - Byte 0, bit 1 (U/L bit)
            - Indicates whether the address is universally administered or locally administered

- The L-PDU field contains logical link control and user data
    - Data field within the L-PDU can be in a variety of formats (e.g. SNA, TCP/IP, NETBIOS or user written)
    - Maximum frame size is determined by the 'token holding timer'
        - Standard allows >4,000 bytes at 4 Mbps; current IBM adapter implementation is about 2,000 bytes
        - Standard allows > 17,000 bytes at 16 Mbps; IBM adapter implementation (16/4 Adapter) is about 18,000 at 16 Mbps and 4,000 at 4 Mbps
        - Actual frame size will be dependent upon amount of available buffer space on the adapter card

- The Frame Check Sequence (FCS) field contains check characters used to validate data integrity (4 bytes)
    - Data integrity is checked between the FC field and the FCS field inclusive

- ED is a 1 byte ending delimiter
    - ED is the same as for the free token format
    - The E bit is used as an error indicator and is set to 0 by the sending station
    - Although only the 'addressed' stations on a ring will copy the data, all stations will check the validity of the data (i.e. FCS check)
    - The first station recognizing an error (e.g., calculated FCS does match the one in the data frame) will turn the E bit on and log an error in the adapter card error log of that station
    - Other stations recognizing the error will check to see if the E bit is on, and if so, will do nothing
    - The sending station will detect the E bit being on when the frame is removed from the ring, and will notify the LLC layer of the condition; the next action is up to the LLC layer (or layers above)

- FS field is a 1 byte frame status field

- When the frame is sent by the sending station, the A and C bits are all set to 0
- The target station (or first station copying a group address message) will turn on the A bits to indicate 'address recognized'
- If the frame is successfully copied, the C bits will likewise be turned on
- When the sending station removes the frame from the ring, the status of the A and C bits are recorded and turned over to the LLC layer; it is the responsibility of the LLC layer or above to take the appropriate action if the bits are not set correctly

Note: When information is copied by a bridge, the A and C bits in the original frame will be set; ; therefore, the sending station knows that another station copied the frame, but does not know if it was the target station.

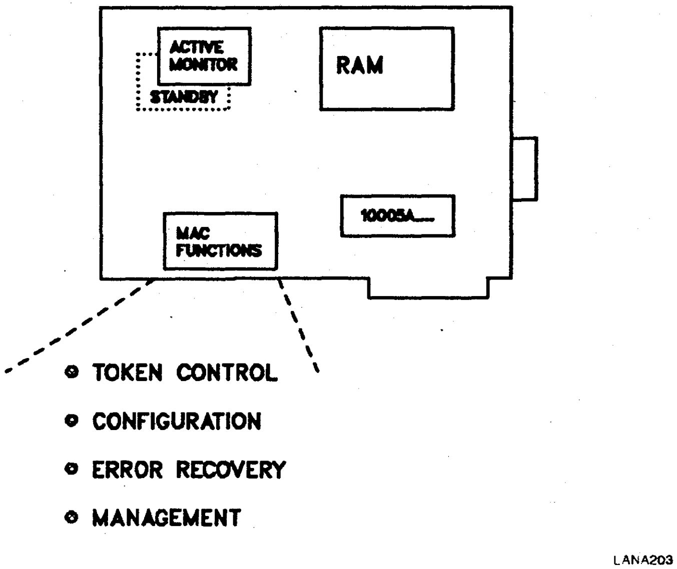

MAC Type Frames

As noted earlier, the FC field in the MAC frame indicates the frame type: either LLC or MAC. Please note the opportunity for confusion here. All busy token frames can be referred to as 'MAC frames'. However, the bit configuration in the FC field determines the 'type' of MAC frame. Why do you even need this special MAC type frames?

Before we answer that, a brief word about LLC type frames. LLC frames are used to pass LLC envelopes around the ring. These L-PDUs as they are called may contain LLC control information and nothing more. They may also contain control information to be processed at layers 3 through 5 (e.g., session setup) or user (application) information. The contents of the L-PDU are transparent to the MAC layer.

In a WAN, there is normally one point of control on any given link (usually called the primary or master station). It is this station's responsibility to ensure correct operation of the link including flow control, and to initiate error recovery as needed. In LANs, there is no one point of control; the functions have been distributed throughout the network (LAN segment). In order to coordinate this distributed function, a 'protocol' had to be implemented at the MAC layer; this protocol implementation is through the use of MAC type frames.

MAC type frames are passed back and forth between adapters for the purpose of coordinating activity related to the physical media access (hence the name media access control). Some of the functions performed by these frames include:

- Token control (e.g., issuing free tokens)
- Configuration management
- Error notification and recovery

- Network management

It is not the intent here to address all 25 of the MAC frames in detail, but rather to expose you to their existence and purpose, and to give you some idea as to the activity that occurs at this level totally independent of and transparent to any upper layer functions.

Examples of MAC Type Frames

Ref.: Chapter 5 of the Architecture Manual

- Claim Token
  - Used by the adapter cards to negotiate who will be the active monitor and be responsible for issuing new tokens
- Active Monitor Present
  - Used to inform adapters that there is an active monitor on the ring
  - Periodically issued to initiate NAUN notification
- Beacon
  - An indication sent by an adapter card indicating that it no longer sees a valid incoming signal
- Duplicate Address Check
  - First frame sent by an adapter card after ring insertion to ensure that the unique address it is using is, in fact, unique
- Report Soft Error
  - Used by an adapter to indicate that it is logging recoverable errors

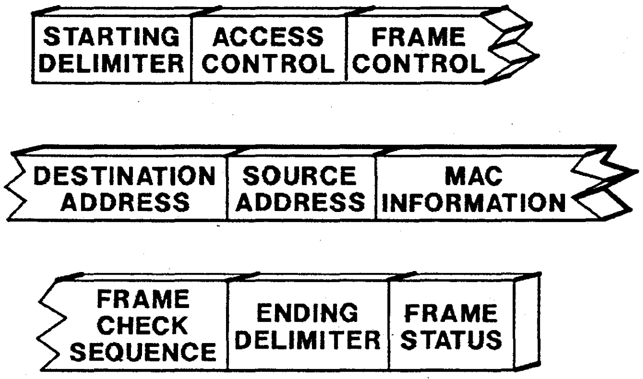

BUSY TOKEN WITH DATA

LANA276

MAC Frame Format

The actual MAC type frame format is almost identical to the data frame; the only difference is that the L-PDU portion of the data frame has been replaced by the MAC information fields.

Overview of MAC Type Frame Format

- First 4 bytes = MAC LLID (length and ID)
  - Bytes 1 & 2 = major vector length
  - Byte 3 = destination and source class
    - Classes include Ring Station, DLC.LAN.MGR, Configuration Report Server, Ring Parameter Server and Ring Error Monitor
  - Byte 4 = specific code point of vector
  - Bytes 3 and 4 identify the function (purpose) of the vector
- Remainder of vector consists of one or more variable length subvectors
  - Subvectors contain additional information about the function
  - For example, a 'Beacon' major vector would include subvectors identifying the NAUN, Beacon type and physical location

ABORT FORMAT

STARTING DELIMITER	ENDING DELIMITER

LANA220

Abort Frame

An abort frame is used to clear the ring of any data after an error has occurred.

- An abort frame consists of 2 bytes: an SD followed by an ED

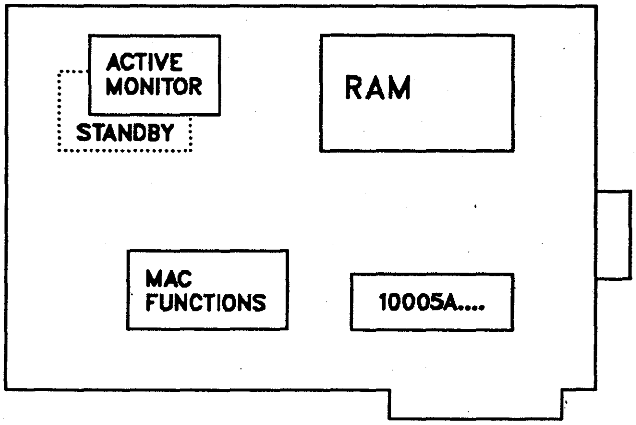

LANA201

IBM Token-Ring Adapter Cards - Overview and Operation

All IBM Token-Ring adapter cards, regardless of the device into which they are installed, have the following capabilities

- Implement IEEE standards 802.5 (MAC sublayer functions) and the interface to the IEEE 802.2 sublayer
- Adapter addressing

  Note: All addresses are 6 bytes in length.

  - Universally Administered (burned-in) - unique to every card
  - Locally Administered - software override of UA when adapter opened
- Provide significant ring level management capabilities
  - Internal adapter diagnostics
  - Lobe fault detection
  - Active and passive monitor capabilities
  - Error logging and reporting
- Active monitor functions (one per ring):
  - Monitors ring for token failures (none, too many, continuously busy), resolves failure (ring purge) and issues a new token
  - Provides master clock synchronization (eliminates jitter)
  - Handles NAUN notification every 7 seconds
  - All adapters have function; only one has it active at a time
  - All other adapters operate in standby monitor mode; this means that they monitor the activity of the active monitor, and if the active monitor does

not perform its functions, will attempt to become the the active monitor through a contention process

- Adapter cards are device specific; that is, PC adapter cards will not work in other hardware such as the RT or ES/9370

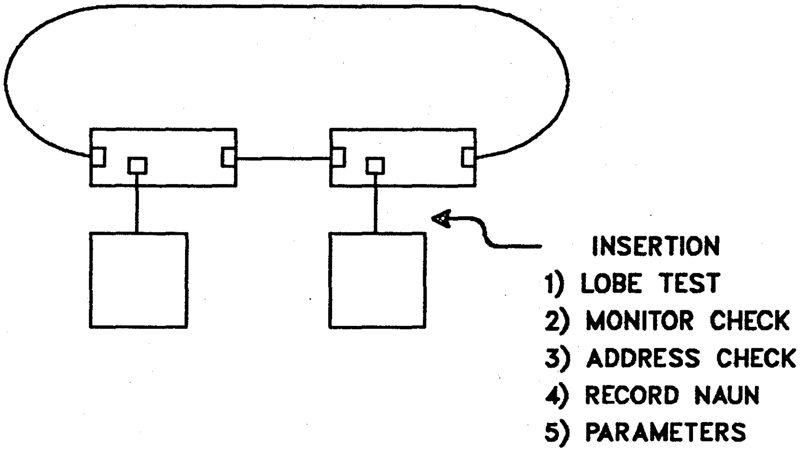

LANA202

Inserting the Adapter into the Ring

All adapters attach to the ring through a 5 step process:

1. Lobe test
   - Adapter transmits over the lobe to the 8228 Multi-Station Attachment Unit (MAU) port where the signal is wrapped back to the adapter
   - Started as a result of an 'open' command to the adapter
   - 2K frames sent and received
   - If 2 errors - test fails
2. Monitor check
   - Adapter activates relays in 8228 physically attaching the station to the ring (phantom voltage from adapter)[5]
   - INSERTION ONTO THE RING WILL CAUSE ERRORS - NO CAUSE FOR ALARM
   - Waits up to 18 seconds to see if an active monitor is on the ring (looks for Active Monitor Present or Ring Purge MAC frames)
   - If no active monitor, adapter enters monitor contention
3. Duplicate address check
   - Ensures hardware address (locally administered or burned-in) is unique on THIS ring (frame does not cross bridges)
   - Adapter transmits message to itself using a special MAC type frame

[5] The phantom voltage is a +5VDC signal superimposed across the wire pairs (i.e. one pair of wires at +5VDC and the other at 0VDC) used to activate the relay in the wiring concentrator.

- If another station has recognized and/or copied frame, this adapter closes with an 'open' failure

4. Participate in ring poll

   - Ensures station knows its NAUN and is known by its nearest active downstream neighbor
   - Station waits to participate in ring poll process

5. Request parameters

   - Station can request parameters from the Ring Parameter Server:
     - Ring number
     - Soft error report timer
     - Maximum authorized ring access priority

IEEE 802.5 Token-Ring Standardized Media

Several different types of media are identified in the standard including coaxial cable; IBM does not implement the use of coaxial cable in its Token-Ring LAN implementation.

- Shielded twisted pair wire
- Unshielded twisted pair wire

  Note: IBM supports the use of this type of medium at 4Mbps only; some vendors support use at either 4Mbps or 16Mbps
- Fiber optic cable

IBM Products That Implement This Standard

ANY IBM PRODUCT THAT SUPPORTS A TOKEN-RING LAN ATTACHMENT WILL IMPLEMENT THE ABOVE STANDARDS AND PROTOCOLS.

- IBM PS/2 - all models
- IBM PC
- Communication Controllers - 3745, 3720 and 3725
- 3174 Control Units
- 9370 systems
- System/36
- System/88
- AS/400
- IBM RT
- Series/1
- 3172 and 8232 LAN Channel Stations
- 4702 Finance Communication System
- 8209 Bridge

Exercise 1

Please go to Appendix B and complete Exercise 1.

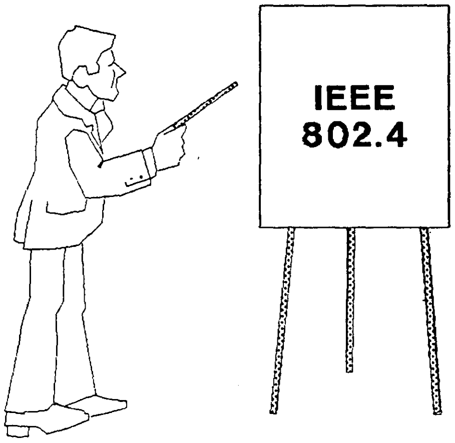

Session 3 - IEEE 802.4 Token-Bus Standard

Objectives

Upon completion of this session, you should be able to:

- Describe the operation of the IEEE Token-Bus LAN
- Identify the purpose of the MAC fields in the IEEE 802.4 standard
- List the IBM products that attach to the Token-Bus LAN.

IEEE 802.4 - Token Bus

- Standard applies to a LAN using a token-passing protocol on a physical bus topology
  - Token-passing is a method in which stations share equal access to the network, but can transmit only when in possession of the 'token'
- Standardized for both baseband and broadband applications
- Standardizes the details of the token-passing method, protocols, physical media, and physical conditions

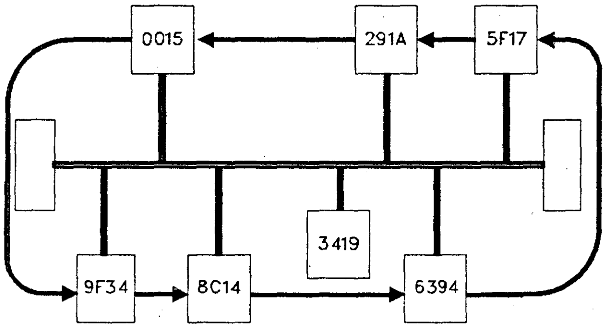

LAN4272

Token Bus Operation

- Configuration is a physical bus treated like a logical ring
  - Sequence of token passing on 'ring' is typically in descending address order (highest to lowest active address)
- Device with information to send waits for a token
- When the device receives the token, it holds it for a predetermined amount of time during which it can transmit and receive data, and poll other devices on the bus (called "token holding time")
- When the time expires, the token is passed to the next device in the sequence
- Devices with no data to send will send the token to the next station immediately
- An optional priority scheme can be used to prioritize messages in each station, allowing high priority traffic to be sent first
- Token management:
  - Periodically, each station will suspend token passing, and send a special 'solicit successor' frame to determine if any new devices have become active on the network within a given address range (terminal sending frame and its present successor)
  - Devices will be 'inserted' into the logical ring based on their address
  - When a device is removing itself from the ring, it will send a special frame to its predecessor indicating that the predecessor should update its successor address to the next station on the network

- Other control frames are defined to address lost token conditions and removal of stations from the network in other than an orderly manner (e.g., station loses power)

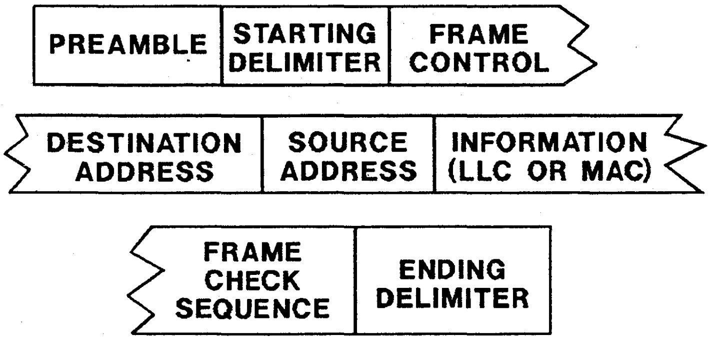

LANA215

Token Bus MAC Frame

There are three types of frames defined for the token bus network:

- LLC frames containing an LLC header and may contain user information (data from the upper layers)
- Station management data frames (adding, deleting stations, etc.)
- Special purpose data frames

Frame Format

- Preamble - 1 or more bytes
  - Purpose of the preamble is to provide a synchronization pattern for the network stations on the bus
  - Size and specific pattern will be determined by the physical network implementation (i.e. speed, modulation technique, etc.)
- Starting delimiter - 1 byte
  - Purpose of the starting delimiter is to provide an indication that data follows; the SD will always be a non-data pattern (e.g., unique coding pattern)
- FC = frame control - 1 byte
  - The purpose of this field is to identify this frame as an LLC data frame or some station management or MAC function
  - Examples of station management and MAC functions include Claim Token, Who Follows, Token, Solicit Successor, etc.
- DA is a 6 byte (or 2 byte) destination address (target or receiver)

- SA is a 6 byte (or 2 byte) source address (sender)
  - Destination and source address fields are defined the same as in the 802.5 Token-Ring standard
- LLC or MAC information - variable size
  - This field contains logical link control and user data, station or token management information or special purpose data as indicated in the FC field
  - Size will vary depending on data content but can not exceed 819 bytes as defined by the standard

  Note: Unlike CSMA/CD which requires a minimum frame size, this field can be any value between 0 and 819 bytes inclusive.
- The Frame Check Sequence (FCS) field contains check characters used to validate data integrity (4 bytes)
  - Data integrity is checked between the FC field and the FCS field inclusive
- ED is a 1 byte ending delimiter
  - Ending delimiter is used to indicate the end of the data and the location of the FCS bytes
  - Like the SD field, ED will always be a non-data format

IEEE 802.4 Physical Network Implementations

The 802.4 standard allows for three different transmission techniques:

- Broadband transmission
  - Data is transmitted using a combination of amplitude modulation and phase shift keying (AM/PSK)
  - Data rates supported are: 1M bps, 5M bps and 10M bps
  - Specific channel bandwidth requirements will depend on the speed
  - Standard CATV type coaxial cable is used as the transmission medium (75 ohm coax with F-type connectors)
- Baseband using single-channel (single user) phase-coherent frequency shift keying (FSK)

  Note: The term "carrier band" is sometimes used to refer to baseband FSK transmission implementations.

  - Data is transmitted by switching between two discrete frequencies
  - Data rate is either 5M bps or 10M bps
  - Standard CATV type coaxial cable is used as the transmission medium (75 ohm with F-type connectors)
- Baseband using single-channel phase-continuous FSK
  - Like phase-coherent FSK, the data is transmitted by using two different frequencies; however, in this implementation, the transitions are smooth (continuous phase)
  - Trunk cable is 75 ohm coax; drop cable is 35 to 50 ohm coax using BNC-type connectors
  - Data rate is 1M bps

Media and Connectors

- Trunk cable is 75 ohm coax (e.g., RG 6/U, RG 11/U)
- Drop cable may be 35 to 50 ohm, or 75 ohm, depending upon transmission technique being used
- Station uses a 50 ohm male BNC-type connector
- Trunk to drop connector is a 75 ohm 'T' connector

IBM Products That Implement This Standard

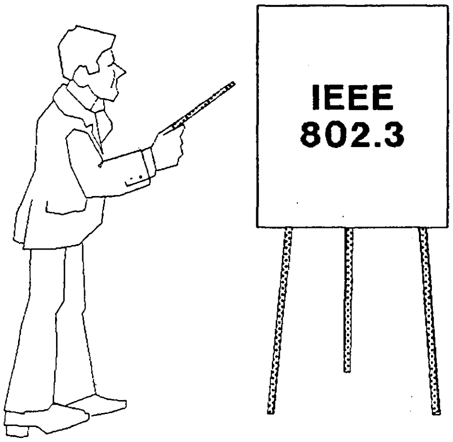

Session 4 - IEEE 802.3 CSMA/CD Standard

Objectives

Upon completion of this session, the student should be able to:

- Identify the characteristics of a CSMA/CD Bus Architecture
- Describe how Carrier Sense Multiple Access with Collision Detect works to allow access to the bus
- Explain the various fields in the IEEE 802.3 MAC frame
- State the difference between baseband and broadband implementations
- Describe the various hardware components in an IEEE 802.3 network
- List the IBM products that attach to CSMA/CD LANs

IEEE 802.3 Standard

- Standard defining CSMA/CD (Carrier Sense Multiple Access/Collision Detect)
  - Based on the random access technique conceived by Xerox Corporation in 1973
- Coaxial Cable is the most widely installed medium but some users are now implementing twisted pair cable and fiber optics
- Standardized for both baseband and broadband implementations
- 1 to 20 Mbps Transmission Speeds
- NOT Ethernet (see unit on Ethernet V2.0 for details)
- Data Encoding Scheme - Manchester Encoding

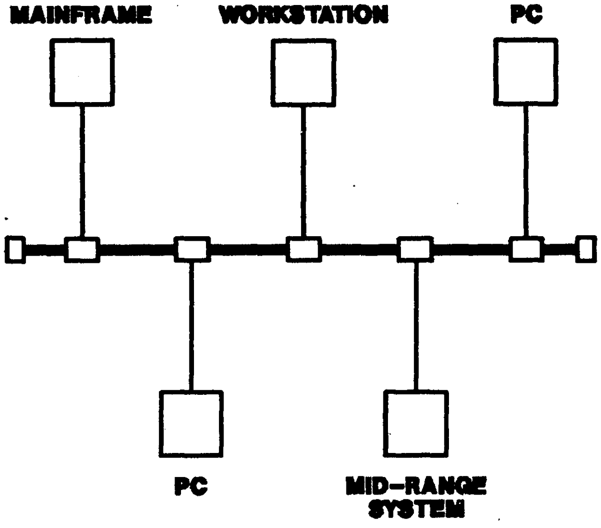

LANA282

CSMA/CD Operation

"Listen before transmitting and listen while transmitting"

- Analogy - Party Line
  - MA (Multiple Access) - Every person on the party line has access to the same wire
  - CS (Carrier Sense) - Person listens first to see if line is being used.
  - CD (Collision Detect) - When two listen at same time and hear nothing so that both ring operator. Must have technique to resolve this.
- CSMA/CD on LAN
  - MA (Multiple Access) - Devices share same wire (bus)
  - CS (Carrier Sense) - Before transmitting, must listen for activity. If there is activity, must continue listening until no activity.
  - CD (Collision Detect) - When two devices begin transmitting at the same time. Both continue listening while transmitting and, if they hear a collision (distortion in their transmitted signal), then they begin a resolution process as follows:
    - Transmit Short Jamming Signal
    - Calculate a random time to wait before trying again
    - Listen. If no activity, then retransmit, else wait until line is not active.
  - The CD resolution process will be repeated 16 times before the inability to transmit is reported; this technique is called "truncated binary exponential backoff"
  - Collision rate increases as network loading increases, potentially impacting performance
  - Packet size must be at least 64 bytes to ensure collisions are detected. This size is based on the propagation time from one end of a 2500 meter LAN to the other end and back with four repeaters in-between.

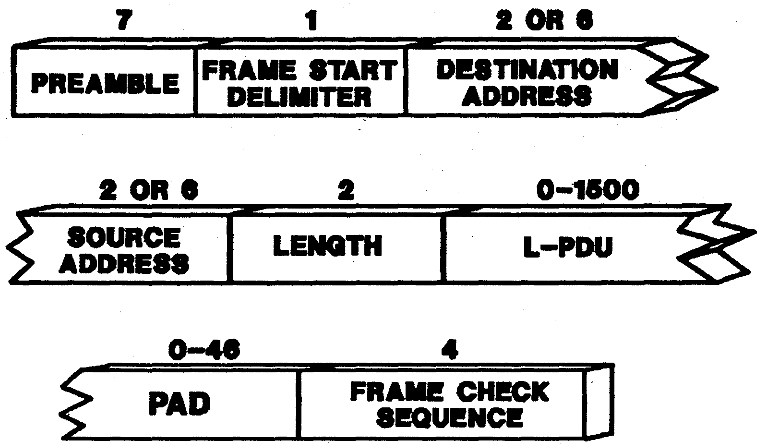

LANA284

IEEE 802.3 CSMA/CD MAC Frame

The data frame consists of a minimum of 64 bytes

- Preamble (7 bytes)
  - Alternating 1 and 0 bits
- Start Frame Delimiter (1 byte)
  - Alternating 1 and 0 bits
- Destination Address (2 or 6 bytes)
  - All adapters must use same address size
  - Singlecast, Multicast, or Broadcast Address
    - Singlecast - individual adapter's address
      - If 6 byte addresses are being used, this is either the universal address (permanently stored in the adapter) or a locally-administered address.
    - Multicast - group of adapters will receive the packet
    - Broadcast - all adapters receive the packet
- Source Address (2 or 6 bytes)
  - Singlecast address only
    - If 6 byte addresses are being used, this is either the universal or locally-administered address of the sending adapter.
  - Length (2 bytes)
    - Number of bytes in the L-PDU
    - 0 to 1500 bytes

- L-PDU (0 to 1500 bytes)
  - The L-PDU field (Logical link control - Protocol Data Unit) contains logical link control fields and user data
  - Data field within the L-PDU can be in a variety of formats (e.g. SNA, TCP/IP, NETBIOS or user-defined)
- Pad (0 or more bytes)
  - Filler so that packet will be 64 bytes long. If packet is already 64 bytes or more, then this field is not used.
- Frame Check Sequence (4 bytes)
  - Cyclic Redundancy Check computed on fields from the destination address thru the pad. Recomputed at destination and compared to this field to verify integrity of the packet.

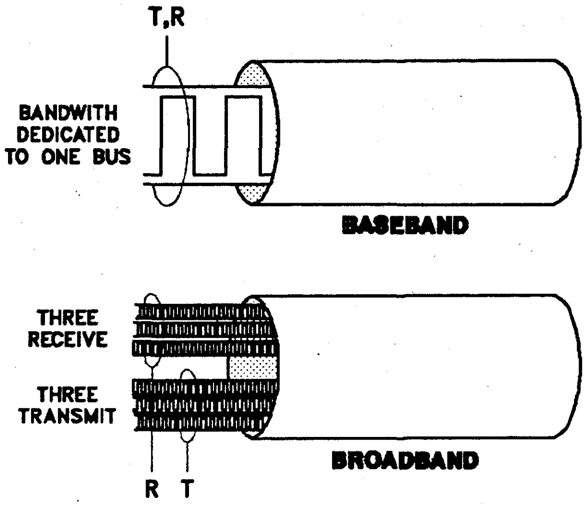

LANA285

Baseband and Broadband Implementations

The IEEE 802.3 standard recognizes the use of media for both baseband and broadband signalling. Traditionally, baseband has been used but there is a growing interest in implementing broadband CSMA/CD networks since they can use available bandwidth on existing cable TV networks. This could make it possible to reach the home user market.

Differences:

- The primary distinction between these methods is that baseband LANs use the entire bandwidth of the medium for both transmission and reception. In contrast, broadband LANs use three 6MHz channels for transmission and another three 6MHz channels for reception.
- Baseband segments are limited to 500 meters each whereas broadband segments can be up to 3600 meters per segment.

10 BASE 2

10 BASE 5

10 BASE T

10 BROAD 36

1 BASE 5

LANA204

Hardware Descriptions

The physical and MAC layers are implemented using a number of hardware components. These include cabling, terminators, transceivers, transceiver cables, adapters, repeaters and bridges.

- Baseband Cabling
  - IEEE 802.3 defines thick or thin coaxial cable and telephone twisted pair for the bus. Although it is common to refer to the coaxial cables as thick Ethernet and thin Ethernet, it will not be used here to minimize the confusion between IEEE 802.3 and Ethernet (see the unit on Ethernet for details on differences).
  - Characteristics of thick coaxial cable:
    - 0.405 inches if PVC is used
    - 0.375 inches if Teflon is used
    - Maximum length per segment of 500 meters.
    - If repeaters are used, the maximum distance between two devices is 2500 meters.
    - No more than 100 devices per segment.
    - Up to 10 Mbps
    - Known as 10Base5
  - Characteristics of thin coaxial cable:
    - RG58R/U (Cable TV cable)
    - Maximum length per segment of 200 meters.
    - No more than 30 devices per segment.
    - Known as 10Base2
  - Characteristics of Telephone Twisted Pair (TTP):
    - Unshielded
    - Maximum length per segment of 100 meters.

    - Known as 10BaseT

- Broadband Cabling
  - RG58R/U (Cable TV cable)
  - Uses 36Mhz for LANs (3 transmitting and 3 receiving channels)
  - Two cabling options:
    - Use two cables, one for transmission and one for reception. The cables are connected at one end through a head-end amplifier. In this case, the transmission and reception frequencies are the same.
    - Use one cable with a head-end amplifier and frequency translator at one end that repropagates the low-frequency channel (transmission channel) onto the high-frequency channel (reception channel).
    - Known as 10Broad36
  - Terminators
    - A terminator is a resistor placed at the end of a cable which matches the impedance of the cable. Its purpose is to conduct any energy to ground so that no reflection occurs when a signal reaches the end of the cable. Each end must have a terminator.
  - Up to 20 Mbps

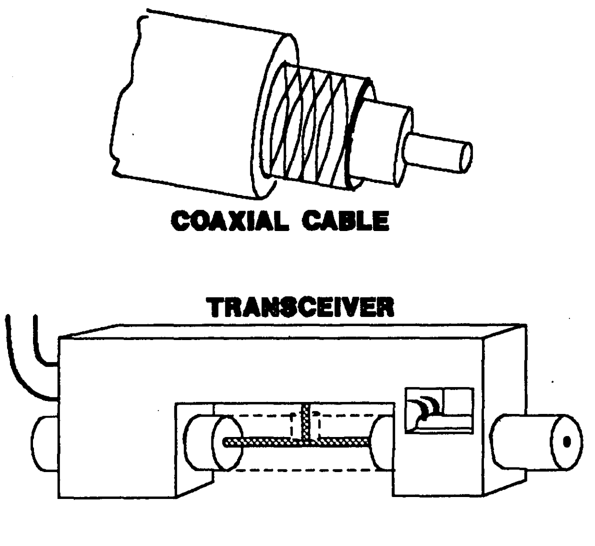

- Baseband Transceivers
  - For thick coaxial cabling, a ″vampire″ tap transceiver is used. It is attached to the bus by drilling a hole to the center wire of the coax and then by clamping the two halves of the transceiver onto the cable. A pin extends into the hole and makes contact with the center wire. At the same time, ″teeth″ on the transceiver penetrate the outer insulation to contact the ground sheathing of the coax. Hence the name ″vampire″. Cable is marked every 2.5 meters. By placing transceivers on these marks, ensures a minimum distance between transceivers.
  - For thin coaxial cabling, the bus is physically run to each device. The cable is cut at that point and a BNC connector is attached to each end. A coax tee is connected to the adapter and each end of the cable is connected to the tee. Since the tee provides continuity between the two bus segments, the coaxial cable is still physically one continuous wire.
  - For telephone twisted pair (TTP), the adapter has two RJ11 or RJ45 jacks for the TTP wire to plug into. TTP is run between each device in a serial fashion.
  - For both thin coax and TTP, the transceiver is built onto the adapter.
- Broadband Transceivers
  - These radio frequency transceivers connect to adapters using the same transceiver cable and electrical interface as with baseband transceivers.
  - Modulates the digital data and puts it on one channel of a multichannel network. Low-end frequencies are used for transmission and high-end frequencies are used for reception.

- Transceiver Cables
  - Transceiver cables connect the transceiver on a thick coaxial cable to the adapter in the device. The cable has 15 wires and can be no longer than 50 meters.
- Adapters
  - An adapter is a board inserted in the LAN-attached device. Its functionality varies from vendor to vendor but it typically implements the physical layer and MAC-sublayer of the data link layer. If thin coaxial cable or telephone twisted pair is used, then a transceiver is built-in. Otherwise, the transceiver is a separate box (see the description of transceivers above).

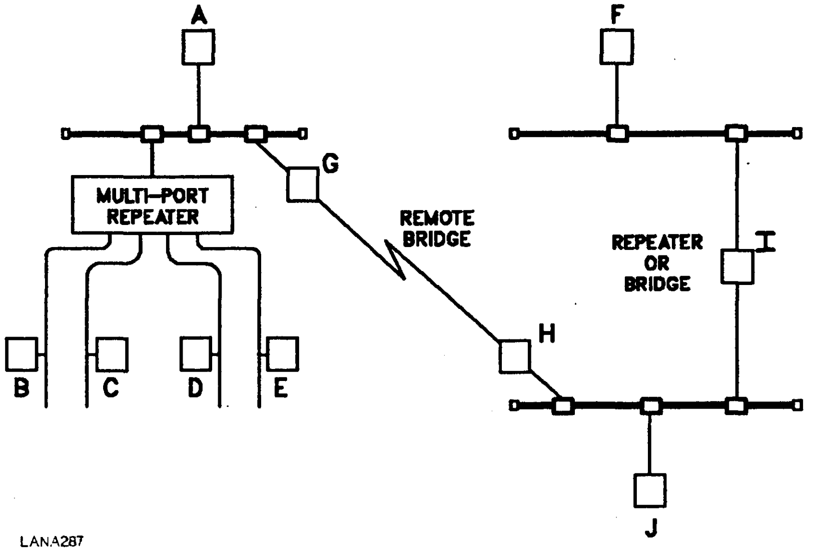

- Repeaters
  - In order to expand an 802.3 LAN to more than one bus, one way is to repeat all signals from one bus to another. A repeater is a device that attaches to two busses and copies bits in both directions at the same time.
  - It is recommended that a packet not traverse more than two repeaters before reaching its destination.
  - Multi-port repeaters allow multiple segments to be connected to a central point. This simplifies fault isolation since problems can be isolated to a single segment by only manipulating the connections to the multi-port bridge.
- Bridges
  - To isolate segments of a LAN, bridges can be used. Allows much larger LANs.
  - Often uses a Spanning Tree algorithm for routing of packets in multi-segment networks.
  - Buffers entire packet before transmitting it on the next segment.
  - Remote bridges allow segments physically removed from each other to be connected via leased lines. A remote bridge is used at each segment to make this connection.

IBM Products That Implement This Standard

- 8232 LAN Channel Station
- 3172 LAN Station
- E/S 9370 Systems
- IBM RT
- IBM PS/2 (Mdl 50 and above)
- 8209 LAN Bridge

IBM PC Network (CSMA/CD)

- CSMA/CD protocols very similar to IEEE 802.3 standard
- Baseband and Broadband transmission (data rate 2 Mbps)[6]
- Bus topology
  - Coax cable medium used for broadband
  - UTP (TTP) medium used for baseband

[6] Broadband transmission uses two 6MHz CATV channels, one for transmitting and one for receiving; this is different from the IEEE 802.3 standard. Also, note speed differences.

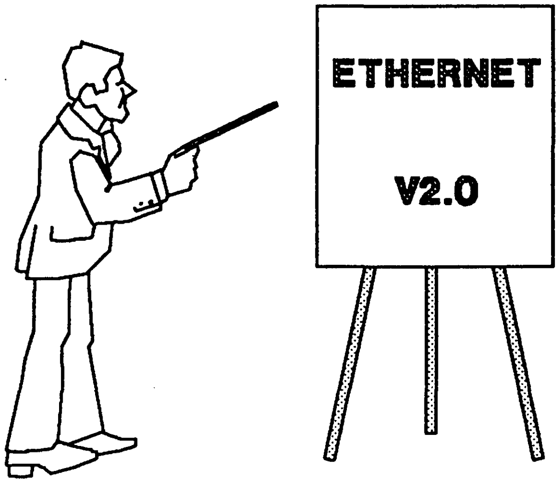

Session 5 - Ethernet V2.0 MAC Frame

Objectives

Upon completion of this session, the student should be able to:

- Describe the history of Ethernet
- State the differences between IEEE 802.3 and Ethernet (DIX V2.0)

  Note: Ethernet V2.0 is also known as DIX V2.0 where DIX stands for the major corporations involved (Digital, Intel and Xerox).

History of Ethernet

- Began with network developed at University of Hawaii called ALOHA. Used ground-radio based terminal network. Included multiple-access with collision detection. The central facility in Oahu would send an acknowledgement if received successfully. Otherwise, would be retransmitted. NO carrier sense so collisions occurred often. On an average, 18% of the available bandwidth was used effectively due to the large number of retries.
- Ethernet developed at the Xerox Palo Alto Research Center (PARC). The definitive paper on Ethernet was published in 1976 by R.M. Metcalf and D.R. Boggs.
- Original Ethernet functioned at 3Mbps and allowed 100 connections.
- Through a consortium in the early 1980's, Xerox, Intel, and Digital Equipment Corporation published what is now known as the Ethernet Blue Book. This was further evolved and eventually produced a cooperative standard known as Ethernet Version 2.0 in 1982.

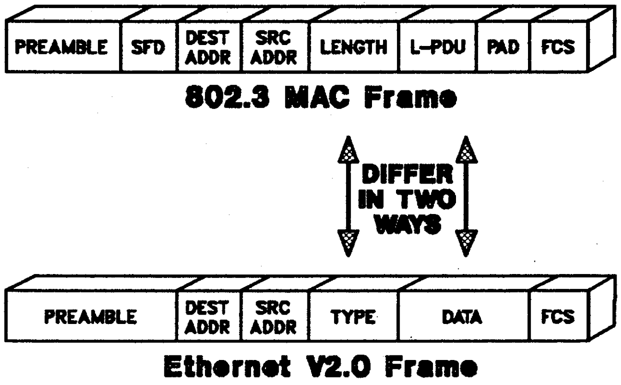

LANA289

Differences Between Ethernet and IEEE 802.3

Ethernet has four distinct differences from the standard defined by the IEEE 802.3 committee. Any one of these differences make them incompatible and one of the differences has no acceptable work around.

- **Protocol Differences (two)**
  - NO LLC CONTROL FIELDS IN ETHERNET
    - Since Ethernet was implemented before standards came out, it didn't include the Logical Link Control fields that IEEE Project 802 decided must envelope any data coming from software in layers 3-7. The LLC fields allow a machine, once it has successfully received a a packet, to send the data to the appropriate communications software. Since there is no LLC in Ethernet, the information about which communications software gets the data was imbedded in the TYPE field of the Ethernet packet.
    - *Workaround:* NONE. There is no way for communications software to know whether it is running on an Ethernet or 802.3 implementation of Layer 2 so it could not build the LLC envelope. This would be a major violation of the OSI concept of higher layers having no knowledge of lower layers. Any later change to layer 2 (like changing adapter vendors) could cause the communications software to no longer work.
  - TYPE AND LENGTH FIELD CONFLICT
    - Ethernet's TYPE field is in the same location as IEEE 802.3's LENGTH field. Types can range from 0 to 65535 whereas the length of data is from 0-1500.

  - *Workaround:* Xerox has not assigned any protocol types from 0 to 1500. Therefore, an intelligent controller can look at the type/length field and determine, by the value, if the packet is in Ethernet or 802.3 format.

- APPARENT DIFFERENCE IN PREAMBLE

  - Last byte of first 8 bytes appears to be different. Actually, both formats use alternating ones and zeroes in the first 8 bytes even though 802.3 chooses to call the last byte a Frame Start Delimiter so there is no actual difference in the first 8 bytes.

- **Hardware Differences**
  - **CABLE SIZES DIFFER**
    - The thick Ethernet cable is 0.395 inches in diameter. In contrast, the thick IEEE 802.3 cable is 0.405 (PVC) or 0.375 (Teflon) inches in diameter.
    - This poses a problem since a transceiver has a pin which is supposed to just make contact with the center wire of the cable. If the diameter of the cable changes, the transceiver pin won't make contact.
    - *Workaround:* Some vendors of transceivers have adjustable center pins so that they will work with any of the three cable diameters.
  - **SIGNAL QUALITY ERROR SIGNAL USAGE DIFFERS**
    - One of the 15 signal lines between a transceiver and an adapter is called Signal Quality Error. Ethernet designers decided that signal quality wasn't a problem so they used this signal for something more important, determining if the transceiver is working or not. Periodically, the transceiver sends a pulse down this line to indicate that it is alive.
    - Unfortunately, the IEEE 802.3 committee chose to use this signal line for its original intention - as a signal quality error indicator. This results in two bad scenarios:
      - Scenario 1 - 802.3 Adapter with an Ethernet transceiver. Periodically, the adapter gets an error indication and retries the current transmission. Results in serious degradation of performance.
      - Scenario 2 - Ethernet Adapter with an 802.3 transceiver. Adapter doesn't get the expected heartbeat and so assumes the transceiver is not working.
    - *Workaround*
      - Scenario 1 - Cut the SQE line so no error indication occurs. Of course, if signal quality does degrade, there will be no indication except that packets will sometimes come in mangled and be rejected.
      - Scenario 2 - Not as easy a solution. Must artificially generate the pulse so the adapter is happy.

IBM Products That Implement This Standard

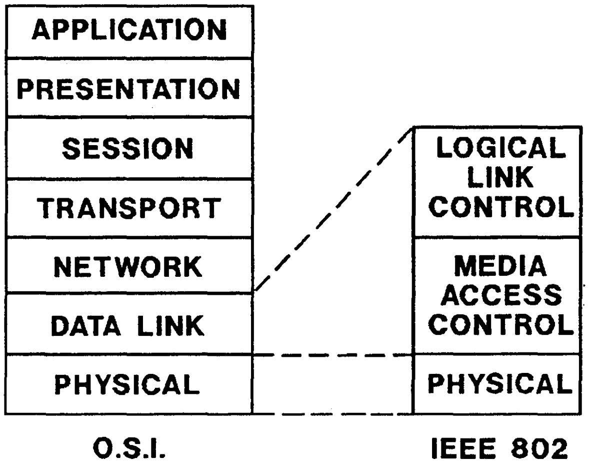

LANA290

Session 6 - IEEE 802/O.S.I. Relationship

Objectives

- The 802 Reference Model consists of two (2) layers, with one of the layers subdivided into two sublayers
  - Physical Layer
  - Media Access Control (MAC) Sublayer (part of DLC layer)
  - Logical Link Control (LLC) Sublayer (part of DLC layer)
- The Physical layer of both IEEE 802 and O.S.I. are defined the same.
- The IEEE 802 MAC and LLC layers relate to the O.S.I. Data Link layer
  - The MAC sublayer is defined as part of IEEE 802.3, 802.4, and 802.5
  - The LLC sublayer is defined in IEEE 802.2

  Note: A MAC sublayer interface is implemented at the physical layer.
- Conversion of the IEEE standard numbers to OSI equivalents:
  - Add an '8' to the IEEE standard
  - Change the '.' in the IEEE standard to a '/'
  - Example: IEEE 802.2 = OSI 8802/2

    Note: This conversion only applies to IEEE 802.1 to 802.5 standards.

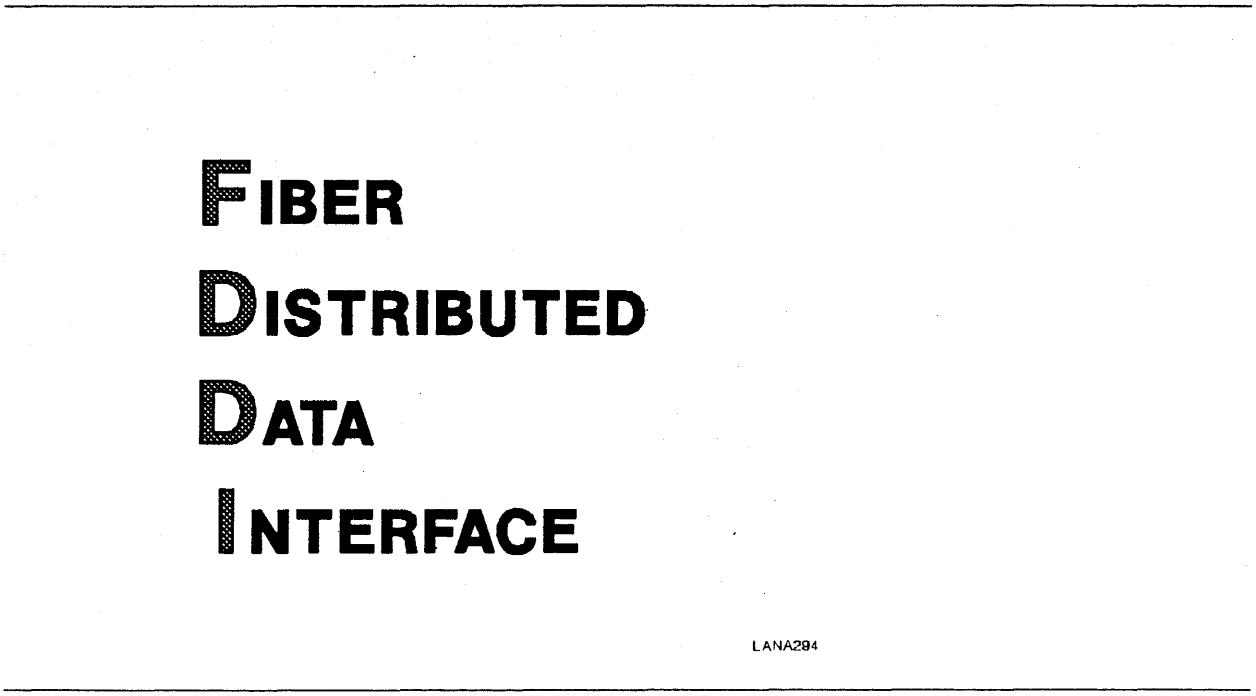

FIBER

DISTRIBUTED

DATA

INTERFACE

LANA294

Session 7 - Fiber Distributed Data Interface (FDDI)

Objectives

Upon completion of this session, you should be able to:

- Briefly describe FDDI
- State the functions of the 4 sublayers in an FDDI implementation
- Briefly describe how bandwidth management is accomplished in an FDDI LAN
- State the purpose of the MAC layer fields in an FDDI frame
- Briefly describe the physical topology used in an FDDI LAN and state the differences between the classes of stations as they relate to the topology.

What is FDDI?

- Standard under development by American National Standards Institute (ANSI) X3T9.5 group
- Defines protocols and standards to be used on a token-passing ring configuration
  - Speed = 100 Mbps data rate[7]
  - Medium = Fiber optic cable (standardized for several fiber sizes)
  - Topology = ring
  - MAC protocol = timed token (see "Token Management" below)

[7] Actual bit rate is 125 Mbps; data rate is 100 Mbps because of the encoding scheme being used (4B/5B).

- Standard addresses the physical and MAC sub-layers only and is written to the 802.2 LLC interface
- Token-passing protocol implements the 'early token release' (similar to the 16 Mbps Token-Ring)
- Stations can have one or two physical connections to the ring
  - Single ring connected stations can go through a wiring concentrator
  - Stations with physical connections to both paths can automatically bypass link (ring) failures
- Rings are dynamically built as stations become active

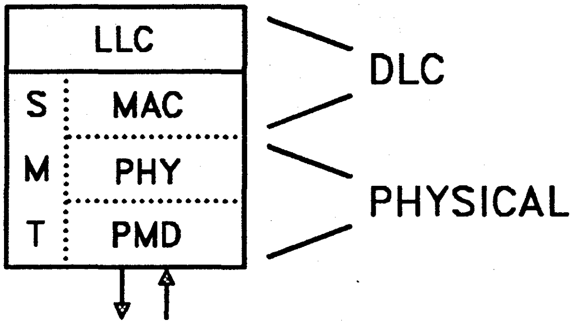

LANA206

Network Sub-Layers

Because of the additional complexity of the network, the functions have been divided into sublayers.

- PMD and PHY are sublayers within the physical layer (OSI layer 1) implementation
- PMD - physical media dependent
  - Defines physical fiber connections, driver and receiver characteristics, optical coding techniques and hardware (i.e. connectors, etc.)
- PHY - physical layer protocol
  - Defines synchronization and physical encoding/decoding techniques
- MAC and LLC are sublayers within the data link control (OSI layer 2 DLC) implementation
- MAC - media access control
  - Defines the media access control protocols (i.e. token management, error recovery, etc.) the same as done for the IEEE MAC standards
- LLC - logical link control
  - Standard is expected to be implemented on LANs using 802.2 LLC protocols
- There is an additional implementation of function called station management (SMT)
  - SMT is responsible for coordinating the activity between the various sublayers including insertion of a station into the ring, error recovery (bypass of a failed ring segment) and station de-insertion

Token Management

Token management is similar to but not the same as the 802.5 standard.

- Token is passed around the ring (same as 802.5)
- Stations with information to send will make the token busy and insert the data (same as 802.5)
- Station can send data until the 'token holding timer' expires (multiple frames)

  Note: This is the same as is defined in the 802.5 standard; IBM's token-ring implementation calls for sending one frame and then releasing the token.
- Sending station will append a free token to last data frame (similar to early token release used for 16M bps rings)
- Target stations copy the data and turn the appropriate status bits on
- Sending station removes the data
- Confirmation of delivery, retransmissions, etc. are not the responsibility of the MAC protocol (same as 802.5)
- FDDI does not use token priorities like those found in the 802.4 and 802.5 standards; instead it controls bandwidth allocation using token rotation timers for synchronous transmission and a combination of priorities and restricted tokens for asychronous transmission

Bandwidth Allocation

- Each station will determine its bandwidth requirements and negotiate with other stations on the ring to determine the target token rotation time (TTRT) for the ring
- The TTRT represents the requirements of all active stations on the ring plus ring transmission latencies and is recorded by all stations
- The TTRT value then determines time allotted to both synchronous and asynchronous transmission
  - Bandwidth requirements represent priority data from that station; synchronous transmissions are used for priority data
  - Lower priority data is sent using asynchronous transmissions
  - Each station keeps track of how long it takes for the token to go around the ring in a token rotation timer (TRT) and compares this value with the TTRT value
  - When a station receives the token, it uses its synchronous time slot to send priority data
  - If the token was received prior to the expiration of the TRT timer (TRT less than TTRT) the station can also send asynchronous frames for the time difference (asynchronous transmission is only allowed when the token rotates around the ring in less than the maximum time)
  - Asynchronous transmission can be further controlled through the use of priorities and ′restricted tokens′ where a connection between two ring stations (called a dialogue) can be established

Token Recovery

- Ring recovery for token failures works somewhat differently than in the 802.5 implementation
- All stations monitor activity on the ring and all stations will attempt to recover the network when a failure is recognized
- Stations attempt to reinitialize the ring by issuing special 'claim token' frames; the frame includes a value for the target token rotation time (TRT)
- The station with the lowest target TRT value will win the contention process and will be responsible for initializing the ring

  Note: The contention process is similar to the 802.5 active monitor selection process where the highest station address wins.
- Ring breaks will result in a beaconing process conceptually the same as defined in the 802.5 standard

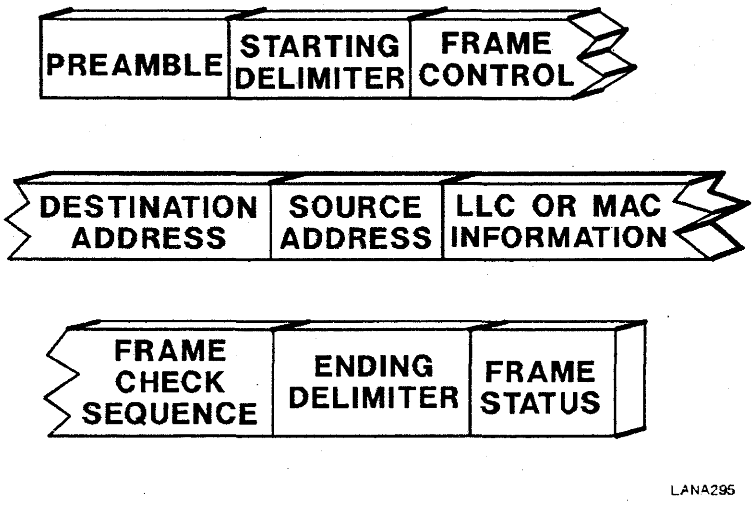

FDDI MAC Frame

The FDDI MAC frame is similar to the 802.5 MAC frame. One major difference is that the term 'symbol' is used to represent 4 bits. The reason for this is the 4B/5B coding scheme used to send the MAC information around the ring.

Several frame formats are defined:

- LLC frames (may or may not contain user data)
- Control frames which include MAC and token formats
- Two formats reserved for future use

Frame Format

- Preamble - 16 symbols (8 bytes)
  - Purpose of the preamble is to provide a synchronization pattern for the network stations
  - Synchronization is needed since there is no master clock on the ring; each station reclocks the data and uses a latency buffer to compensate for jitter
- Starting delimiter - 2 symbols (1 byte)
  - Purpose of the starting delimiter is to provide an indication that data follows; the SD will always be a non-data pattern (e.g., coding violation)
- FC = frame control - 2 symbols (1 byte)
  - The purpose of this field is to identify the frame type and address format being used
  - Frame type may be LLC or MAC control

  - One bit is used to indicate whether this is a synchronous or an asynchronous transmission
  - One bit is used to indicate whether this particular DA and SA field is 4 or 12 symbols (2 or 6 bytes)

- DA is the destination address (target or receiver)

- SA is the source address (sender)

  - Destination and source address fields are defined the same as in the 802.5 Token-Ring standard
  - As indicated above, the DA and SA fields may be either 2 or 6 bytes
  - Address formats may be mixed on the ring although two stations must use the same format to communicate with each other

- LLC or MAC information - variable size

  - This field contains logical link control and user data, or MAC data as indicated in the FC field
  - Size will be dependent on data content but maximum total frame length should not exceed 4000 bytes
    - Frame length is limited because each station must reclock the data (different from 802.5 where there is one master clock per ring)

- The Frame Check Sequence (FCS) field contains check characters used to validate data integrity (8 symbols - 4 bytes)

  - Data integrity is checked between the FC field and the FCS field inclusive

- ED is the ending delimiter - 1 or 2 symbols

  - Ending delimiter is used to indicate the end of the data and the location of the FCS bytes
  - Like the SD field, ED will always be a non-data format
  - ED for a token contains 2 symbols; all other frames are 1 symbol

- Frame status field - 1 symbol (4 bits)

  - Purpose of this field is to indicate to the sending station whether the receiving station recognized its address and copied the data
  - An error bit is included in this field and is used to flag frames as containing errors

Data Encoding

- Data is encoded on the LAN using a scheme where 4 bits of information are coded and sent as 5 data bits (4B/5B)

- This implementation is used to guarantee bit transitions without a requirement for increased bandwidth necessitated when using Manchester coding schemes

- Once the 5 bit substitution is done, the data is sent on the LAN using the NRZI (non-return to zero inverted) format

- This implementation results in an actual bit rate 125% of the data rate (i.e. 125M bps for a 100M bps data rate)

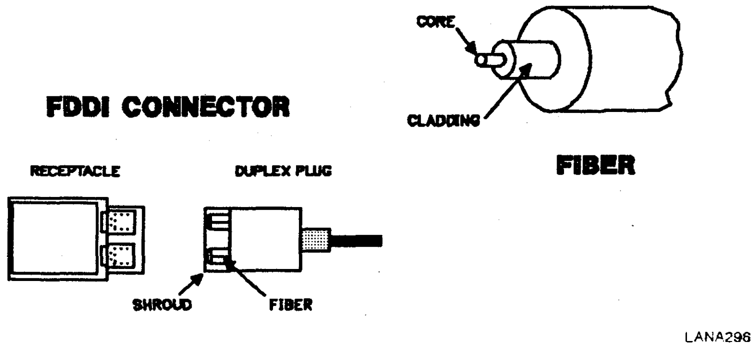

ANSI X3T9.5 FDDI Physical Network Implementations

Fiber Optic Cable

- FDDI standard recommends use of 1300 nanometer (nm) multimode fiber in one of two sizes:
  - 62.5/125 micron (most common size)
  - 85/125 micron
- Standard allows for use of two additional fiber sizes although they may limit maximum cable length between stations and maximum ring size
  - 50/125 micron
  - 100/140 micron (ICS Type 5 cable)
- Actual installation will consist of two fiber rings:
  - Primary ring is used for normal data traffic
  - Secondary ring may be used for data traffic although it is expected that initial implementations will use it for back up and recovery only (very similar to the 802.5 main ring path implementation)
- Maximum fiber length between nodes is 2 kilometers
- Maximum recommended ring length is 200 kilometers

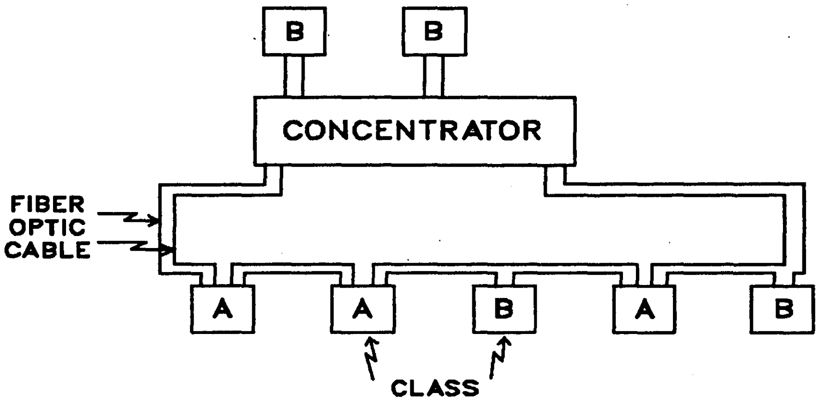

LANA297

Types of Stations

- Standard Defines two types of stations to the fiber rings, Class A stations and Class B stations.
  - Class A stations connect to both the primary ring and the secondary ring (requires two duplex plugs)
  - Class B stations connect to the primary ring only and will usually be connected through wiring concentrators
- Wiring concentrators (WC) are also defined in the standard
  - Wiring concentrators are Class A stations on the main fiber network that allow the connection of Class B stations in a star-wired ring topology
  - Class B station connection to the WC may be via a fiber cable pair or coaxial cable pair
- The connector to be used is also defined in the standard
  - Connector is called the Media Interface Connector (MIC)
  - It is a keyed rectangularly shaped connector with two fiber ferrules protected by a shroud
- When a station is not active on the ring, it will connect the incoming signal to the output (using optical bypass relay techniques)

Ring Recovery

- When a failure occurs between stations, the Class A stations adjacent to the fault (before and after) will wrap the ring onto the secondary ring (backup path) bypassing the fault
- Class B stations do not have this capability

- Class B stations between the Class A stations and the fault will no longer be part of the network until the fault is fixed
- Class B stations on the operational portion of the recovered ring will not be impacted

Feedback Questions

Directions:

First select the correct answer for each of the following questions by circling the appropriate letter. In a few minutes you will be asked to input your answers, one at a time, using the Student Response Unit.

1. The function on the Token-Ring LAN that is responsible for managing the token (e.g., generating new tokens, purging busy tokens, etc.) is called:

   a. Ring Error Monitor

   b. Token-Ring LAN Manager

   [✓ c.] Active Monitor

   d. Station Master

2. (True/False) For all three IEEE 802 MAC standards, the MAC address field must be 6 bytes long.

   a. True

   [✓ b.] False

3. Which of the following statements about LAN performance will always be true?

   a. CSMA/CD (IEEE 802.3) LANs are better than 4Mbps Token Ring (IEEE 802.5) LANs because they can operate at 10Mbps

   b. Token-Ring LANs are better than CSMA/CD LANs since they are deterministic (i.e. predictable performance)

   [✓ c.] Performance can be impacted without the user's knowledge if a poorly installed wiring system is causing MAC level FCS errors and data retransmissions

   d. Performance will not be impacted by anything at the MAC or physical levels; performance is a function of higher level software layers

4. Which of the following statement is true relative to Ethernet LANs?

   [✓ a.] There is no LLC layer in Ethernet

   b. Although there is an LLC layer in Ethernet, it is not completely compatible with the 802 standards

   c. Ethernet LANs are implemented at the physical level the same as the IEEE 802.3 LANs allowing interchangeability of media and connectors

   d. Because of the significant amount of Ethernet installations, the IEEE is expected to adopt the Ethernet LAN implementation as part of the IEEE 802.3 standard within a few years

UNIT 03: Logical Link Control Standards

Session 1 - IEEE 802.2 LLC Standards

Objectives

Upon completion of this session, you should be able to:

- Briefly explain the purpose of the IEEE 802.2 LLC layer
- State the 3 type of services supported by the IEEE 802.2 standard
- Describe the purpose of SAPs and link stations and how they relate to the LLC layer
- Interpret the meaning of the bits in the 3 fields of the LLC layer
- Identify the LLC control field commands and briefly explain the purpose of each command
- Describe how sequencing of packets is accomplished at the LLC layer
- Discuss the performance implications of the LLC layer parameter values

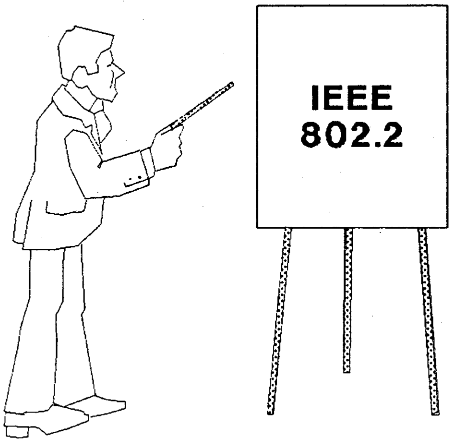

IEEE 802.2 - Logical Link Control

LLC Overview

- Logical Link Control (LLC) layer is responsible for information delivery (and delivery verification if required)
- 802.2 standard defines a common logical link protocol based on the existing CCITT High-Level Data Link Control (HDLC) protocol[8]
  - LLC protocol defines the "packet" used to transport data within a LAN independent of the type of LAN (MAC and physical layers) being used
  - LLC packet includes the data which will be visible to the higher OSI functional levels (e.g., SNA, TCP/IP, NETBIOS)
  - The LLC packet is called a LLC Protocol Data Unit (L-PDU)
- Link control parameters for each link station allow some network tuning at this level
  - For example, the number of frames to be sent and/or received prior to acknowledgement can be selected, along with the number of retries that should occur when an error is experienced
  - Timers determine how long the machine will wait for various conditions to occur
- The L-PDU is contained within the Media Access Control (MAC) frame

[8] The flag field and check fields in a standard HDLC frame are not used in the L-PDU because the MAC frame provides these services.

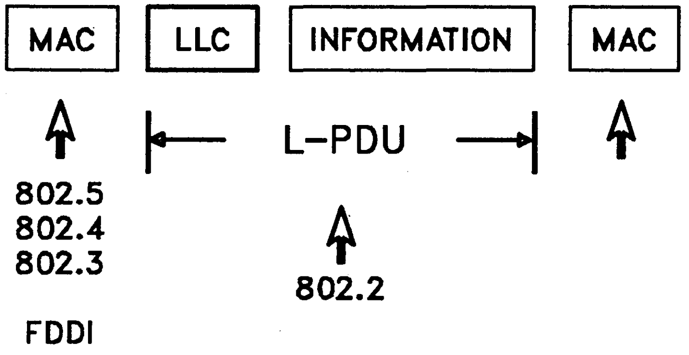

LANA301

Types of LLC Service

As part of 802.2 protocols, the IEEE has defined several classes of service.

- Type 1: Connectionless-oriented (broadcast) service
  - LLC frames are sent independently from one another without acknowledgement by the receiver; no guaranteed sequential delivery or error recovery at this level
- Type 2: Connection-oriented service
  - LLC frames are sent in sequence and acknowledgements are returned to verify correct delivery; automatic retransmission of errored frames is supported
- Type 3: Acknowledged connectionless service
  - Connectionless (datagram) type of service
  - Immediate acknowledgement of each data unit by receiver
  - Acknowledgement required by sender prior to sending next data unit

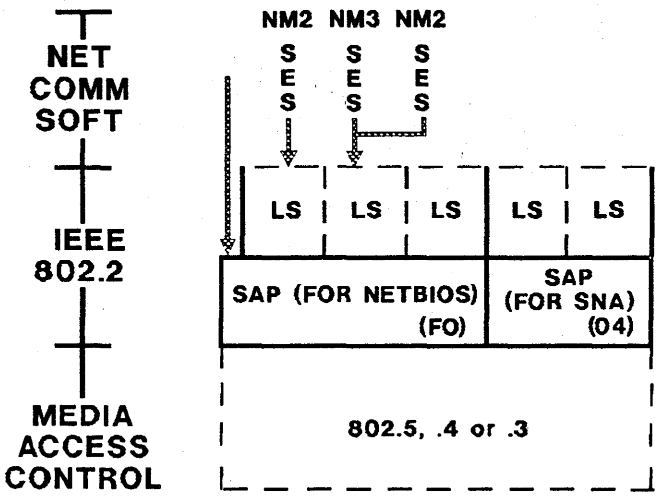

SAPS and Link Stations

Service Access Points

- Service Access Points are used by the applications (through one or more intermediate levels of software) to gain access to the network and the network to gain access to the applications[9]
- SAPs provide a method of multiplexing many logical links (sessions) on a single physical access
- Each SAP will provide support for one (and only one) communication protocol
  - Examples of protocols include NETBIOS, SNA and TCP/IP
- One protocol can use multiple SAPs to access different applications (or hosts)
  - For example, an SNA host Gateway may use multiple SNA SAP addresses (i.e. 04, 08, 0C) to access three different hosts through the same Gateway
- A total of 254 SAPs are possible although the actual number will be device dependent
- Using today's technologies, the maximum number of SAPs in a network station will be two or three

Link Stations

[9] In the visual - NM = Network Name, SES = Session and LS = Link Station.

- Link stations are used to maintain a formal connection between network nodes and are needed to support sequential delivery of frames
- Required for connection-oriented (Type 2) services only
- Link stations are associated with and used by one SAP
- Each SAP will require one link station for each unique combination of network station and SAP it will connect to
- For example, a network station with both SNA applications and NETBIOS applications communicating with SNA and NETBIOS applications in another network station will require:
  - One SAP in each network station to support the SNA applications
  - One SAP in each network station to support the NETBIOS applications
  - One link station for each SAP in each network station (total of two link stations per network station)
- One SAP can support up to 255 link stations
- Defined during initialization and allocated dynamically

SAPs, Link Stations and Sessions

- A logical connection between two network stations is referred to as a session
- In order to support a connection-oriented session to a specific application, a SAP and link station are required
- Network resources can be identified to the network using names (as is done in NETBIOS implementations)
- When a request is made for a session with a specific resource name, the session will be associated with a link station
- Requests for additional sessions from the same remote network node can use the same link station since one link station can support multiple sessions between the same two network nodes
- Refer to the visual for this example:
  - Network node has two SAPs, one for SNA applications and one for NETBIOS applications
  - NETBIOS SAP has 3 link stations associated with it
  - One link station is supporting one session to a remote node for resource NM2
  - Another link station is supporting 2 sessions to another network node, one for resource NM3 and one for NM2

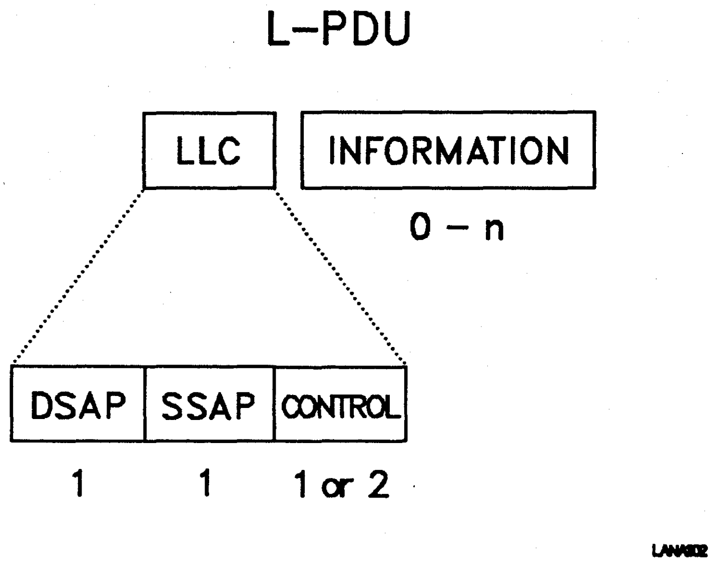

LLC Protocol Data Unit Format

Ref.: *Token-Ring Network Architecture Reference* (SC30-3374), Chapters 8 and 9.

The L-PDU consists of two address fields (1 byte each), a one or two byte control field and an optional information field. The length of the information field will be dependent on the specific LAN implementation.

- Similar to the MAC layer address protocol (which contains both a destination and source physical address), the L-PDU will contain two addresses:
  - Destination Service Access Point (DSAP) - the destination SAP of the L-PDU
  - Source Service Access Point (SSAP) - the source (origination) SAP of the L-PDU
- Some SAPs have been predefined by the IEEE and some by 'users'
  - IEEE SAP examples include:
    - Global SAP: FF
    - LLC Management SAP: 02
    - Null SAP: 00
  - User (IBM) defined SAPs include:
    - SNA Path Control: 04, 08, 0C ... (any multiple of 4 up to 'EC')
    - NETBIOS: F0
    - LAN Management: F4
    - RIPL SAP: F8
  - Because of the continuous increase in the number of applications written for LANs (i.e. 802.X protocols), the 802.2 SAP addressing had to be modified:
    - 'AA' was added to the list of defined SAP addresses
    - 'AA' is used for Sub-Network Address Protocol (SNAP) addressing

    - Presence of address 'AA' means that additional addressing will be contained following the LLC control field
    - SNAP address format will be 5 bytes long and will usually be found in Ethernet TCP/IP implementations[10]

- DSAP - Destination SAP (1 byte)
  - Purpose is to identify the receiving station's SAP being used to access a particular application
  - Low order bit (bit 7) indicates whether this is an individual or group SAP
    - Group SAPs are used to address multiple SAP addresses in a common group (bit 7 = 1 for group addresses)
  - Bit 6 indicates whether the SAP address is IEEE defined (bit 6 =0) or user defined (bit 6 = 1)
  - Remaining bits make up the actual address

- SSAP - Source SAP (1 byte)
  - Purpose is to identify the sending station's SAP being used by a particular application
  - Low order bit (bit 7) indicates whether this is a command or response L-PDU (not considered part of the address)
  - Bit 6 indicates whether the SAP address is IEEE defined (bit 6 =0) or user defined (bit 6 = 1)
  - Remaining bits make up the actual address

- Control Field (1 or 2 bytes)
  - Purpose is to control the flow of information at the LLC level, including confirmation of delivery and packet sequencing

  - Field contains LLC commands and responses, sequence numbering (as appropriate), and the poll/final bit

  - Size of the field will be determined by the specific command being transferred

    - All commands which do not contain sequence numbers will be 1 byte in length; includes all data transfers in connectionless services mode
    - All commands using sequence numbers will be two bytes in length

    Note: Two bytes are required for sequencing because the LLC protocol is based on HDLC, modulo 128 format.

- Information field (0 to n bytes)
  - Optional field, but if present, data must be in octet format (multiples of 8 bits)
  - Purpose is to transport information from the source to the destination
  - Maximum allowable length of the I field will be determined by the specific LAN implementation
    - For example, most 802.3 networks will allow 1500 bytes; 802.5 networks operating at 4M bps will allow 4000 bytes
  - Actual I field length will be determined by node resources (e.g., size of transmit buffer, tuning parameters)

[10] SNAP addressing will be covered in the unit on TCP/IP.

- Information types:
  - User data

    Note: Format of the user data will depend on the protocols being used above the LLC layer. Formats include SNA, NETBIOS, TCP/IP and various vendor proprietary implementations.
  - LLC information such as XID parameters or error data

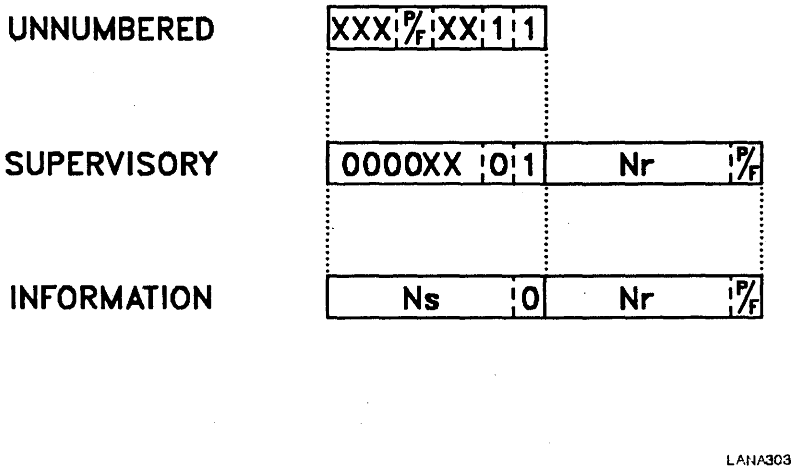

Control Field Formats

Control field formats will be dependent upon whether the nodes are using connectionless or connection oriented services. Before we take a look at specific control field formats, there are some common bits in all frames that can be discussed. Before we do that, let's take a look at the types of control field formats.

There are three control field types:

- Unnumbered frames
  - Used for link management (connection establishment, link test and error recovery)
  - Also used for information flows in connectionless services
  - Always one byte in length
  - Identified by the 2 low order bits of the control field (bits 6 & 7) being a 1
- Supervisory frames
  - Connection-oriented services only
  - Used to assist the flow of information between two link stations
  - Provides verification of delivery of frames and request for retransmission of frames received in error
  - 2 bytes in length (modulo 128)

    Note: Modulo 8 would use a single byte control field.
  - Identified by bits 6 & 7 of byte 0 in two byte control fields; bit 6 = 0 and bit 7 = 1 respectively
- Information transfer frames
  - Connection-oriented services only

- 2 bytes in length (modulo 128)

  Note: Modulo 8 would use a single byte control field.
- Identified by bit 7 of byte 0 being set to 0

- In addition to the bits that indicate the frame format, there is one other bit that is always present in a control field: the P/F bit
  - Poll/Final (P/F) bit can be used to control flow of information between two network nodes
  - Bit 3 of 1 byte control fields and byte 1 bit 7 of 2 byte fields
  - Bit is called the 'poll' bit in command frames (sender to receiver) and the 'final' bit in response frames (receiver to sender)
  - Frame direction (command or response) determined by the low order bit in the SSAP field
  - When the P/F bit is on in a command frame, the receiver must respond to the frame
  - Receiver will set the P/F bit based on the command being received along with other considerations
  - When the bit is off, receiving stations will usually not send a response until some other action occurs such as a timer tripping or receipt of the maximum number of frames agreed upon (window size)

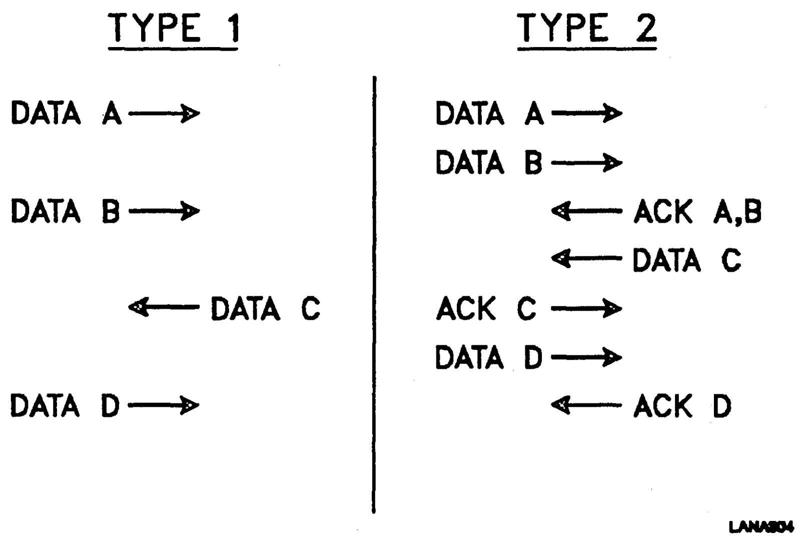

Connectionless Services (Type 1) Formats

- General rules for connectionless services:
  - Network node requires a SAP but no link station
  - Only three formats are used to support all activity on the link; no mode setting commands are used
- XID (Exchange Identification) - C/R[11]
  - At the link level, the XID is used to exchange characteristics of the sending and receiving stations
  - 802.2 XID will consist of a 3 byte I field used to identify the type of service supported and the receive window size to be used by each station
  - Higher level XID formats can also be exchanged, such as those used in an SNA link establishment; SNA exchanges would include network node identifiers such as product ID and PUID
  - I field size and content will be implementation dependent
  - Only valid response to an XID is an XID
- Test - C/R
  - Used to test the station to station transmission path
  - An I field may be included in the command; if it is, the receiver must return the same I field in the TEST response
  - Only valid response to a TEST command is a TEST response

Note: XID and TEST commands are used by SNA network nodes to establish link connections across bridges (command/response flows will collect 'source routing' information).

[11] C/R identifies whether the format is a command (C), a response (R) or both (C/R).

- UI (Unnumbered Information) - C
  - Used to transport unsequenced data

Connection-Oriented Services (Type 2) Formats

- General rules for connection-oriented services
  - Network nodes require both SAP and link stations to support the connections
  - Information frames are assigned sequence numbers which are used to confirm delivery
  - Link stations maintain information about the connection, including the current send and receive sequence numbers

Prior to transferring information using Type 2 services, a logical connection must be established between two network nodes. When the connection is established, each node will assign a link station to the logical connection. A number of commands are used to establish and maintain this connection.

- SABME (Set Asynchronous Balanced Mode Extended) - C
  - Used to establish a logical connection with a remote station using Type 2 services[12]
  - Receipt of a valid response (UA) will result in the sequence counters being set to '00'
- UA (Unnumbered Acknowledgement) - R
  - Used as an acknowledgement to the SABME command
  - Also used as an acknowledgement to the DISC (Disconnect) command
- I (Information transfer) C/R
  - Used to transfer data between network nodes
  - 2 byte control field includes both send (Ns) and receive (Nr) sequence number counters
  - Each counter (Ns and Nr) is 7 bits long allowing for a maximum count of 128 (0 to 127); hence the term modulo 128
  - Counters are used to guarantee sequential delivery of the information, and by the receiver to request retransmission of frames received in error
  - Counters are reset to '00' by during link initialization (SABME)
  - A receiving station can respond to a command I frame with a response I frame (although not full duplex, data can go both ways during on a link connection)
- S (Supervisory) C/R
  - Used to perform supervisory functions such as acknowledgement of information frames, request for retransmission of information frames and to maintain link synchronization
  - When a receiving station is required to respond (received a frame with the P/F bit on, maximum receive window reached or timer has expired) it will send a Supervisory frame if it has no data to send
  - S frames contain Nr counter values only

[12] Somewhat equivalent to an SDLC Set Normal Response Mode (SNRM) command.

  - S frames will be passed back and forth periodically between two stations that do not have any information to transfer

- Other C/R formats used to support Type 2 connections

  - XID: Can be used prior to the establishment of a Type 2 connection to validate the identity of the stations and negotiate link level parameters

  - TEST: Can be used prior to link establishment to check the link level connection

  - FRMR (Frame Reject) is used during the connection when a severe link level error has occurred such as incorrect frame sequencing or receipt of an invalid command

  - DISC (Disconnect) is used to disconnect the link connection

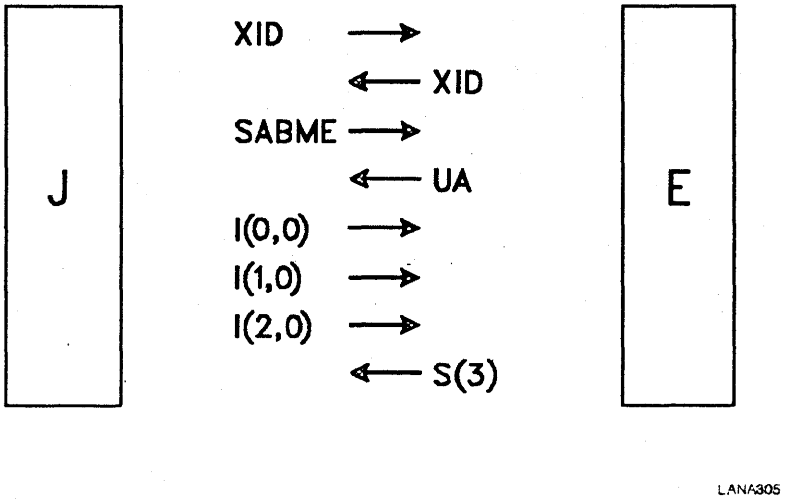

Frame Sequencing and Acknowledgement

The send (Ns) and receive (Nr) sequence numbers allow the link stations to verify the sequencing of frames and to support retransmission of frames received in error. An understanding of how these counters work will help you understand the concept of link level verification and the concept of windowing or pacing.

In this scenario, we have two stations, Jake and Elwood (in honor of those great musicians, the Blues Brothers).

1. Jake and Elwood exchange IDs using XIDs where they agree that each has a receive window size of 3 frames.
2. Jake establishes the connection with Elwood using the SABME and Elwood responds with a UA (both stations now have a count of '00').
3. Jake sends the first I frame (Ns=0, Nr=0).

   Note: The reason that the Ns count is sent as '0' is that Jake does not update his Ns counter until after the frame has been sent.
4. Elwood receives the first frame and updates Nr=1.
5. Jake sends the next frame (Ns=1, Nr=0).
6. Jake sends the next frame (Ns=2, Nr=0); after sending the frame, Jake's counters will be Ns=3, Nr=0.
7. Since Elwood has received the frames without errors, his Nr=3.
8. Elwood has a receive window of 3, and since he has now received three frames, he will respond. He can send an S frame (Nr=3) or return an I frame (Ns=0, Nr=3).
9. Jake can now send more frames to Elwood.

10. If there had been an error in one of the frames going to Elwood, Elwood would have stopped listening and counting and not respond until either a receive time out occurred, or Jake asked for a response using the P/F bit. When Jake responded, the Nr count would indicate the last good frame received prior to the error; Jake would then resend the information, starting with the error frame.

11. The information transfer would continue until the counters reach 127 when they would roll over to 00 and start all over.

12. Although it may not be obvious to you yet, the implementation of windows and timers to support pacing can have a significant impact on network performance.

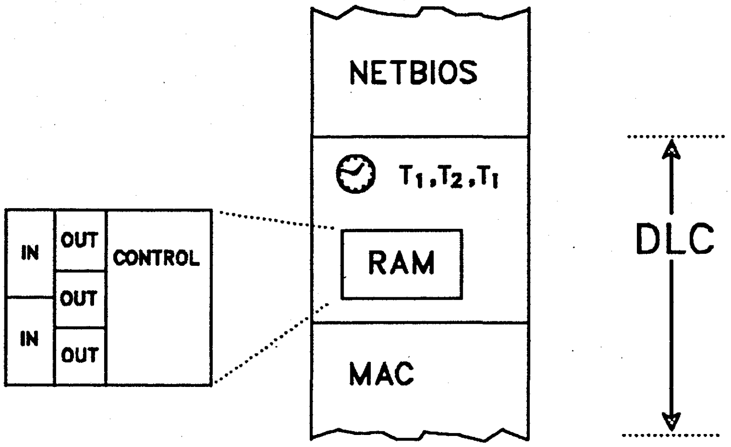

LANA306

Data Link Control (DLC) Interface Parameters

When the higher levels of software issue an open command to the LAN adapter, the command will include a number of parameters which will determine the environment in which the adapter will operate. The software issuing the open may choose to use existing defaults designed into the products, or may choose to alter the parameters. In many cases, the LAN adapter card resources (e.g., RAM buffer size) will impose limitations on the performance of the LAN. One of the major purposes of LAN tuning is to select the right parameters for optimum performance at the DLC level.

Examples of DLC Parameters

- SAP IDs
- Timers (transmit timeout, receive timeout, inactivity)
- Transmit and receive window sizes
- Transmit and receive buffer sizes and number of buffers
- Retry count

Some Things That Can Impact Performance at the Link Level

The intent of the following information is not to make you tuning experts. It is to alert you to the fact that architecturally, there is designed into most adapters the ability to control link level flows. Performance can be enhanced, or negatively impacted, by selection of various values during the set up of the adapter operating parameters.

- Transmit window size (also called MAXOUT)

  - Determines how many frames will be sent prior to the sender turning on the P/F bit and waiting for a response
  - The larger the value, the more adapter buffer space required to store frames for possible retransmission
  - If the value is too large, you could experience delays during error recovery and data retransmissions and could cause overruns at the receiving end

- Receive window size (also called MAXIN)
  - Determines how many frames will be accepted by the receiver prior to sending an acknowledgement
  - Larger window size requires more adapter buffer space
  - Must be equal to or less than MAXOUT at sender's end; otherwise, time-outs can occur

- There are also several timers involved
  - Transmit timeout (T1) timer determines how long a sending station will wait for a response prior to resending a frame
  - Receive timeout (T2) timer determines how long a receiving station will wait before acknowledging a received frame, if the receive window size has not yet been reached
  - Inactivity timeout (Ti) timer determines how long a station will wait before soliciting the status of a remote link station

Feedback Questions

Directions:

First select the correct answer for each of the following questions by circling the appropriate letter. In a few minutes you will be asked to input your answers, one at a time, using the Student Response Unit.

1. (True/False) The LLC control field is not used for connectionless datagram services.

    a. True

    [✓ b.] False

2. Which of the following is a true statement about SAPs and link stations?

    a. One link station is required for each active session in a network node

    b. SAPs are required for SNA implementations, but are not needed for NETBIOS or TCP/IP

    c. SAPs are only required for connection-oriented services

    [✓ d.] Link stations are only required for connection-oriented services

3. Which LLC control field is used prior to connection establishment to validate the identity of the calling station?

    a. TEST

    b. UA

    [✓ c.] XID

    d. Supervisory (S)

UNIT 04: NETBIOS and SMB

Session 1 - NETBIOS

Introduction

The 802.2 LLC layer provides a consistent interface to the layers above it (OSI layers 3 through 7). Layers 3 through 5 of the OSI model provide network, transport and session services respectively. Starting with layer 3 and working up to layer 7 we find a variety of implementations, most of which are vendor specific (proprietary) protocols.

One way of implementing OSI layers 3 through 5 is through one (or possibly two) program products; another way is to incorporate the functions into the application package.

Even within IBM, we have multiple ways of communicating across LANs using different protocols.

- NETBIOS (currently implemented using the LAN Support Program or the OS/2 Communications Manager)
- SNA (implemented using VTAM, VTAM and NCP together, microcode in specific products or a 3270 emulation program)
- TCP/IP (implemented using specific program products or the OS/2 Communications Manager)

Each of these implements a unique way of communicating over the token-ring LAN so that applications written to one interface can not communicate directly with applications written to a different interface. For example, NETBIOS applications can not communicate directly with an SNA host. Gateway products must be used to perform the protocol conversion from one network protocol to another.

Objectives

Upon completion of this session, you should be able to:

- Describe NETBIOS
- Identify the services provided by NETBIOS on a LAN and explain each
- State the purpose of an SMB and its relationship to NETBIOS
- List the IBM products that implement NETBIOS

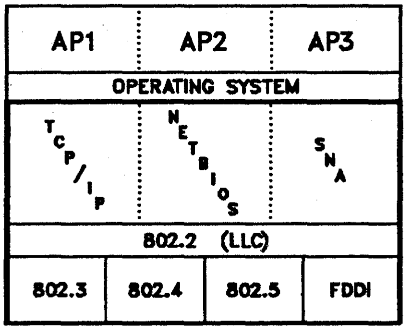

LANA235

What is NETBIOS?

- NETBIOS is an acronym for Network Basic Input/Output System
- NETBIOS was originally developed for the IBM PC Network (announced in 1984)
- Because of its relative ease of implementation, it quickly became very popular and is now a de facto standard[13]
- Provides a high level software interface to allow applications running on different computer systems to communicate
- Implements functions similar to the functions of levels 3 through 5 of the ISO OSI model

  Note: Although a comparison is made here between NETBIOS and the OSI layers, it is important that you fully understand that NETBIOS *is not compliant with and has not been accepted as* an OSI standard.
- Allows applications to communicate to other applications using network (logical) names (NETBIOS Name Service function)
- NETBIOS also provides session level services (Session Service function)
- Application interface is the same regardless of the type of LAN being used (applications written to NETBIOS should run on any LAN using NETBIOS communication)

[13] The term "de facto standard" is used here to indicate that, although the protocol is not part of any documented standard, its wide acceptance and use in the marketplace positions it as an industry standard.

- Supports connection-oriented (Type 2) services and connectionless (Type 1) datagram services
- NETBIOS uses one SAP address ('F0') and can coexist with other network software protocols using different SAPs (e.g., SNA)

NETBIOS and bridges:

- NETBIOS implements source routing
- Broadcasts are sent with the SA field high order bit set to 1
- Source Routing field indicates single-route broadcast (no route information recorded during search)
- Target station(s) respond with an all-routes broadcast and so responses received by the source will contain route discovery information

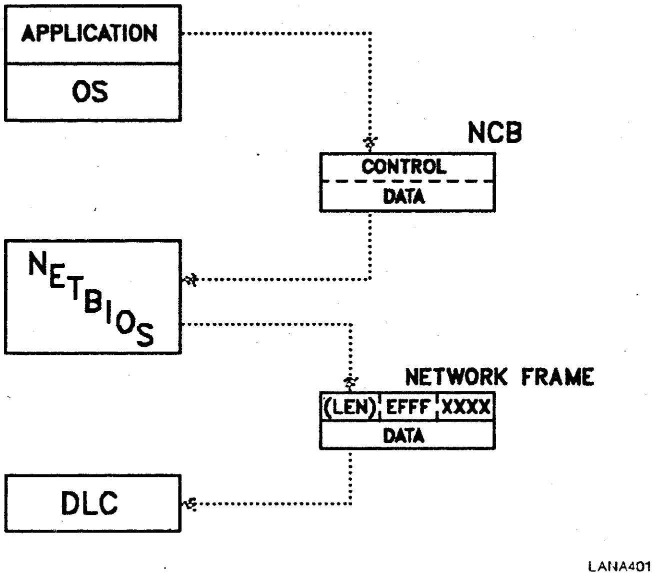

NETBIOS NCBs

Note: Details of the NETBIOS Network Control Blocks (NCBs), commands and formats and protocols will be found in the *Local Area Network Technical Reference* manual (SC30-3383).

- NETBIOS application interface is implemented using a format called the Network Control Block (NCB)
- NETBIOS supports 24 commands used by the application to start the network, establish sessions, transfer data, etc.
- Size and specific content of the NCB will be command dependent
  - For example, when an application needs to establish a session with a remote resource, it will first issue a command called NCB.FIND.NAME
  - Contents of the NCB will include information such as which adapter should be used and the specific names being sought

NETBIOS Formats & Protocols

- NETBIOS converts the NCB to a network frame and delivers it to the LLC (i.e. 802.2) sublayer within the DLC layer
- Network frame will always have a header field followed by data (from the NCB)
- Header format:
  - 2 byte length field
  - 2 byte NETBIOS format identifier X'EFFF'

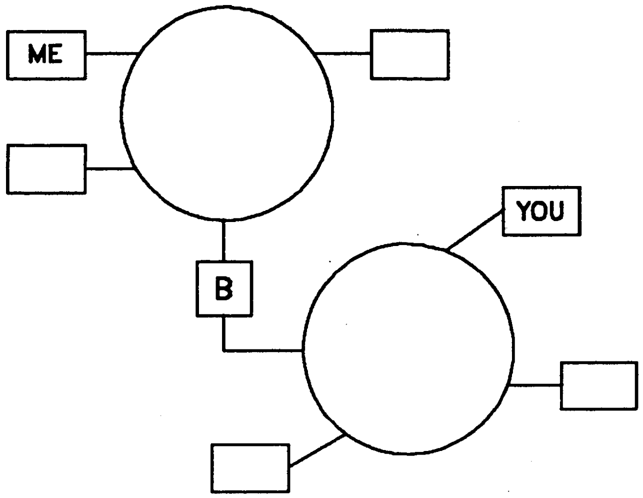

LANA403

NETBIOS Services

Name Service

The NETBIOS Name Service function allows applications to communicate across the network to other applications through the use of names without any knowledge of the underlying addresses or network implementations. Since names are not permanently associated with any address, flexibility is provided to the users.

- NETBIOS names must be unique on the network, not just on a given LAN segment (does not apply to group names)
  - When a name is added to the network (e.g., the application is adding another resource), NETBIOS will broadcast the name across the network to make sure that it is unique; any responses represent duplicate names
- Names are kept in a table so they can be found by other network nodes
- Group names can be used to address multiple nodes simultaneously

The Name Service function performs address resolution as follows:

1. NETBIOS receives the name request from the application
2. NETBIOS sends a broadcast to all stations with NETBIOS active
3. MAC frame destination address is the functional address for NETBIOS (C00000000080) and the imbedded NETBIOS information contains the target network name
4. All stations in the network with NETBIOS active will copy the frame
5. Station which recognizes the name (target) will respond with an acknowledgement frame

6. When a response is received by the sender, the sender can correlate the MAC physical address with the target name
7. A session is established between the two names and will remain active until terminated (NETBIOS Session Services)

Session Service

- NETBIOS provides support for reliable data transfers between two network names across a logical (or virtual) connection[14]
- Sessions are associated with names; one name can support multiple sessions through the use of session identifiers
- NETBIOS will assign sequence numbers to data blocks for delivery confirmation and error recovery at the session level
- An application can send up to 64K bytes in one block; NETBIOS will segment the data into block sizes based on the agreed upon parameters established at session initialization time with the remote node
  - When a session is established between 2 nodes, the maximum data frame size will be determined by the sender's transmit buffer size, receiver's buffer capabilities and any limitations imposed by bridges
- NETBIOS supports 802.2 flow control using pacing (windows)
- NETBIOS will use the 802.2 Type 2 service to support session service

  Note: An additional level of block sequencing (i.e. Nr and Ns counters) will occur at the 802.2 level

Datagram Service

- NETBIOS Datagram Service allows messages to be sent without first having to establish a session
- Messages can be sent to individual or group names, or broadcast across the entire (NETBIOS) network
- No acknowledgement is received for these messages

NETBIOS and the DLC Interface

The DLC interface allows NETBIOS to alter the parameters used by the LAN adapter and supporting DLC code (802.2 in most cases). We have already seen an overview of these parameters in Unit 3. Some network (i.e. layers 3 through 5) software does not give the user much flexibility in customizing the network for individual requirements. However, NETBIOS provides the user with access to 25 alterable parameters; if the user doesn't choose to set them, NETBIOS will supply default values. These parameters passed by NETBIOS to the DLC and adapter interfaces can have a significant impact on the operation of the node.

[14] Do not confuse the term 'logical' here with SNA logical unit (LU) implementations; NETBIOS is not an SNA product.

NETBIOS and Other Vendors

IBM implements the NETBIOS interface along with all of the underlying services provided by the complete protocol suite. Other vendors, recognizing that NETBIOS is a de facto standard, have implemented support for the NETBIOS interface, but are not using all of the services. These vendors will allow the user (application) to build the NCB according to the NETBIOS interface standard. However, these vendors will then transport this information using some other set of network protocols. These protocols may be vendor specific (e.g., Novell NCP) or a different standardized protocol stack such as TCP/IP. Implementations such as this expand the interoperability of the network systems.

IBM Products That Implement NETBIOS

- IBM LAN Support Program
- OS/2 Communications Manager (V1.1 and V1.2)

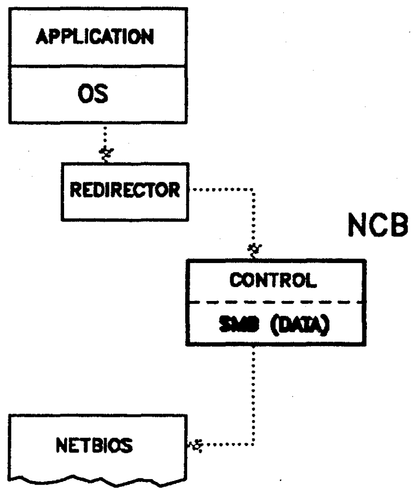

LANA404

Session 2 - Server Message Blocks (SMB)

Introduction

NETBIOS provides a network software interface to the user or application; it does not provide any application or presentation layer services. A number of vendors, including IBM, have implemented a presentation layer NETBIOS compatible protocol that provides the application with a very valuable service: LAN Server/Redirector functions. These services are implemented through the use of server message blocks (SMBs). Although completely separate from NETBIOS (SMBs can be used with other network software interfaces), most current LAN implementations will use the NETBIOS interface.

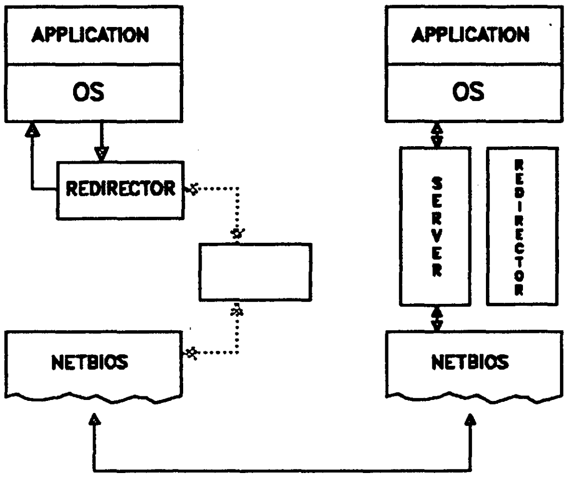

LANA405

Server/Redirector Functions and SMBs

Before we take a look at what an SMB is, we must first look at how the redirector function works.

- DOS was originally designed to operate in a single user environment
- When version 3.0 of DOS was announced, it included the ability to share resources known to DOS with other users (i.e. remote nodes)
- Some function was then needed to decide whether an application request for a resource should be directed to the local DOS resource, or if it should be 'redirected' to a network resource (that's where the name 'redirector' comes from)
- Redirected functions were packaged into another protocol (SMBs) and sent to the network software (i.e. NETBIOS)
- At the receiving end of the network, the SMB is passed to the 'server' portion of the code, unpackaged and sent to the appropriate DOS resource in that node
- Output generated by the request will be redirected back to the original requester in another SMB

What exactly is an SMB

- An SMB is a protocol that supports communication between a redirector function in one network node and a server function in another network node
- It is currently implemented using NETBIOS as the network software; however, since the SMB appears to be just plain vanilla data to the lower (network) software layers, it is technologically feasible to use other network implementations (e.g., TCP/IP or vendor specific protocol)

- SMBs can be divided into four main categories:
  - Session SMBs used to control session establishment and termination
  - File access SMBs used to send file control commands across the network
    - File control commands include opening, closing, reading, writing, creating, deleting and locking files (or portions of files)
  - Print control SMBs for sending data to network print queues
  - Message SMBs used for sending and forwarding messages through the network
- Redirectors implementing the SMB protocol (e.g., IBM LAN applications or the Microsoft Redirector) must communicate with a server that also implements the SMB protocol

IBM Products That Implement SMBs

Exercise 2

Please go to Appendix B and complete Exercise 2.

Feedback Questions

Directions:

First select the correct answer for each of the following questions by circling the appropriate letter. In a few minutes you will be asked to input your answers, one at a time, using the Student Response Unit.

1. NETBIOS provides which of the following services

   a. Name Services

   b. Session Services

   c. Datagram Services

   [✓ d.] All of the above

2. Which of the following is a true statement relative to NETBIOS?

   [✓ a.] Allows applications to communicate to other applications using network names

   b. Must be configured with the correct network address prior to use

   c. Uses a central name server to associate names with network addresses

   d. Allows an SNA application to communicate on a LAN

3. (True/False) IBM is the only vendor currently using SMBs to implement server/redirector functions?

   a. True

   [✓ b.] False

UNIT 05: SNA LAN Implementations

Session 1 - Implementation of SNA on LANs

Overview of SNA

SNA has gone through some dramatic changes since its inception in 1974. When first announced, it was a single host, hierarchical-based architecture. Today, it supports multi-host and multi-network connections, and allows both vertical and horizontal network integration. What does all of this have to do with LANs you ask yourself? Basically, LANs provide the physical and data link services for SNA, complementing the WAN implementation using SDLC and traditional telecommunications network schemes. LANs do not alter the SNA logical connectivity in any way; they merely provide an alternative way of getting from one logical unit to another.

Before we look at how LANs provide support for SNA, we need to take a brief look at how SNA networks can be implemented. For the purposes of this course, we have divided them into three categories.

Objectives

Upon completion of this session, you should be able to:

- Briefly describe how SNA supports hierarchical and peer-to-peer networking
- State how SNA can be implemented on LANs
- Describe how specific IBM hardware and software products are implemented on LANs.

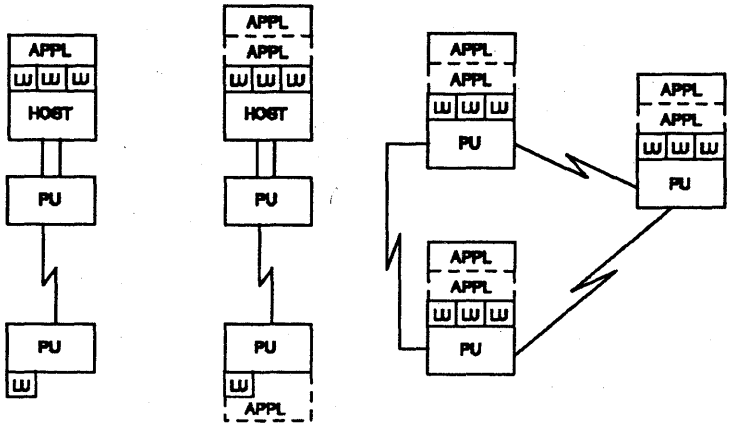

SNA Network Implementations:

- Traditional hierarchical implementation
  - Network implementation where all network flows are controlled by one or more hosts (domains)
  - LUs (applications) at the host were designated as primary LUs and end users or remote applications were secondary LUs; primary LUs controlled the logical connectivity
  - Supporting physical unit (PU) services also implemented in a hierarchical fashion
    - PU T4 controlled the PU T2.0s in the network
    - PU T2.0s could not communicate directly to other PU T2.0s
  - This is what we called a vertically integrated network
- Advanced Program to Program Communications (APPC)
  - An LU (called type LU 6.2) interface that allows connection of LUs (applications) together on a more equal basis
  - In many cases, either LU can be the primary and the other the secondary; decided through negotiation
  - Might be called 'logical peer to peer networking' (you won't see this term anywhere but here)
  - APPC applications can be implemented on either vertically or horizontally integrated networks
- Low Entry Networking (LEN)
  - A physical network implementation that allows multiple mid-range (and now large host) systems to be tied together in a physical network where each node participates on an equal basis with other nodes at the PU level

- In a LEN implementation, network control functions are distributed among the nodes and the PU services are enhanced to support the connectivity requirements (PUs are T2.1 and can communicate directly to other PU T2.1s)

  Note: Advanced Peer-to-Peer Networking (APPN) is an enhancement to the LEN implementation providing a wider range of functions.
- Horizontally integrated network

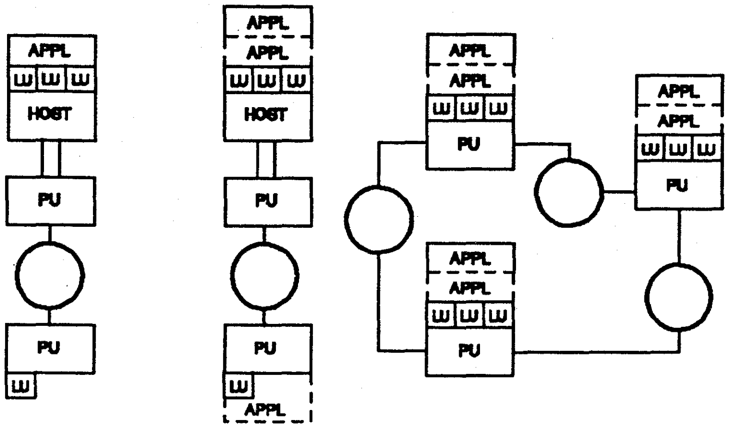

LANA502

SNA and LANs

LANs can be used to replace the WAN connection for SNA nodes which support this configuration. Since the LAN is only replacing the physical and DLC layers, there is no impact on the higher level SNA functions. The ability of the attaching product/system will be determined by that specific product's capabilities.

SNA PIU Enveloping

The SNA path information unit (PIU) consists of three distinct pieces of information.

- Transmission Header (TH) which contains the SNA addresses, pacing and routing information and other information that enables SNA to route the information from its source to its destination
- Request/Response Header (RH) which identifies the direction of the flow, whether or not a response is required and the type of information contained in the RU
- Request/Response Unit (RU) which contains network control commands (used for session establishment, maintenance and termination) or user data
- The TH, RH and RU will be passed to the DLC layer on a LAN; there it will be placed in an LLC envelope (called an L-PDU) (SNA loses its identity at this level; it's just another 'I' field)
- From there, the L-PDU will be passed to the MAC layer for transmission across the LAN

  Note: The only currently supported MAC implementation for SNA networks is the 802.5 Token-Ring

A Comparison of SNA and NETBIOS

- Like NETBIOS, SNA resources are known by their network (LU) names on the network
- Unlike NETBIOS, the SNA architecture requires all of these network names to be mapped to specific addresses, both logical and physical (where NETBIOS dynamically translates a network name to a physical address at the time of session establishment, SNA requires that all of these correlations be previously defined)
- NETBIOS does a (NETBIOS) broadcast search to locate the network name it is seeking; the MAC frame includes routing information designating the frame as a single route broadcast
- SNA, although it already knows the destination address of the target node, must also send a search out to determine whether the target is on the local ring (this will be done using either an LLC XID or TEST frame); the frame is sent without routing information the first time, and if no response is received within a given time, the frame will be resent with routing information indicating an all stations broadcast

IBM SNA LAN Products

At the end of most sessions in this course, a list is given identifying the IBM products which implement the architectures and protocols discussed in that session. Because of the extensive support for LANs in the SNA products, and the range of connectivity alternatives, we will go into a little more detail here about how the SNA products fit into the LAN architectures.

- VTAM (V2.X, V3.X)
  - As a general rule, VTAM supports LAN attachments transparently; that is, the physical LAN network implementation is unknown to VTAM
  - Level of VTAM determines level of services provided (e.g., earlier levels of VTAM were strictly hierarchical; support for LEN network attachments is now possible)
  - How VTAM sees the device attachment will be determined by the type of Token-Ring Gateway being used:
    - Devices (PUs) attached through a 3745 NTRI attachment will appear as switched (dial in/out) nodes (applies to both local and remote 3745s)
    - Devices (PUs) attached through a 3174 X1L will appear to be channel attached (local) devices
    - PUs attached through a remote 3174 Gateway will appear to be multidrop leased line connections through the local 3745
    - In a 9370 configuration with a direct attached ring, PU T2.0 devices on the ring are treated as switched nodes; other 9370s (PU T5 devices) are configured as peer devices
- NCP (Gateway)
  - Minimum NCP level required was V4.2 for 3725s and V5.0 for 3745s; it supported strictly hierarchical connections (PU T2.0) for LAN devices
  - Latest level, NCP V5.3, supports both PU T2.0 and PU T2.1
    - PU T2.0 nodes are used for the host (VTAM) connections
    - PU T2.1 nodes on the ring can be configured as APPN nodes in the NCP and the NCP can be used as a connection point to an APPN network
- 3174
  - All SNA and LAN support code is implemented in the 3174 microcode (Licensed Internal Code)
  - 3174s are currently implemented as PU T2.0 nodes only
- 3270 Emulators
  - All 3270 emulation programs that support LAN connections come with the SNA support for this connection; they need external device drivers for the DLC layer implementation (e.g., LAN Support Program)
  - 3270 Emulation Program V3.0 and Workstation Program V1.1 support PU T2.0 functions only
  - Personal Communications/3270 in a Network Station configuration works the same as the emulators above
  - Personal Communications/3270 in a Gateway configuration will communicate to the host as a PU T2.0, but will appear to be a PU T4 to the downstream PUs it is supporting
  - OS/2 V1.2 using the 3270 emulation mode will operate the same as Personal Communication/3270

  - OS/2 V1.1 will operate the same as Workstation Program V1.1

- AS/400
  - Workstations on the ring can communicate to the AS/400 in either PU T2.1 mode or PU T2.0 mode depending on the specific code running in the workstation
  - AS/400s communicate to other AS/400s in PU T2.1 mode
  - AS/400s communicating to a 370 host in 3270 Emulation mode will operate in PU T2.0 mode

- APPC/PC (DOS)
  - This program when loaded into a PC (or PS/2) will support either PUT2.0 or PU T2.1 depending on the configuration
  - Program allows an application to communicate to another application using an LU 6.2 interface
  - OS/2 V1.2 and V1.1 contain the same code functions as APPC/PC

- System/88
  - Using the appropriate software in the S/88, it can participate in both hierarchical and peer networks, and supports APPC applications

Feedback Questions

Directions:

First select the correct answer for each of the following questions by circling the appropriate letter. In a few minutes you will be asked to input your answers, one at a time, using the Student Response Unit.

1. (True/False) SNA network nodes can use either an IEEE 802.3 CSMA/CD or 802.5 token-ring LAN for communications.

   [✓ a.] True

   b. False

2. Which of the following is NOT a true statement relative to the implementation of SNA on Token-Ring LANs?

   a. SNA can use the Token-Ring to support both hierarchical and peer-to-peer networking

   b. SNA PU T2.1s can communicate to other PU T2.1s on a Token-Ring LAN

   [✓ c.] The SNA path information unit (PIU) format is altered slightly when using a Token-Ring network connection

   d. SNA can coexist in the same machine as, but can not communicate directly with, NETBIOS

3. Of the following products which implement SNA, one of them can not use a LAN for network communication. Which one is it?

   a. VTAM V3.1.2

   b. NCP V5.1

   [✓ c.] 3274 Cluster Controller

   d. PS/2 running OS/2 V1.1

UNIT 06: TCP/IP LAN Implementations

Session 1 - TCP/IP Overview

Introduction

As was pointed out in the previous unit, Transmission Control Protocol/Internet Protocol (or TCP/IP as it is usually called) represents one way of implementing the layer 3 and 4 functions. The IP portion is a network layer implementation and TCP a transport layer implementation. Session layer functions are the responsibility of the application using TCP/IP.

Objectives

Upon completion of this session, you should be able to:

- Briefly describe the evolution of TCP/IP and state why it is so popular today
- State the relationship between the TCP/IP layers and the OSI 7 layer network model
- Explain internetwork addressing as it applies to TCP/IP
- Discuss IP routing techniques and its uses of gateways.

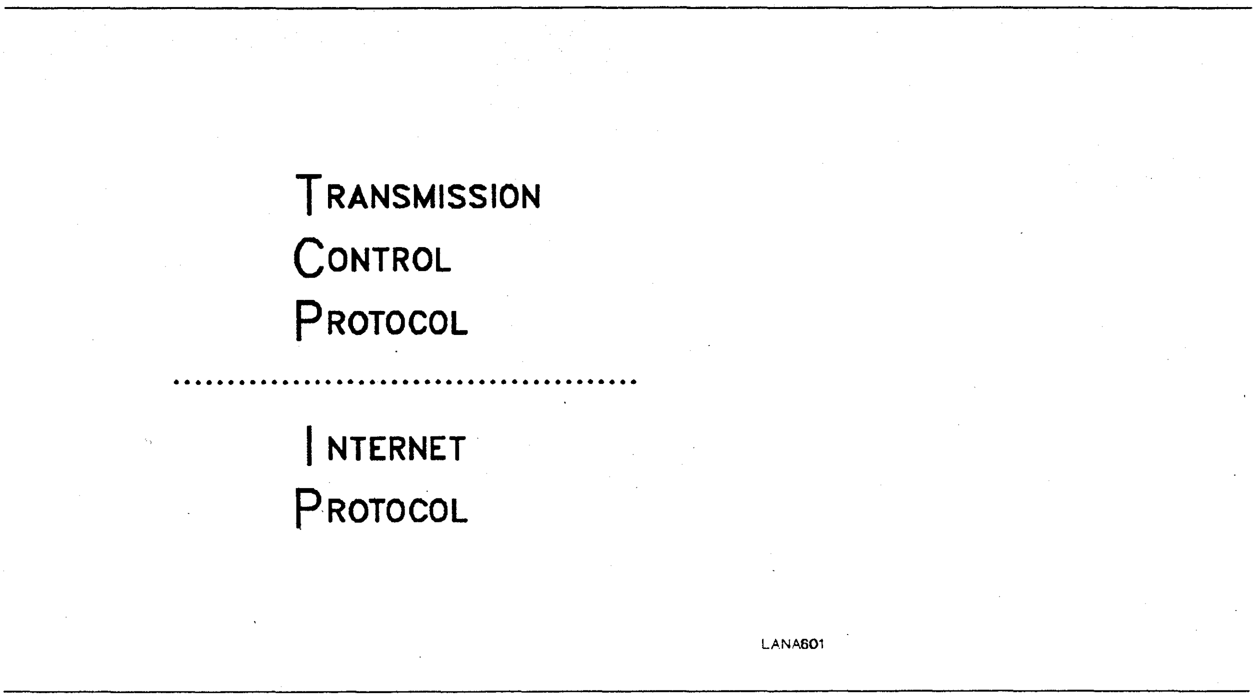

TRANSMISSION
CONTROL
PROTOCOL

INTERNET
PROTOCOL

LANA601

What Is TCP/IP and Where Did It Come From?

- TCP/IP is a result of work begun in the early 70's by the Defense Advanced Research Projects Agency (DARPA) along with a number of vendors
- DARPA needed a set of protocols that would provide transparent network services to UNIX based applications for the Advanced Research Projects Agency's network (known as ARPANET)
- TCP/IP as we know it today was officially implemented starting in the late 70's
- DARPA also funded the implementation of TCP/IP; the University of California at Berkeley was commissioned to distribute the code along with the UNIX operating system
- The result was the very quick proliferation of the code throughout the university and scientific communities; TCP/IP has become the de facto standard for communication subsystems operating in UNIX environments

Request For Comments

- TCP/IP is a dynamic set of protocols; that is, it is continuously being updated and expanded
- The update process is called 'request for comments' (RFC); it allows users to recommend changes and additions to the protocol
- RFC activity is coordinated through the Network Information Center (NIC) in Menlo Park, Ca. (NIC will provide a complete list of RFC numbers along with their descriptions upon request)

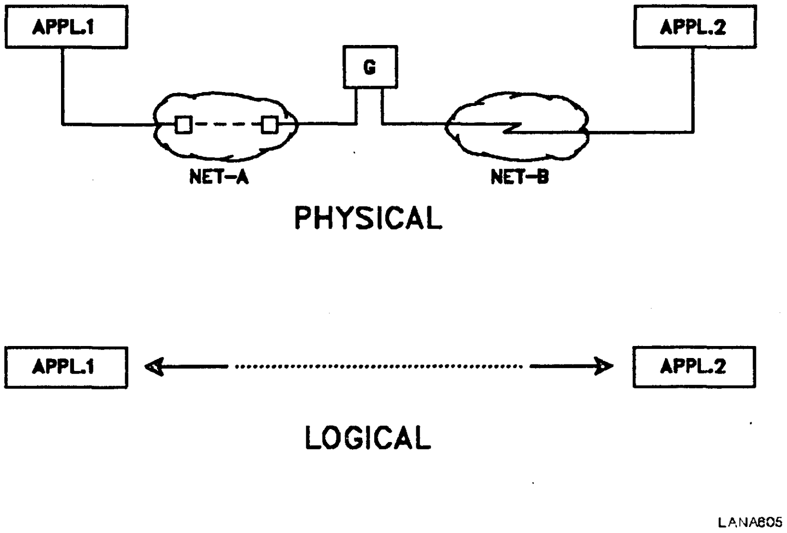

TCP/IP Architecture

Ref.:*TCP/IP Tutorial and Technical Overview*, Chapter 2

Internetworking

- One of the design points of TCP/IP was universal communications services, or internetworking as it is called
  - In other words, the set of protocols, implemented through software, must provide networking services to the layers above (i.e. applications) transparently[15]
- Another design point was the ability to interconnect different physical networks, including the ability to route data from one network to another, again transparently to the upper layer functions

[15] The word 'transparently' here means that the application does not need to be concerned with the underlying physical network implementation.

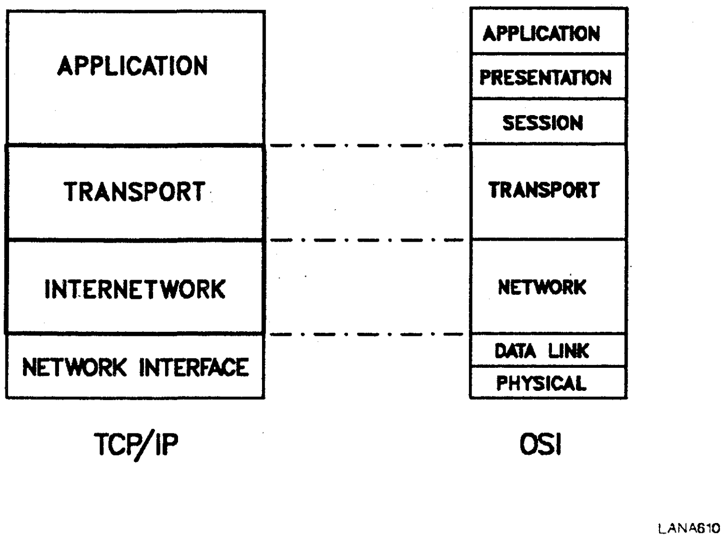

TCP/IP vs. OSI Layering

Although TCP/IP can not be compared directly to the OSI network model, the functions of the TCP/IP layers are similar to those found in some of the OSI layers. The attempt here is to make this comparison. Please remember that TCP/IP and the OSI model are two very different ways of implementing the same function (application to application connectivity across one or more networks).

The ISO OSI model (and SNA) sees the networking world as consisting of 7 layers while the TCP/IP view is of 4 layers.

- Network Interface layer
  - Roughly equivalent to OSI layers 1 and 2

    Note: In some implementations (i.e. X.25) some layer 3 functions will be included in the network interface layer).
  - This layer includes the physical interface and the logical link connection
  - Logical link connections supported include 802.2 and X.25 along with a number of other link level protocols and implementations
- Internetwork layer
  - Roughly equivalent to OSI layer 3
  - Implemented by protocols such as the Internet Protocol (IP)
- Transport layer
  - Roughly equivalent to OSI layer 4
  - Implemented by protocols such as the Transmission Control Protocol (TCP)
- Application layer

- Roughly equivalent to OSI layers 5, 6 and 7
- Implemented by a set of 'applications' which sit on top of the TCP interface (e.g., XWINDOWS, FTP, SMTP, etc.)[16]
- User applications can be implemented on top of the TCP applications

A good way to summarize the differences would be to say that in the OSI implementation, it is expected that the functionality is clearly assigned to specific layers and that all of the layers will be implemented in a network node. TCP/IP provides the users much more flexibility in designing their implementation to meet specific needs. All of the layers do not have to be in any given implementation, and multiple protocols may be found within one layer (e.g., network layer functions might include IP and X.25 in the same node).

[16] The word 'applications' here means a set of defined protocols designed to provide specific sets of services such as file transfer and messaging functions. Another way of looking at these might be to call them utilities.

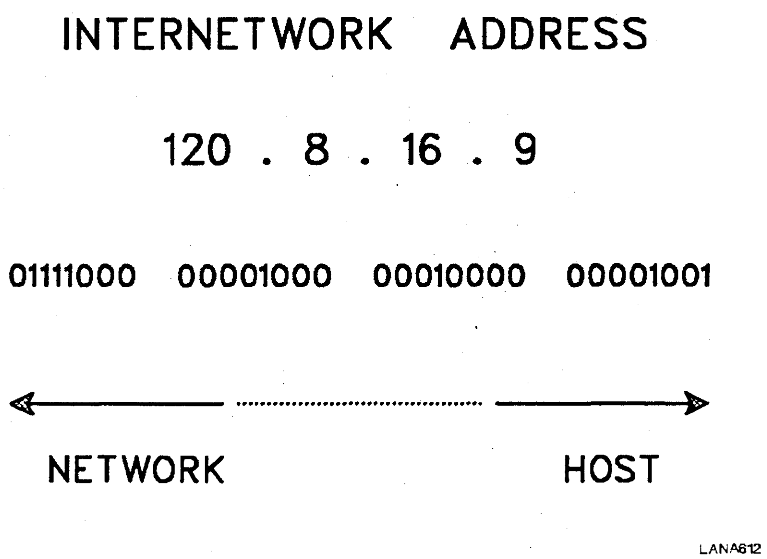

INTERNETWORK ADDRESS

120 . 8 . 16 . 9

01111000 00001000 00010000 00001001

NETWORK HOST

LANA612

Internetworking Addressing

Because of the ability to interconnect network nodes across multiple networks, an appropriate addressing scheme had to be designed along with the networking protocols. All users of the network (including applications) are located at *hosts*. The word 'hosts' here does not mean the same as what we traditionally have called hosts (i.e. large system controlling a portion of the network). Host here means any network node, whether it be a large system, a small PC or anything in between.

- Each host is assigned a network (IP) address, 32 bits long, called 'internet' address in TCP/IP terms
- The address is usually represented in dotted-decimal form
  - 120.8.16.9 (01111000 00001000 00010000 00001001)
- The IP internet address consists of two major parts
  - < network address >
  - < host address >

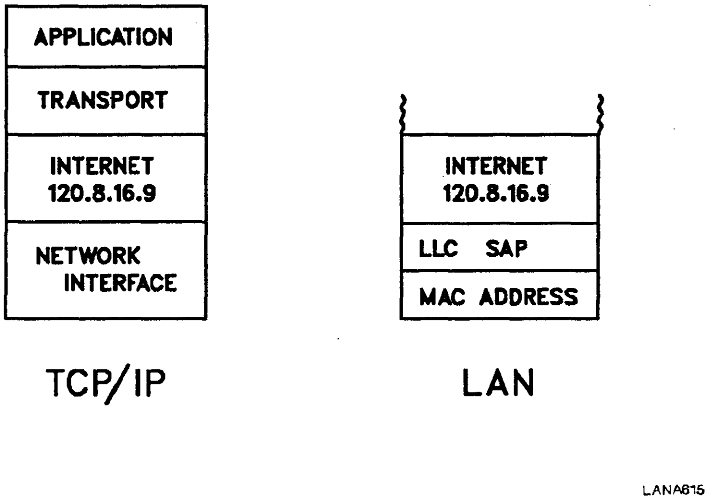

- The <network address> portion of the address is assigned by a central authority (the NIC), and must be unique throughout the interconnected network
- The <network address> will be one of 4 classes which determine the division between the <network address> and the <host address>[17]
  - Class A allows 7 bits to be used for the <network address> portion (128 networks); remaining 24 bits used for hosts
  - Class B allows 14 bits to be used for the <network address> portion; remaining 16 bits used for hosts
  - Class C allows 21 bits to be used for the <network address> portion; remaining 8 bits used for hosts
  - Class D is used for 'multicast' addresses (limited broadcast addresses)

  Note: IMPORTANT. This address is the NETWORK or IP layer address; in order to communicate to another host, the address still must be resolved into the appropriate lower level address(es) as required by the physical network being used to transport the packet. For example, a LAN network node will require a MAC address. Details of how this is done will be discussed later in this Unit.
- The <host address> must be unique within a given network
- Network nodes moved from one network to another must have their <network address> changed, and will probably also require a change in the <host address>

[17] Class of address is determined by the first several bits in the address field (refer to Chapter 2 of the TCP/IP reference manual).

Subnets and Subnetwork Addressing

- Large networks can be broken into smaller networks called subnets
- Subnets are implemented by maintaining the integrity of the <network address> but dividing the <host addresses> portion into two separate pieces: <subnetwork address> <host address>
- Specific implementation of the address split is left up to the user
- Implementation of subnets requires that all hosts have the IP code supporting subnetworking

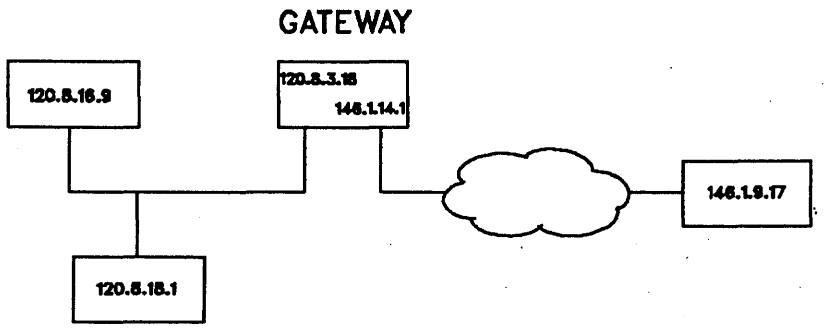

LANA618

Interconnecting TCP/IP Networks - Overview

Interconnect Gateways

Details of LAN interconnectivity will be covered in a later Unit; only IP routing will be discussed here. The two functions have different purposes in life (network vs. LAN segment interconnection). What we would like to emphasize here is the fact that TCP/IP can be found in many different types of networks and, in fact, was designed to support the interconnection of these networks using what is known as the *gateway* function.

- TCP/IP 'gateways' can connect networks with different physical implementations (e.g., LAN and X.25), as long as the same higher level network implementation (i.e. IP) is used[18]
- This function is also referred to as a 'router' function
- Packets are routed through the network based on their target <network address>; the <host address> is ignored during inter-network routing
- Routing is one of the more important functions of the IP layer; therefore, basic gateway function is part of the base IP code; it is available in any node running TCP/IP

[18] This differs from how IBM envisions gateways (connecting two independent or dissimilar networks together, or LANs to hosts).

An Overview of IP Routing

All network nodes using IP have a basic routing (gateway) function which allows packets to go from one network to another. In order to implement a 'full-function' gateway, additional protocols are needed. Let's take a brief look at the basic and full-function gateways.

Basic Gateway Functions

- Basic gateway function provides the ability to connect two networks together
- Gateway has limited knowledge of the networks; it only knows about the hosts on the networks to which it is connected
- 'Direct routing' is used if the destination and source hosts are on the same physical network (same <network address>)
- 'Indirect routing' is used if the destination and source hosts are on different networks (routing must be done via one or more gateways)
  - Source only needs to know about the address of the first gateway
  - Each gateway will forward the packet until the destination network is reached; the packet is then sent using direct routing
- Each host must keep an IP 'routing table' which maps:
  - Local host addresses as direct routes
  - Off-network destination host addresses as indirect routes through a locally attached gateway
  - A default route for packets with a destination network address not listed in the routing table

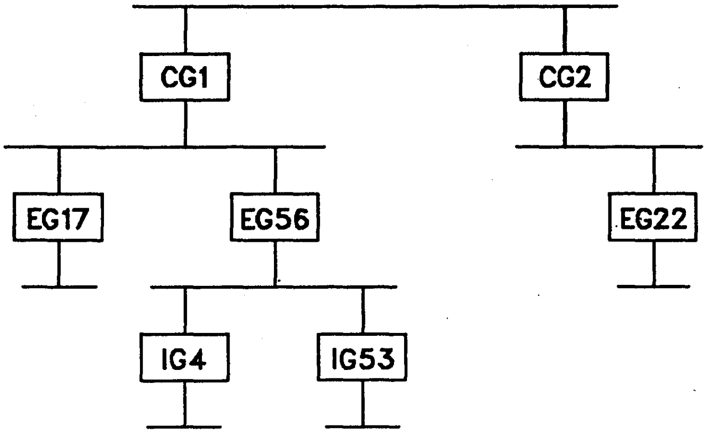

LANA620

Full-Function Gateways

More sophisticated gateway functions are needed in situations where the gateway needs to know about all interconnected networks and where routing tables need to be built dynamically due to continuous network topological changes. The gateways are divided into two classes: core and non-core gateways.

- Core gateways are maintained by the NIC and are used to interconnect major user networks to ARPANET
- Non-core gateways are used to interconnect multiple user networks together
- Both gateway types must know about all of the networks that they support; this is accomplished by transferring information using one of the gateway protocols

Gateway Protocols

There are several protocols associated with the transfer of information between gateways (routers). The specific protocol(s) being used will be determined by the gateway type.

- Gateway-to-Gateway Protocol (GGP)
  - Used to transfer information between 'neighbor' gateways
  - Each gateway will update its tables based on the received information and will send information about the changes to its neighbor
  - Information transferred includes networks reachable by this gateway and the path cost for reaching the network
- Exterior Gateway Protocol (EGP)

  - Used to transfer information between 'exterior' gateways (exterior gateways are ones that connect networks in different autonomous networks)[19]
  - The EGP protocol supports passing of route information as well as other dialogs between the exterior gateways
  - Exterior gateways collect information from interior gateways (gateways within an autonomous network)

- Interior Gateway Protocols (IGP)

  - A set of protocols used to transfer information between 'interior' gateways
  - Examples of IGP protocols include the Routing Information Protocol (RIP) and the HELLO protocol
  - RIP is a class of protocols based on Xerox network XNS routing protocols, and is often found in UNIX environments

[19] An autonomous network is a group of physical network connected together and administered by a single authority.

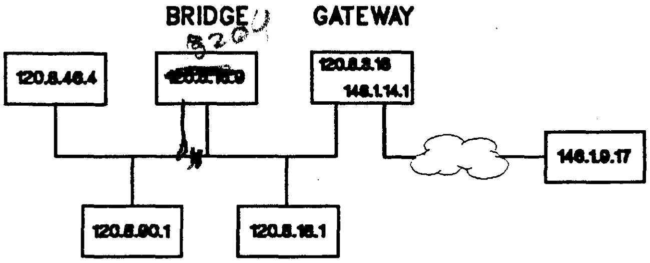

LANA822

TCP/IP and IBM Bridges

Details of bridging techniques will be addressed in a later unit.

- In order for TCP/IP to be used in a multi-segment IBM Token-Ring LAN, the code must support source routing
- Source routing is separate from the TCP/IP code and will not be found in many OEM TCP/IP packages
- As pointed out above, bridging and IP routing are DIFFERENT

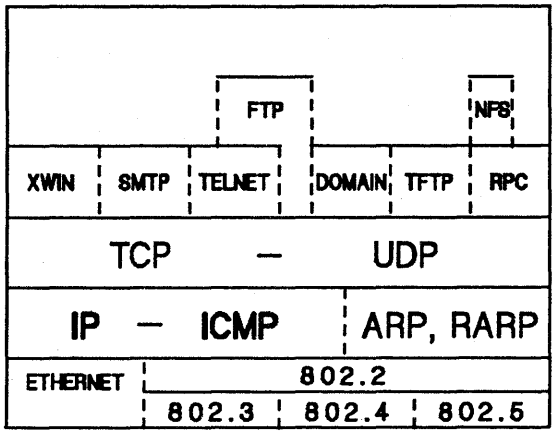

LANA630

Session 2 - The TCP/IP Protocol Suite

IP - Internet Protocol

- Roughly equivalent to the OSI network layer in functionality
- Provides transparent networking services to upper layer software functions
- The data travels across the network in 'packets'; the basic packet used to move information is called the 'IP datagram'
- IP is a connectionless-oriented protocol providing delivery of data packets on a best-effort basis
- Since it is a connectionless protocol, it does not provide delivery verification, guaranteed sequential delivery of packets or error recovery
  - All of these functions will be the responsibility of some higher level software function such as TCP or an application
- Because of physical network limitations, IP will sometimes have to fragment a packet
  - IP expects all physical sub-networks to be able to support a packet size of at least 576 bytes
  - Data packets are assigned a unique identification number prior to fragmentation
  - When the packets are fragmented, each is sent with an offset value
  - Receiver can reassemble the packet using the identification number and the offset values (packets must be received sequentially; the receiving node will dump out of sequence packets)

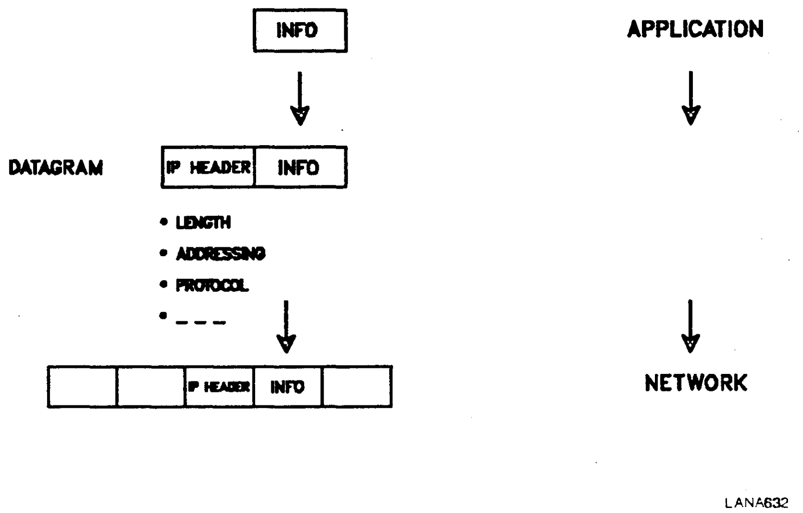

IP Datagram Format

Ref.:TCP/IP Overview manual - Chapter 2

The information from the layer above IP (e.g., TCP) will be passed to the IP layer in a format consistent with the protocol being used. IP will envelope this information in an 'IP datagram' without any alterations to the information.

The IP datagram will consist of the IP header followed by the unaltered information field. Some of the fields contained in the IP header include:

- Version of IP being used
- Lengths of the IP header and total packet
- Identification number (and offset values if fragmentation is being used)
- Higher level protocol (e.g., TCP, ISO-TP4, etc.) to which IP should deliver the packet
- Source and destination IP address

When the IP datagram fields are completed, the packet is passed to the DLC layer where it is further enveloped

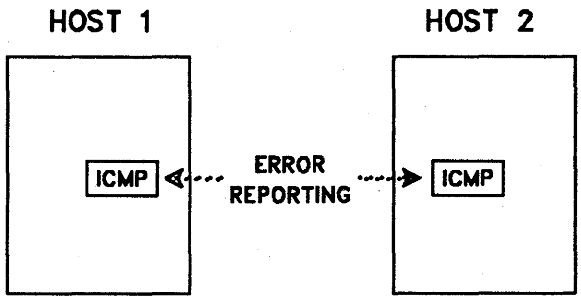

LANA634

ICMP - Internet Control Message Protocol

- ICMP is used for basic error reporting between two hosts at the IP level
- ICMP is an integral part of the IP code (always there); however, use of the code is not required
- The intent of ICMP is to support some error reporting, not to address every error situation (in other words, ICMP does not relieve the upper layer protocols from the responsibility of ensuring error free delivery)
- Some of the types of ICMP messages include:
  - Destination unreachable (e.g., gateway, host, port within the host)
  - Time stamping
  - Overrun conditions
  - Information (address) request/reply

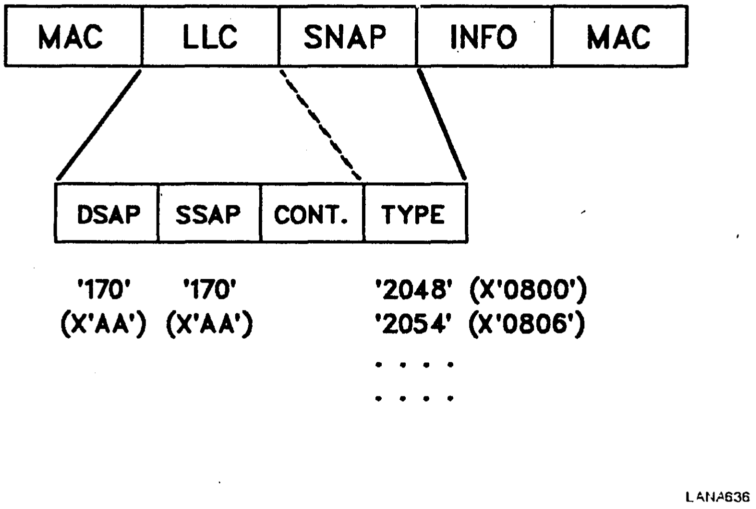

SNAP - Sub-Network Address Protocol

Earlier we said that that the IP address had to be converted to a physical address before it could be sent across the network. The protocol that performs this function is called the Address Resolution Protocol (ARP). Before we discuss ARP, we need to understand still another protocol, called Sub-Network Address Protocol (SNAP).

In the unit on Media Access Control (MAC) frames, we highlighted some of the differences between the 802.2 implementations and Ethernet (or DIX) Version 2 implementations, but pointed out that the two implementations are compatible (can co-exist on the same LAN segment with no problem).

At the level above the MAC level, the 802.2 implementations use Logical Link Control (LLC) protocols. Ethernet, on the other hand, uses a proprietary implementation, incompatible with 802.2 formats; it is not even considered to be a true LLC function. In order for TCP/IP to be used in a LAN implementation, it must support both types (i.e. 802.2 and Ethernet) of implementations.

In addition to the 802.2 vs. Ethernet consideration, there is another one. As more applications use the 802.2 protocols, the need for SAP addresses continues to increase. Since SAP addresses represent a finite resource, something needed to be done to address this situation.

- SNAP is an extension to the 802.2 protocol to accommodate the additional addressing requirements
- SNAP uses the standard 802.2 LLC header (DSAP/SSAP/C Field) and appends a 2 byte field to the header, called the 'Type' field

- SNAP has gone through several iterations; the currently recommended version is designed to work on all three IEEE LANs (i.e. 802.3, 802.4 and 802.5)
- The DSAP and SSAP addresses used to indicate the presence of a SNAP Type field will always be '170' (x'AA')
- SNAP address assignments (related to IP functions):
  - 2048 (X'0800) for IP datagrams
  - 2054 (X'0806) for ARP datagrams
  - 32821 (X'8035) for RARP datagrams

  Note: Although the current level of SNAP supports all of the IEEE LANs, please keep in mind that some earlier levels will only support specific implementations (e.g., Ethernet) and may present compatibility problems on the LAN.

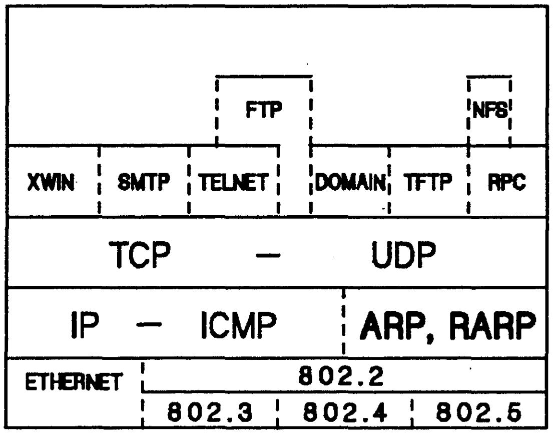

LANA640

ARP - Address Resolution Protocol

- A host (source) wanting to send a packet to another host (destination) will turn the packet over to IP; the IP routing function will attempt to match the destination IP address to a physical address
- Some networks will require a more sophisticated matching scheme than others (bear in mind that the actual network interface will be a function of the device driver being used to support the physical network); this is where ARP enters the picture[20]
- In the case of LAN based networks (i.e. Ethernet and 802.X LANs), ARP is used to perform this translation
- Nodes supporting the ARP function will maintain a table that equates IP addresses to physical addresses

[20] The discussion here will be limited to LANs.

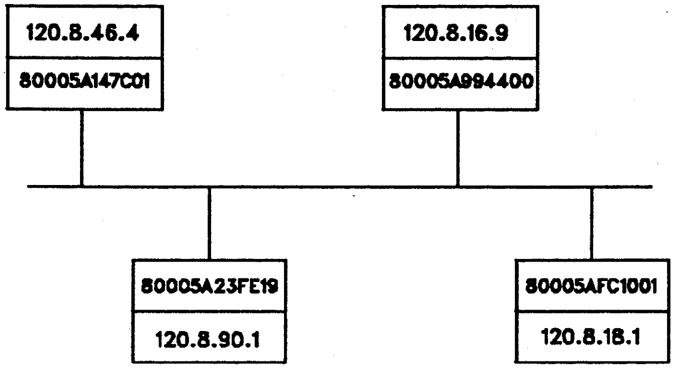

LANA642

- Packets containing an IP address that is not in the table (sometimes called the ARP 'cache') will result in an ARP broadcast being sent across the internet in an attempt to locate the address
- A host recognizing it's IP address in the ARP packet will respond allowing the originator of the broadcast to identify the physical address of the destination (similar to the way NETBIOS works)
- An ARP packet contains information such as:
  - Type of hardware and protocol being used (e.g., Ethernet)
  - Source and target hardware (physical) address (6 bytes - 48 bits)
  - Source and target IP address (4 bytes - 32 bits)
- ARP address resolution is not effected by subnet implementations

RARP - Reverse ARP

- Some hosts (e.g., diskless workstations) do not know their IP address but do know their hardware address
- Using RARP, the host can send a request to a 'RARP Server' to learn the IP address of the host (opposite to ARP where the IP address is the known)
- Requires a specific function called an RARP Server in the network; the Server will have a pre-defined mapping table to support this function
- The contents of the RARP packet are very similar to the ARP packet

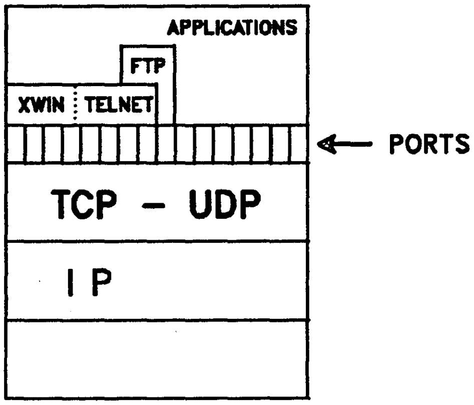

LANA644

Next Stop - Layer 4

The protocols we have just discussed are primarily network layer functions (as they relate to the OSI network model). We are now going to move up a layer to what OSI calls the transport layer. In the TCP/IP world, this consists primarily of 2 protocols: UDP and TCP. We will first look at UDP and then TCP, but before we do that, there's another concept we need to discuss.

Ports and Sockets

- Each host has one or more 'processes' which have a need to communicate to processes on other hosts
- A good example of a process is an application
- In order to use the functions of TCP/IP, each process must identify itself to TCP by one or more ports
- Some commonly used processes (e.g., Telnet and FTP which we discuss later) will always use the same port number; these are called 'well-known' ports
- The range of well-known port IDs is 0-255; user applications will uses IDs above this range
- Ports are used by UDP and TCP, and also by the ISO Transport Protocol - Class 4 (ISO-TP4)
- Processes are known to the network by a concatenation of their port address and the IP address: the total identifier is called a 'socket' or 'socket address'
- Socket = <network address> <host address> <port number>

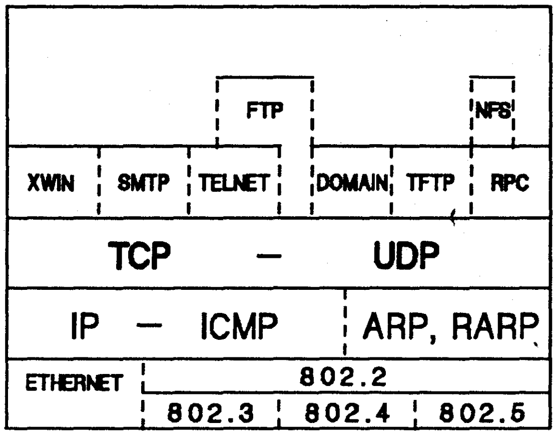

LANA646

UDP - User Datagram Protocol

- UDP provides the interface between ports and IP for connectionless data transfer (i.e. datagrams)
- UDP provides no error recovery, flow control or reliability
- UDP will receive the information from the port (process), and envelope it in a UDP datagram
- UDP datagram will contain the information, along with the source and target port addresses, a length value and a header checksum value
- TCP/IP applications that use the UDP protocol will be identified as we discuss each application

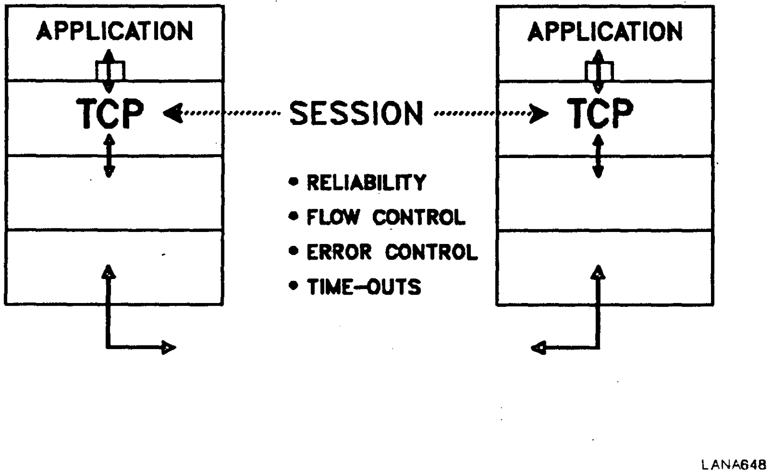

TCP - Transmission Control Protocol

- TCP is a transport layer protocol that provides connection type services, including reliability, flow control and error recovery
- Connection represents a logical circuit between two sockets
  - Logical connection must be established prior to data flow (very similar to the concept of an SNA or NETBIOS session)
- TCP provides a continuous data stream interface to the application; it will take care of segmentation of the data as required by the lower level network functions
- Reliability:
  - TCP assigns a sequence number to each byte; sequence numbers are used at the receiving end to resegment the data and confirm delivery[21]
  - All packets must be acknowledged; if no acknowledgement is received in a given time, TCP will resend the data
- Flow control:
  - Flow control is accomplished using a windowing scheme
  - TCP uses byte oriented sequencing (as opposed to counting packets as is done in other protocols)
  - Window size (number of bytes to be transmitted) is set when the connection is established, but will change dynamically during data transfer

[21] Only the sequence number of the first byte of each packet is sent to the receiver.

  - Windowing scheme allows the receiver to notify the sender as to the maximum number of bytes that the receiver can handle beyond the last successfully received segment

- Error recovery:

  - Error recovery is done using byte sequencing; the receiver sends an acknowledgement indicating the last byte number successfully received; sender will then retransmit any bytes (in the form of packets) not yet acknowledged
  - As in other protocols, such as SDLC, once an errored-packet is received, the receiver will stop counting, even if the following packets are received error free

- Time-outs:

  - TCP has the ability to adjust time-out values related to data transmission by examining the time it takes for an acknowledgement to be returned to packets sent (round-trip delay)
  - How TCP reacts to time-out and error conditions will be determined by how each host (user) implements error recovery

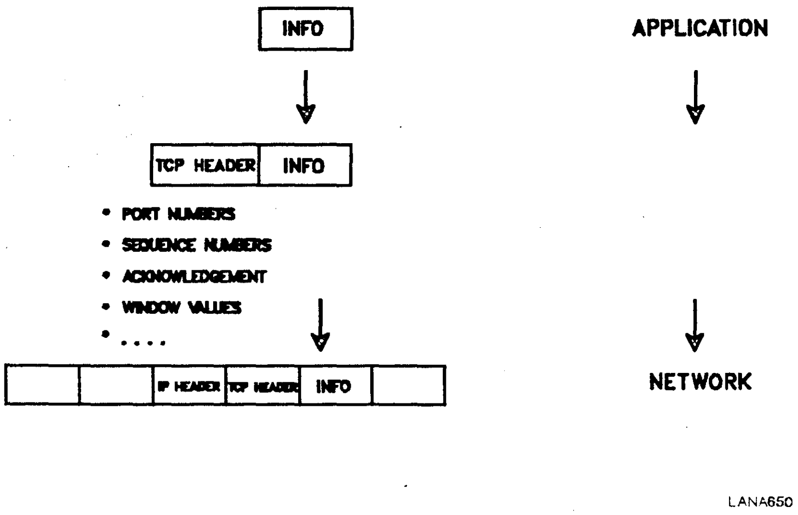

TCP Segment Format

Like just about every protocol we have seen so far, TCP will take the data from the higher level application and envelope it; the TCP envelope is called a 'TCP Segment'. Obviously, the contents will be much more elaborate than that used by UDP. The TCP Segment will contain information such as:

- Sort and destination port numbers
- Byte sequence numbers
- Acknowledgement indicators
- Window values
- Checksum for the TCP header
- Options
- User data

Session 3 - TCP/IP Applications

Objectives

Upon completion of this session, you should be able to:

- Briefly describe the following commonly used TCP/IP 'applications':
  - TELNET
  - File Transfer Protocol
  - Name Servers
  - Simple Mail Transfer Protocol
  - X-Windows
  - Remote Procedure Call
  - Network File System
- List the IBM products that implement TCP/IP protocols and identify which protocols are supported by each product.

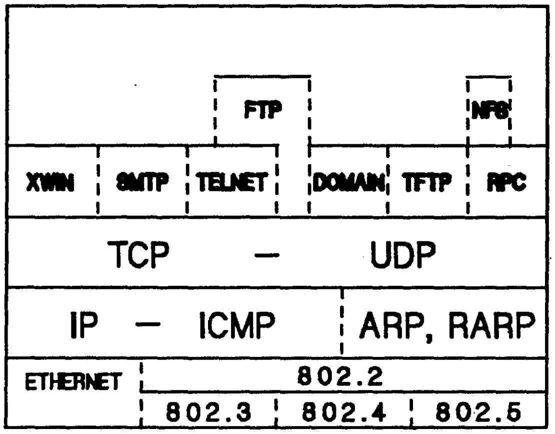

LANA660

Introduction

In addition to the network and transport layer protocols, the TCP/IP protocol suite includes a number of higher level software packages, called 'applications'. Functionally, they provide services such as file transfer, messaging and virtual terminal support. In a sense, they are more analogous to utility programs that real applications as we usually think of them.

These applications interface to TCP (or UDP) and are all optional. All TCP/IP packages will support the network and transport protocols (i.e. TCP/IP and their related protocols); however, inclusion of specific applications will vary widely from vendor to vendor.

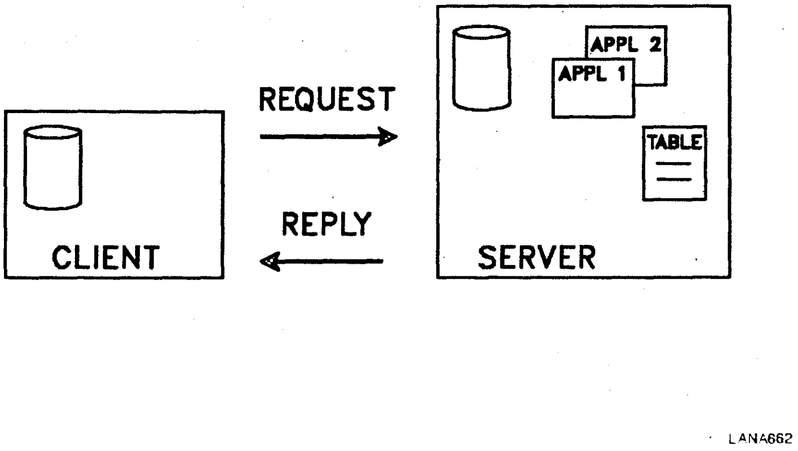

Clients and Servers

Ref.: TCP/IP Overview manual, Chapter 3

Like more and more of today's network implementations, TCP/IP uses a 'client/server' model for communications between applications. Before we look at the TCP/IP applications, we need to briefly review the meaning of the terms client and server.

- A 'server' is a process or application that provides a service to other users (applications) in the network
- A 'client' uses the services of a server by invoking a 'requestor' function
- Requests are sent by the client; replies are sent by the server
- Some servers (i.e. the more popular TCP/IP applications) provide their services through 'well known' ports
- Server functions require considerably more code than do client functions; TCP/IP products providing support for the various applications may provide only the client function, only the server function, or both

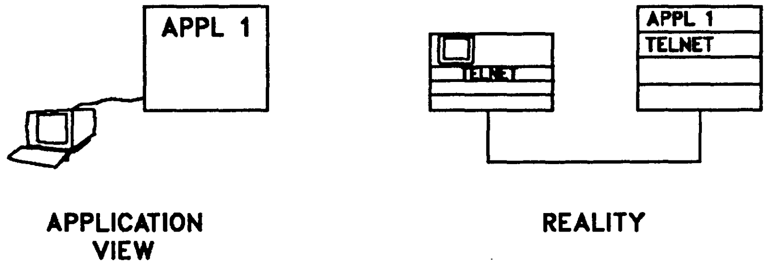

LANA664

TELNET

- Provides a standardized interface to allow a process at one host (using the TELNET client function) to access a process at another host (using the TELNET server function) as if the client was a local terminal directly connected to the host
  - A PC running DOS can log on to a VM host (using TELNET services); to the VM host, the PC looks like a normal CMS user's terminal
- TELNET is based on the idea of a Network Virtual Terminal (NVT)
  - NVT is a virtual device that provides a basic terminal structure for a display and a printer
  - Each host then maps it's unique terminal structure requirements to the NVT
- Support for options over and above the basic structure can be negotiated at connection time
- Examples of some of these options:
  - Support for specific terminal types (e.g., 3278-2 type)
  - Window sizes
  - Translate table selection
  - Full-screen formatting
  - Support for echoing

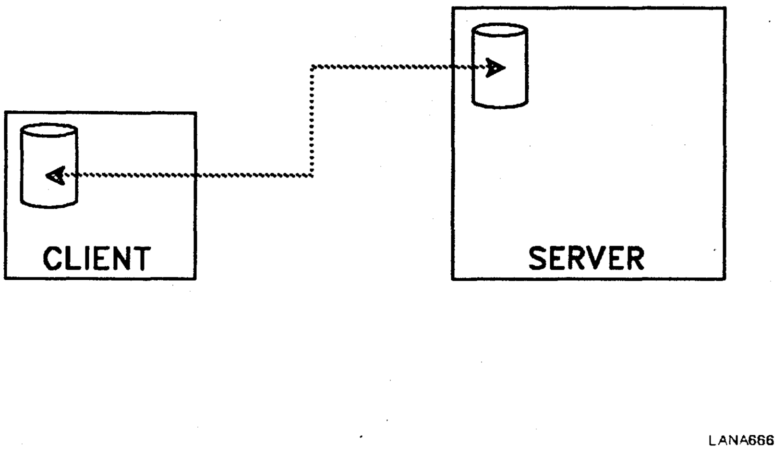

TFTP - Trivial File Transfer Protocol

- Limited function protocol that provides support for disk-to-disk file transfer
  - Limited function means that the protocol provides the ability to read or write a file from or to a server; it is not a full function file transfer program
- It is implemented using UDP (does not provide for reliable transport services)
- No user authentication provided

FTP - File Transfer Protocol

- Full function protocol that provides support for host to host file transfer in either direction
- User requesting the connection becomes the client; user providing the service becomes the server
- FTP is implemented using TCP which provides the end to end connection-oriented services
- User authentication done prior to the data transfer
- Two connections are required between the hosts to support the operation:
  - One connection is used for the logon and to control the operation; it uses the TELNET protocol
  - A second connection (port) is used for the actual data transfer; it uses a Data Transfer Process (DTP) which uses the TELNET conventions
- Overview of sequence of operations for file transfer:
  1. Logon to the remote host (client to server)
  2. Identify the directory to be used
  3. Identify the file to be transferred
  4. Define the mode of transfer (e.g., ASCII, EBCDIC, etc.)
  5. Start the transfer
  6. End the operation

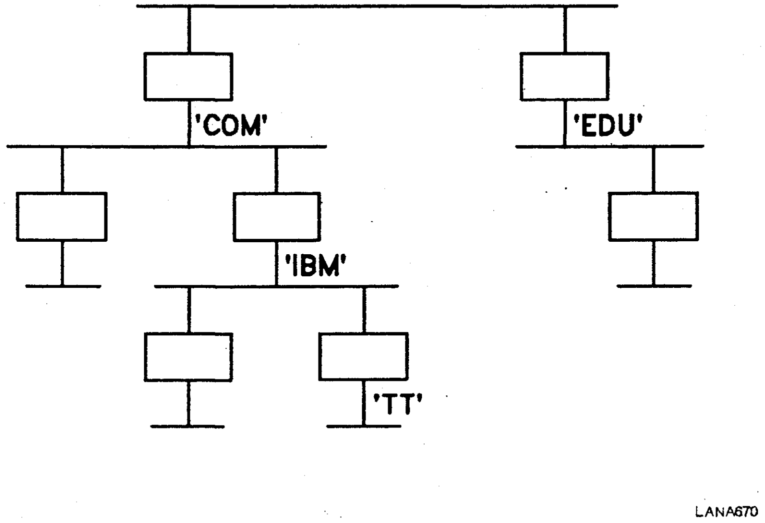

Name Servers

As the acceptance and use of TCP/IP grew, the use of numeric IP addresses was replaced by the use of symbolic names. Names, like the addresses, were concatenated (e.g., nm1.nm2.nm3). Before a packet could be sent through the network, the name had to be mapped to an IP (numeric) address; this function can be done using a table at the source host, or can use the services of a 'name server'.

Today, network names are structured in a hierarchical fashion using a concept of domains. What this means is that the internet is divided into hierarchical domains so that the responsibility for mapping between a network name and an IP address can be done efficiently.

Domain Name System

- Highest level domain will be a major network type
  - *mil* will be used by military groups
  - *edu* will be used by educational institutions
  - *com* used by commercial organizations
  - There are also others; these are given as an example
- Next level will represent major networks within a network type
  - For example, within the 'edu' network would be a domain for each university
- At the next level, the network domains would be divided into individual networks, such as networks within a given university
- The network name would be a concatenation of the various domain names

  - For example, the name ′vm9.tt.ibm.com′ represents a host ′vm9′ on the network ′tt′ which is a domain on the ′ibm′ network which is connected to the top level network ′com′

Name Servers

Before a packet can be sent across the internet, the target name identified by the user sending the packet must be mapped to a valid IP address.

- Client function is called the ′name resolver′ since it has the responsibility for converting the name to an IP address
- If the name resolver function can not locate the name in it′s resident mapping table, it will send a request to a name server
  - The IP address of the local name server must be known to the client
  - The client can request recursive or nonrecursive services to be used in the event that the name server does not know the target name
    - Recursive service means that the local name server will contact other name servers and report the results back to the client (IP address or a negative response)
    - Nonrecursive service means that the local name server will send the IP addresses of the other name servers back to the client; client must then contact them directly
- Another protocol, called the ′Domain Protocol′, is used for communicating between the domains
- Domain Protocol Messages can be sent using either TCP or UDP

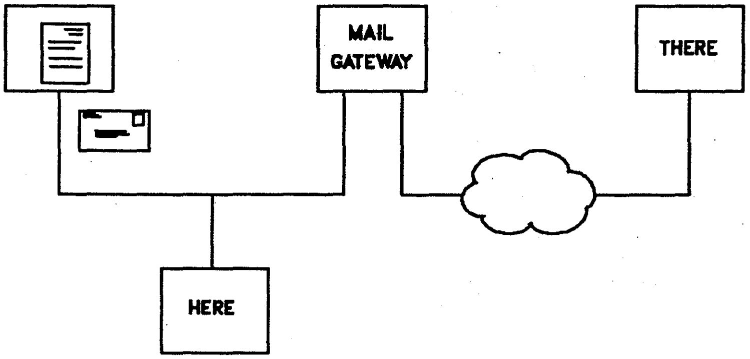

LANA672

SMTP - Simple Mail Transfer Protocol

- Protocol that provides support for messages and note exchanges between two hosts
- SMTP works by spooling the mail, establishing an end-to-end connection using TCP, and keeping a copy of the mail until the receiving host's SMTP function has acknowledged receipt (not a store and forward mail system)
- SMTP mail can be sent across other (local) mailing systems through an appropriately configured gateway; however, SMTP will only guarantee mail integrity from the sender to the gateway
- SMTP will use the name server functions discussed above for address mapping
  - Name format will be <name@host.net1.net2> where 'net1' and 'net2' are the hierarchical network names
- SMTP consists of two standards:
  - Standard for the mail format (layout and contents of the mail envelope)
  - Standard for the host-to-host transfer of the mail
- SMTP is one of the most popular TCP/IP applications and will be supported by most vendor TCP/IP offerings

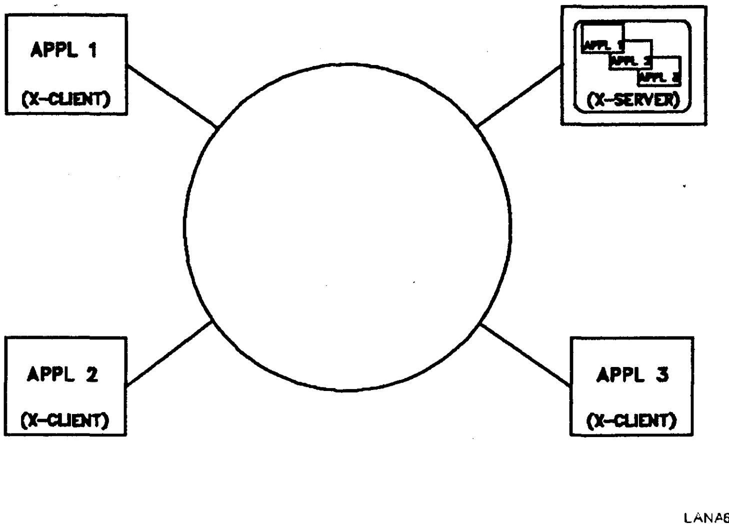

X-Windows

- X-Windows is a windowing system that allows a user to simultaneously display screens from several local and/or remote processes
  - Local windows are managed through sockets
  - Remote windows can use the services of TCP/IP, or any reliable byte stream protocol
- Processes (programs or applications) can be viewed, or run in a virtual mode with only an indicator (icon) displayed on the screen
- Communication is done using a client/server function
  - Application, called the X-Client, receives and sends information to a programming interface called 'xlib' instead of communicating directly with the display
  - X-Client is responsible for window contents, not the actual display on the terminal
  - User's terminal is called the X-Server
  - X-Server keeps track of the windows using a screen hierarchy similar to a DOS disk root and sub-directory structure
  - X-Server works with the display management software at the user terminal to display the actual window(s) on the screen
- Since the application talks to a programming interface ('xlib'), it does not need to be concerned with the type or location of the user display
- XLIB provides an interface supporting more than 100 different application calls
  - Specifics of the programming interface can be found in the appropriate RFCs and product documentation supporting the interface

RPC - Remote Procedure Call

- RPC is an application programming interface (API) that allows a program (client) to call a subroutine executed on a remote host (server)
- Server receives the 'call' message, interprets the request, executes the request and returns the results to the client
- Process at server is physically independent of client
- RPC can use any transport protocol (includes both TCP/IP and UDP/IP)
- RPC supports distributed processing; it allows a process to distribute the work among multiple (remote) hosts - it is not cooperative processing which requires true interaction between the multiple processes
  - An analogy would be sending some of your clothes to the laundry and picking them up a day later as opposed to doing them all yourself

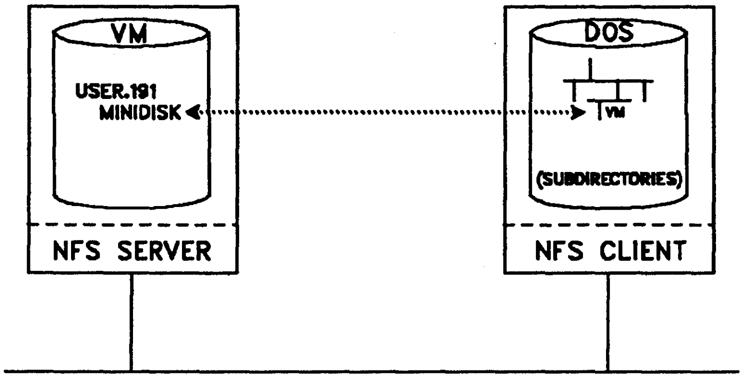

LANA678

NFS - Network File System

- NFS, developed by SUN Microsystems, allows hosts machines to share file systems across a network
- NFS was developed to be machine, operating system and transport protocol independent; this was accomplished by using the RPC interface
- NFS allows remote file systems to have the appearance of being local resources to users (protocol is transparent to the user)
- NFS implements two protocols: Mount protocol and NFS protocol
  - Mount protocol is used to establish the connection and access the appropriate set of files (directory)
  - Options supported include password security, code conversion, remote disk addresses, etc.
  - NFS protocol is used for the actual disk functions
  - Supported functions include searches, reading and writing, and directory maintenance

- NETBIOS Services Protocol
- Ping
- Netstat
- Finger
- Time & Date
- Remote Print Servers

LANA680

Other TCP/IP Application Protocols

The only protocols in the TCP/IP suite that are required in TCP/IP implementations are IP and ICMP; all of the others are either recommended or optional.

Recommended protocols include:

- TCP
- TELNET
- FTP
- SMTP
- Domain Name Protocol
- NETBIOS Services Protocol (see below)

In addition to the ones we have discussed in the last two sessions, there are a number of other protocols the will be found in some TCP/IP implementations. We will look at them here very briefly.

- NETBIOS Services Protocol
  - Designed to allow applications written to a NETBIOS interface to use the network services of TCP/IP
  - NETBIOS name services are supported through the use of a NETBIOS Name Server function which works in a manner similar to a TCP/IP Name Server (name to address mapping)
  - NETBIOS session services are mapped to TCP connections
  - NETBIOS datagram services require the use of a NETBIOS Datagram Distribution (NBDD) node
    - IP was not designed to support broadcast messages

- The NBDD node will convert the NETBIOS broadcast messages to UDP packets and send one to each designated receiver of the broadcast message

- Ping
  - Ping is merely an echoing application that uses UDP and ICMP
  - It usefulness is in verifying basic connectivity to another host
- Netstat
  - Netstat is used to query the status of the local host (sockets, links, devices, etc.)
- Finger
  - Used to display information about users of a remote or local host
- Time and Day protocols
  - Used to pass time and date information through the network
- Remote Print Servers (several versions)

New TCP/IP Protocols

As we said before, TCP/IP is a very dynamic set of protocols. Currently, there is emerging two new protocols, one which addresses network management and the other intended to provide more efficient routing. More details will be added here as these protocols are developed, but we wanted to at least expose them to you.

- Simple Network Management Protocol (SNMP)
  - Protocol to be used for exchanging system management information between hosts
  - Provides system monitoring, controlling and reporting services
  - Although SNMP is implemented on TCP/IP, it is really transparent to the network protocol stack and can be implemented on different types of networks (e.g., OSI networks)
- Open Shortest Path First (OSPF)
  - OSPF is a sophisticated routing protocol that provides efficient internetwork routing, particularly in large and medium networks
  - OSPF also provides support for least cost routing and multiple link routing
  - Potential replacement for Routing Information Protocol (RIP)

Cross Reference of IBM Product TCP/IP Protocol Support

Refer to the *TCP/IP Tutorial and Technical Overview* manual (GG24-3376).

Feedback Questions

Directions:

First select the correct answer for each of the following questions by circling the appropriate letter. In a few minutes you will be asked to input your answers, one at a time, using the Student Response Unit.

1. Which of the following statements is NOT true about the TCP/IP internet address?

    a. Addresses are usually written in 'dotted decimal' format (e.g., 120.8.16.9)

    b. Address represents the network address and the host address

    [✓ c. Address is the physical network address used by LAN nodes]

    d. Address can be subdivided to support subnetwork addressing

2. (True/False) TCP/IP gateways actually function as network routers and are not functionally equivalent to an SNA gateway implementation

    [✓ a. True]

    b. False

3. Which of the following is a false statement?

    a. SNAP is used to extend the addressing capabilities at the LLC level

    b. ARP and RARP are used to resolve internet addresses to physical addresses

    c. Processes (applications) communicate to the network through ports

    [✓ d. TCP can be used as either a connection-oriented or connectionless protocol]

4. Which of the following are TCP/IP 'applications'?

    a. TELNET, SMTP, FTP, TFTP, UDP

    [✓ b. TELNET, FTP, TFTP, SMTP, X-Windows]

    c. FTP, TFTP, X-Windows, UDP, ICMP

    d. SMTP, TFTP, FTP, RPC, ICMP, ARP

UNIT 07: LAN Interconnection Architectures

Session 1 - Interconnection of Token-Ring LANs

Introduction

Earlier in this course, we defined the terms bridges, routers and gateways. In this Unit, we will be taking a much closer look at how these functions are implemented, particularly how LANs are bridged together. With this knowledge, you should be able to decide which LANs can be connected together, and the implications of making such connections.

Before we can discuss how LANs are bridged together, it is important that you first understand the two different ways of routing information across multiple LAN segments. In the past, multiple networks were interconnected using either routers (which performed the routing at the network level such as is done with TCP/IP) or gateways (which performed the function at some level above the network level). Both routers and gateways required routing tables to support these interconnections. Because of the higher speed and increased performance expectations for LANs, they must be connected together at the lowest possible functional layer for optimum performance. In addition, since LANs can be very dynamic networks, traditional methods using predefined tables do not provide the flexibility required in this environment.

Objectives

Upon completion of this session, you should be able to:

- Explain the difference between source routing and transparent bridging
- Describe IBM's source routing implementation, and the difference between the two types of broadcasts
- Briefly explain IBM's implementation of the spanning tree approach to single path routing
- Explain how a bridge determines whether or not to forward a frame
- Interpret the bits in the IBM Routing Information fields

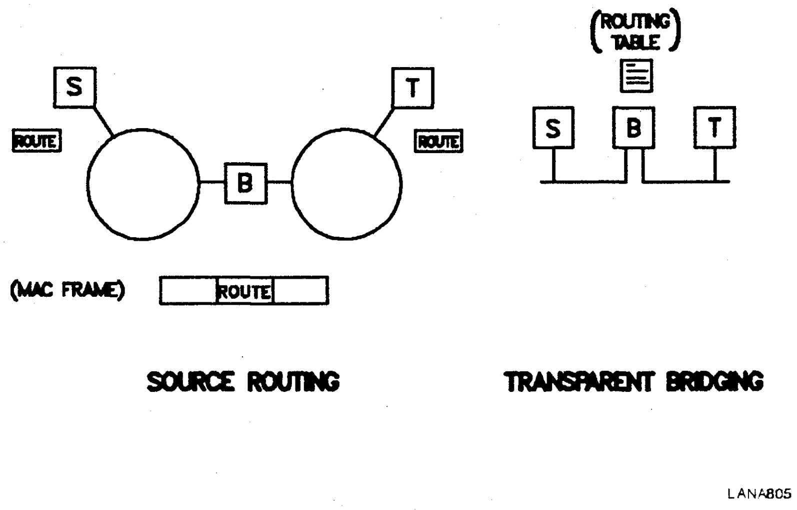

Source Routing vs. Transparent Bridges

Source Routing

- Each frame contains the required routing information
- Bridges do not contain any routing information; they examine incoming frames and determine if the frames should be forwarded
- Bridges are transparent to the source and target stations
- Routes or paths between a target and source station are 'discovered' during session establishment (see below)
- Source will select the best path during the route discovery or session establishment process (i.e. "source routing")[22]
- Source and target stations will both save the image of the path between the two nodes for the duration of the logical connection (i.e. session)
- Any other paths discovered during the discovery process will be discarded

  Note: During the session, loss of a LAN segment or bridge in the path can result in session failure; at a minimum, a new route will have to be discovered,

[22] In today's implementations, the first response back during a route discovery sequence is considered to be the best path; it may not necessarily be the shortest physical path between the the nodes.

Transparent Bridges (Learning Bridges, Spanning Tree Bridges)

- No routing information is contained in the frames; the appearance to the network nodes is that all nodes are on the same LAN segment (thus the name 'transparent' bridging)
- Bridges build routing tables dynamically
- Tables are built during the bridge 'learning' time when it listens to the attached LAN segments and determines the active addresses on each segment
- Note that any bridge will see traffic from non-attached segments as belonging to one of the attached segments
- Bridges also communicate to other bridges in the network
  - Bridges must communicate to one another to ensure that redundant paths are not permitted in the configuration (bridge forwarding can be enabled and disabled dynamically)
  - A special protocol is used for this bridge communication[23]
- The topology for non-redundant paths between all segments is called a "spanning tree"; the bridges build the 'tree' based on the specific capabilities of each bridge
- Once the topology (tree) is built, the bridges will continue to communicate with each other to accommodate any changes in the network
- Routes are transparent to the network nodes (nodes do not store any route information and are not even aware that the logical connection passes through one or more bridges)

[23] This protocol, as implemented on IBM bridges, will be discussed later in this Unit.

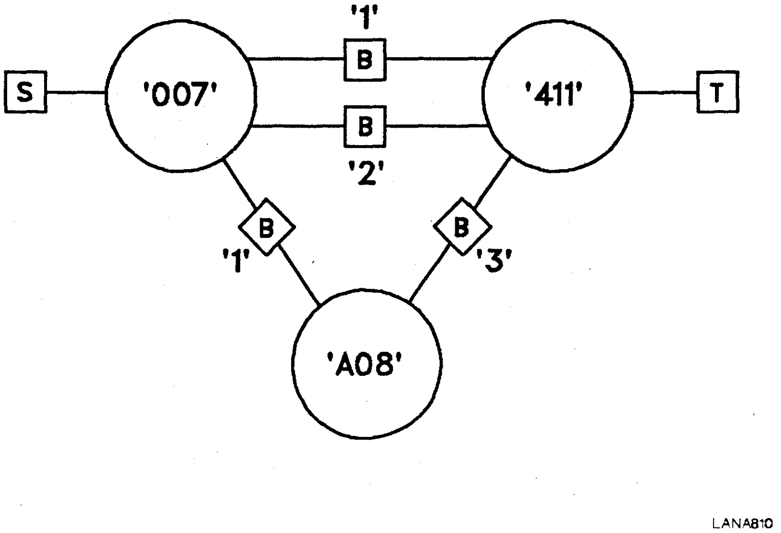

IBM Source Routing Implementation

Ref.:*Token-Ring Network Architecture Reference* manual, Chapter 3.

There are several different ways to implement the actual route discovery process. Before we take a look at how various IBM products implement this process, we first need to look at the types of broadcasts supported in source routing implementations.

Note: The following discussion assumes that all bridges have the frame forwarding function enabled.

All Routes Broadcast

- A source node will send a route discovery frame (e.g., SNA devices will usually send an LLC XID or TEST frame) with a routing information field included and indicating an all-routes broadcast

  Note: Specific MAC frame formats will be covered later in this Unit.[24]

- All routes broadcasts will be propagated across all bridges in the network

  Note: As a broadcast frame is forwarded through a bridge, the bridge will append its path identifier in a routing information field (to be discussed later in this Unit).

[24] Do not confuse an 'all-routes' broadcast with an all stations broadcast address in the DA field of the MAC frame.

- Target will receive one or more route discovery frames (LANs with multiple paths between a source node and target node results in the target receiving multiple copies of the frame sent by the source)
- The target will respond to each copy it receives; the response will include the routing information collected during the search
- The source will select the best path (the first response from the target station)
  - Bridge and/or LAN segment congestion in the shortest physical path between the source and target can result in the source receiving the first response via another path; at the time of the response, this is considered to be the 'best' path
- Next frame sent by the source will include the path selected; target will 'learn' the route when it receives the frame

Single-Route Broadcast

- A source node will send a route discovery frame (e.g., NETBIOS devices will usually send a NETBIOS broadcast) with a routing information field included and indicating a single-route broadcast
- Only bridges with the single-route broadcast (SRB) forwarding option turned 'on' will propagate the search
- In a properly designed and implemented network, there will only be one SRB path between any two nodes in the network
- Target should receive only one route discovery frame
- Target response to the query will include a routing information field indicating all-routes broadcast

  Note: An alternative is for the target to send the response as a non-broadcast frame using the routing field received in the original frame; implementation of how the responses are sent are product specific.
- The source will receive one or more responses depending upon whether or not multiple paths exist between the source and target
- As in the situation above, the first response back wins
- Next frame sent by the source will include the path selected; target will 'learn' the route when it receives the frame

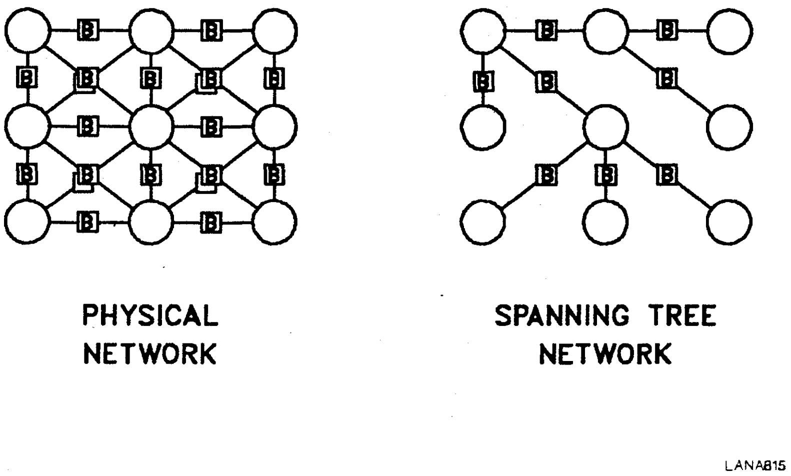

Use of the Spanning Tree Concept with Source Routing

- IBM recommends the implementation of a spanning tree approach to bridge implementation to ensure that all nodes have connectivity through no more than one path
- Spanning tree configuration can be done manually or, with specific IBM bridge products, automatically
- Automatic configuration uses a special protocol between the IBM bridge products called the 'Spanning Tree Protocol'; frames are called Bridge Protocol Data Units or BPDUs
- Each bridge is configured during installation with a set of parameters that uniquely identify it and identify its relative 'cost' in the overall network (e.g., 16Mbps adapters will have a lower cost than 4Mbps adapters)
- In this implementation, a bridge will assume one of three roles:
  - One 'root' bridge is selected and becomes responsible for coordinating the configuration process
  - A 'designated' bridge is one that will have SRB forwarding turned on (bridge has no parallel paths or initially has the lowest path cost in a parallel configuration)
  - A 'stand-by' bridge is one that exists in a parallel path; it will not have SRB forwarding on, but will monitor the BPDU activity in the event that it needs to become a designated bridge because of topological changes (e.g., designated bridge failure)
- The root bridge coordinates the configuration process
  - It does this by sending out special 'Hello BPDUs'

  - These BPDUs are propagated throughout the network through designated bridges, with each bridge updating the information (such as path cost) as appropriate
  - Stand-by bridges receive the BPDUs and update their information, but do not forward the BPDUs

- In automatic configuration mode, the SRB topology will adapt to existing network conditions while maintaining connections using the lowest path cost

  - For example, when an SRB bridge leaves the network, another bridge (or bridges) must be reconfigured to maintain connectivity without redundancy

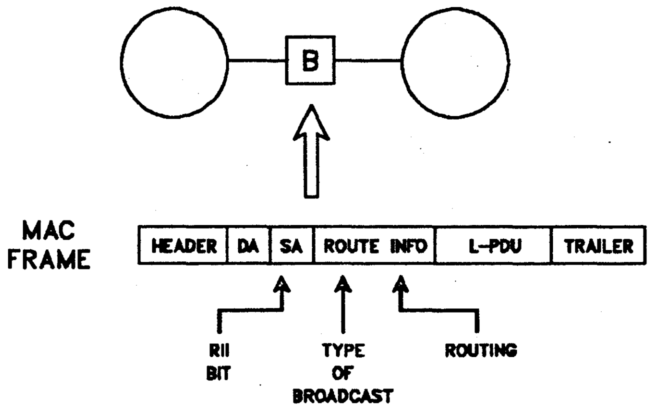

LANA820

When Does the Bridge Copy (and Forward) Data Frames?

The Routing Information (RI) field is contained within the MAC frame, immediately following the 6 byte source address (SA) field. How does the bridge know enough to copy only those frames with an RI field present in the frame?

- Network nodes that implement source routing will include RI information as needed (e.g., route discovery, session flows)
- The presence of the RI field will be indicated by the sending node setting SA field byte 0, bit 0 to a '1' (called the 'RII' bit)[4]
- Bridges will look for this condition (RII = 1) on all frames passing by on attached LAN segments
- If RII = 1, the bridge will inspect the frame to see if it should copy and forward it to the next LAN segment
- The decision to copy and forward the frame will be based on several conditions

  Note: The bridge will not turned on the 'C' (copied) bits at the end of the MAC frame unless it is copying and forwarding the frame.

[4] RII bit stands for 'Routing Information Indicator' bit.

The Bridge Decision to Copy and Forward a Frame

Note: The assumption is made here that the bridge has frame forwarding enabled. Bridges that do not have this function on will not forward any frames.

Note: The word "forward" means that the bridge will copy the frame from one LAN segment into its buffers, and then send it out on its other attached LAN segment (using the appropriate MAC protocol).

1. Is this an all-routes broadcast?
   - If no, go to the next step
   - If yes, examine any existing RI field data (route designators) to determine if the frame has already been on the attached LAN segments
     - If the frame has not been on one of the attached segments, copy it and add an additional route field, then forward it (it obviously has already been on the one it was just copied from)
     - If it has already been on both segments, ignore it
2. Is this a single-routes broadcast?
   - If no, go to the next step
   - If yes, does this bridge have SRB turned on?
     - If SRB forwarding is on, copy the frame, add an additional route field and forward it
     - If SRB forwarding is off, ignore the frame
3. Non-broadcast - examine the RI field to determine the route
   - If this bridge is part of the route, forward the frame
   - If not, ignore the frame

It should be obvious to you after looking at this process that, although bridges do not copy and forward every frame on their attached LAN segments, they do go through a frame examination process which requires processing time and resources in each bridge.

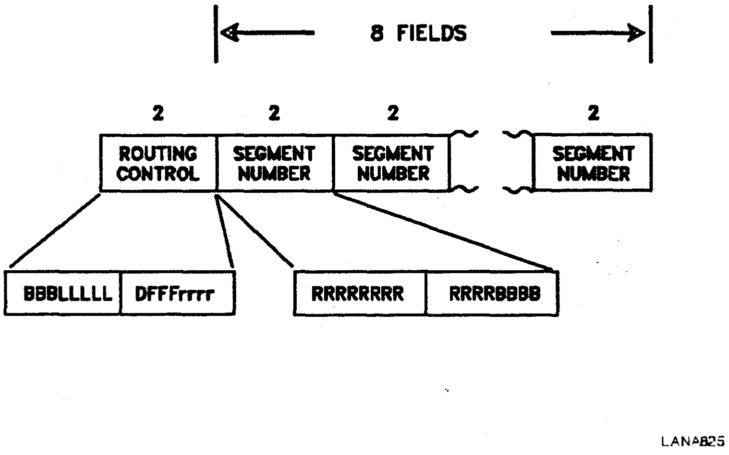

Routing Information Field

Ref.:*Token-Ring Network Architecture Reference* manual, Chapter 2.

As noted above, this RI field will be part of the MAC layer protocol. If present, it will be located between the SA field and the beginning of the L-PDU.

The RI portion will consist of a 2-byte routing control field and up to 8 2-byte route identifiers.

Note: IMPORTANT. During this discussion, please note that the bridge only alters the RI portion of the MAC frame. All other fields, including the destination and source address fields and the entire L-PDU, are left unaltered by the bridges. This is true when interconnecting similar LANs (i.e. 2 token-ring LANs), but will not necessarily be true when interconnecting different LAN types as we shall see in Session 2.

Routing Control Field

- Broadcast indicators
  - 3 bits that identify whether this is a broadcast frame, and if so, the type of broadcast
  - All-routes broadcasts will cross all segments (potential for multiple copies on any given segment)
  - Single-routes broadcast will cross each segment, but will result in only one copy on any given segment
  - Non-broadcast: used once a specific route has been selected for the logical connection

- Length bits
  - 5 bits that indicate the total length (in bytes) of the RI field
  - For initial broadcast frames, the source will set this field to '2'
    - For a single-route broadcast, this value will remain '2' for the entire discovery process since no route information is collected
  - For an all-routes broadcast, the bridge will alter this value if it is forwarding the frame
    - First bridge will increment the count by '4'; all subsequent bridges will increment the count by 2 (we'll see why in a minute)
    - RI field can be a maximum of 18 bytes, allowing for the frame to pass across 8 LAN segments connected by 7 bridges
- Direction bit
  - Indicates to the bridge how the RI field should be interpreted (forward or reverse direction) by the bridge
  - Frames traversing the ring from the originating node to the destination node will set the bit to '0'; frames going in the opposite direction will have the bit set to '1'
- Largest frame (size) bits
  - Indicates the maximum frame size that can transverse that particular portion of the route
  - Maximum frame size is the length of the L-PDU (including LLC fields and the information field)
  - Maximum frame size will be determined by several factors, such as the type of LAN segment between the bridges and segment speeds
  - Frame size values:
    - '000'- 516 bytes (minimum frame size supported for Type 1 and Type 2 services)
    - '001'- 1500 bytes (maximum frame size supported for 802.3 LANs)
    - '011'- 4472 bytes (maximum frame size supported for 4Mbps 802.5 and FDDI LANs)
    - '100'- 8144 bytes (maximum frame size for 802.4 LANs)
    - '101 - 11407 bytes
    - '110'- 17800 bytes (maximum frame size for 802.5 16Mbps LANs)
    - '111'- maximum value set when an all-routes broadcast is started

Segment Number Field(s)

- During the installation process, each segment is given a unique 3 digit number (X'000' to X'FFF')
- All bridges connected to that segment must use the same segment number
- Each bridge connected to a particular segment is assigned a single bridge number (X'0' to X'F'), which may or may not be unique
  - Parallel bridges connecting the same segments must have different bridge numbers
- Together, the segment and bridges number form a route designator
- Segment number fields:

- Each segment number field consists of 16 bits (2 bytes); 12 bits are used for the LAN segment number and 4 bits are used for the bridge number (you might say that this is more than a 2-bit operation)
- As noted above, the fields are filled in by the bridges during an all-routes broadcast route discovery
- First bridge adds 4 bytes:
  - 1st two bytes represent the segment from which the frame was copied and the bridge that copied it
  - 2nd two bytes represent the segment on which the frame is being forwarded (last 4 bits are set to '0')
- Each subsequent bridge adds 2 more bytes
  - Bridge inserts its number in the last 4 bits of the existing field
  - Segment on which the frame is being forwarded goes in the next 1 1/2 bytes (again last 4 bits (1/2 byte) are set to '0')
- This process continues until 7 bridges have been crossed, or all segments have been traversed

Note: IBM bridge programs allow the user to limit the maximum number of segments a frame will transverse to less than 7 (called the 'Hop Count Limit').

IBM Product Implementations of Source Routing

Different IBM products implement source routing differently. Although the net result will be that the source station will select the route to be used, how the route is established will vary by product. Here are a few examples of how it's done.

- OS/2 EE V1.1 & V1.2 Communications Manager (in 3270 Emulation mode)
  - 3 non-broadcast TEST frames from source
  - All-routes broadcast from source
  - Non-broadcast frame(s) back from target
- 3270 Emulation Program V3.0 (using the 802.2 interface)
  - 3 non-broadcast XID frames
  - Single-route broadcast from source
  - All-routes broadcast frame back from target
- Personal Communications/3270 (using the 802.2 interface)
  - TEST frame to self
  - 1 non-broadcast TEST frame
  - All-routes broadcast from source
  - Non-broadcast frame back from target
- NETBIOS applications (using LAN Support Program or OS/2 EE V1.1 & V1.2)
  - Single-route broadcast from source
  - All-routes broadcast frame back from target

IBM Bridge Products

- Bridge Program V2.1
  - Runs in specific models of PS/2s and PC/ATs
  - Interconnects two Token-Ring LAN segments operating at either 16Mbps or 4Mbps
  - Also supports interconnection of 2 remote Token-Ring LAN segments across a telecommunications link
- Bridge Program V2.0 (no longer available)
  - Runs in specific models of PS/2s and PC/ATs
  - Interconnects two Token-Ring LAN segments operating at either 16Mbps or 4Mbps
- Bridge Program V1.1 (no longer available)
  - Runs in specific models of PS/2s and PC/ATs
  - Interconnects two Token-Ring LAN segments operating at 4Mbps
- 8209 Bridge
  - Stand-alone unit
  - Connects two LAN segments, one being a Token-Ring and the other a CSMA/CD LAN implementing either IEEE 802.3 and/or Ethernet protocols

Feedback Questions

Directions:

First select the correct answer for each of the following questions by circling the appropriate letter. In a few minutes you will be asked to input your answers, one at a time, using the Student Response Unit.

1. Which of the following statement is true about LAN bridging?

   [✓ a.] In a source routing implementation, the MAC frame contains the necessary routing information to correctly route the frame
   b. Transparent bridges are compatible with source routing
   c. In a transparent bridge, the user must pre-load all routes to other bridges before starting the bridge
   d. Source routing bridges are better than transparent bridges because no processing is required at the source routing bridge to pass the frame between LAN segments

2. Route discovery frames which need to be propagated through all source routing bridges will have which of the following broadcast indications on (assume bridges are configured according to IBM published guidelines)?

   a. Single-route broadcast
   [✓ b.] All-routes broadcast
   c. All stations broadcast
   d. Spanning-Tree broadcast

3. (True/False) The maximum frame size that can passed through an IBM source routing bridge is 516 bytes.

   a. True
   [✓ b.] False

Additional Bridging Information

Refer to Appendix C for a more detailed discussion on the interconnection of different LAN types such as IEEE 802.3 and 802.5 LANs. An exercise is included in the material.

UNIT A: Appendix A

Appendix A - NETBIOS Commands and Frame Formats

NETBIOS NCB Commands (Application Interface)

This is a summary of Chapters 4 and 5 of the *IBM Local Area Network Technical Reference* manual (SC30-3383-2). Please refer to the manual for details on any of this information.

The following are the names of the NCBs passed back and forth between the application and NETBIOS. Each NCB will contain all of the data necessary to perform the task (i.e. buffer addresses of data, network names, session numbers, etc.)

- ADD.GROUP.NAME
  - Adds a group name to the table of names by which this node is known on the network
- ADD.NAME
  - Adds a network name to the table of names by which this node is known on the network
  - Name musty be unique throughout the network (LAN)
- CALL
  - Used to open a session with another network name
- CANCEL
  - Used to cancel outstanding commands
- CHAIN.SEND
  - Used to send data to a session partner
- CHAIN.SEND.NO.ACK
  - Used to send multiple ('chained') buffers to a session partner
  - Does not require an acknowledgement
- DELETE.NAME
  - Used to remove a name from the name table
- FIND.NAME
  - Used to locate a network name on the LAN
- HANG.UP
  - Used to end a session with a partner
- LAN.STATUS.ALERT
  - Used by the application to indicate that it wants to be made aware of temporary LAN errors lasting longer than one minute
- LISTEN

  - Used to allow a session to be established with a local name (i.e. CALL from a remote node)

- RECEIVE
  - Allows receipt of data from a session partner for certain SEND commands
- RECEIVE.ANY
  - Allows receipt of data from any session related to a specific name
- RECEIVE.BROADCAST.DATAGRAM
  - Allows receipt of a broadcast datagram message from any network name
- RECEIVE.DATAGRAM
  - Allows receipt of a datagram message from any network name
- RESET
  - Used to reset the NETBIOS interface (e.g., ends sessions, deletes names, etc.)
  - Results of this command will vary depending on specific software being used (e.g., OS/2 will work differently than LAN Support Program)
- SEND
  - Used to send data to a session partner
- SEND.BROADCAST.DATAGRAM
  - Used to send a broadcast (all stations) datagram message
- SEND.DATAGRAM
  - Used to send a datagram message to any unique or group network name
- SEND.NO.ACK
  - Used to send data to a session partner without an acknowledgement required at the NETBIOS level
- SESSION.STATUS
  - Used by the application to obtain the status of sessions for one or all local names
- STATUS
  - Used by the application to request the status of NETBIOS
  - Information returned includes adapter status and contents of counters
- TRACE
  - Activates (and deactivates) a trace of all NCBs issued to NETBIOS
- UNLINK
  - Used to provide compatibility (considered as no-op by NETBIOS)

NETBIOS Frame Protocols

When an NCB is sent to NETBIOS by the application, NETBIOS will either perform the requested action internally and return the results to the application, or will assemble a NETBIOS Frame which will be passed to the LLC layer for transmission onto the network. The same will hold true for stations receiving NETBIOS frames (i.e. either information will be passed to the application or a response NETBIOS Frame will be assembled and passed down to the LLC for transmission across the LAN.

The following are the NETBIOS Frame types.

- ADD_GROUP_NAME_QUERY
  - Checks for duplicate group names on the network
- ADD_NAME_QUERY
  - Checks for a duplicate name on the network
- ADD_NAME_RESPONSE
  - Negative response from a remote node; the name being added is already on the network
- DATA_ACK
  - Data acknowledgement for last or only frame sent
- DATA_FIRST_MIDDLE
  - Send session data; first or middle frame(s)
- DATAGRAM
  - Send application generated datagram
- DATAGRAM_BROADCAST
  - Send application generated broadcast datagram
- DATA_ONLY_LAST
  - Send session data; only frame or last frame
- NAME_IN_CONFLICT
  - Duplicate names detected in remote nodes during name searches
- NAME_QUERY
  - Request to locate a network name on the LAN
- NAME_RECOGNIZED
  - Response to a NAME_QUERY (name recognized by a network node)
- NO_RECEIVE
  - No receive command outstanding to hold receive data
- RECEIVE_CONTINUE
  - Indicates that a receive is outstanding
- RECEIVE_OUTSTANDING

  - Re-transmit last data; receive command has now been issued
- SESSION_ALIVE
  - Verify that the session is still active
- SESSION_CONFIRM
  - Acknowledgement to SESSION_INITIALIZE
- SESSION_END
  - Session termination
- SESSION_INITIALIZE
  - A session has been set up
- STATUS_QUERY
  - Request remote node status
- STATUS_RESPONSE
  - Response to STATUS_QUERY
- TERMINATE_TRACE
  - Terminates traces at remote (and local) nodes

UNIT B: Appendix B

Appendix B - Exercises

Exercise 1

Refer to attached Exercise pages 2-1 to 2-5 and answer the questions on page 2-5.

Exercise 2

Refer to attached Exercise pages 2-6 to 2-8 and answer the questions on page 2-8.

Exercise Trace Exercise

In this exercise we will analyze some traces that depict activities that you perform in the Activation Lab. The purpose of this exercise is to establish the relationship between those activities and the MAC and LLC functions discussed in prior lectures. We'll also see some of the facilities of the Token-Ring Trace and Performance Program.

MAC Frame Example

```
**************************************************************************
*  An example of a media access control (MAC) frame shown by the         *
*  Trace Analysis Summary Display of the Token-Ring Trace and            *
*  Performance Program.                                                  *
*                                                                        *
**************************************************************************
          .                .                .
          .                .                .
          .                .                .
 312 400000000002 RS 400000000002 RS <Duplicate Address Test>
 --- ------------ -- ------------ -- ------------------------
  A       A       A       A       A         A
  |       |       |       |       |         |
  |       |       |       |       |         +------ Interpretation
  |       |       |       |       +---------------- Source class -
  |       |       |       |                            RS=Ring station
  |       |       |       +------------------------ Source address
  |       |       +-------------------------------- Destination class -
  |       |                                            RS=Ring station
  |       |                                            REM=Ring error monitor
  |       |                                            RPS=Ring parameter server
  |       |                                            NM=Network manager
  |       +---------------------------------------- Destination address -
  |                                                    C00000000002=RPS
  |                                                    C000FFFFFFFF=All stations
  +------------------------------------------------ Trace record number
```

LLC Frame Example

```
**************************************************************************
*  An example of a logical link control (LLC) frame shown by the         *
*  Trace Analysis Summary Display of the Token-Ring Trace and            *
*  Performance Program.                                                  *
*                                                                        *
**************************************************************************
          .                .                 .
          .                .                 .
          .                .                 .
 164 400000000099 F0 400000000001 F0 <I> 0 0 Session_Initialize
 --- ------------ -- ------------ -- --- --- ------------------
  A       A       A       A       A   A   A   A
  |       |       |       |       |   |   |   |
  |       |       |       |       |   |   |   +--  Interpretation
  |       |       |       |       |   |   |           NETBIOS
  |       |       |       |       |   |   |           SMB
  |       |       |       |       |   |   |           SNA
  |       |       |       |       |   |   +------  Sequence number(s)
  |       |       |       |       |   +----------  LLC command/response
  |       |       |       |       +--------------  Source SAP
  |       |       |       +----------------------  Source address
  |       |       +------------------------------  Destination SAP -
  |       |                                           F0=NETBIOS
  |       +--------------------------------------  Destination address
  |                                                   C00000000080=NETBIOS
  +----------------------------------------------  Trace record number
```

Station Activation Traces - Trace 1

```
*************************************************************************
*                                                                       *
*    A PC LAN Program user is successful in initializing the           *
*    Token-Ring adapter and starting the PC LAN Program                 *
*                                                                       *
*    The user's adapter address is 400000000002                         *
*                                                                       *
*    FOR CLARITY, IRRELEVANT FRAMES WERE FILTERED OUT OF THIS TRACE     *
*                                                                       *
*************************************************************************
          .                .                .
          .                .                .
          .                .                .
 311 C000FFFFFFFF RS 400035690200 RS <Ring Purge>
 312 400000000002 RS 400000000002 RS <Duplicate Address Test>
          .                .                .
          .                .                .
          .                .                .
 319 C00000000002 RPS 400000000002 RS <Request Parms>
          .                .                .
          .                .                .
          .                .                .
 342 C00000000080 F0 C00000000002 F0 <UI>  Add_Name_Query
          .                .                .
          .                .                .
          .                .                .
 383 C00000000080 F0 C00000000002 F0 <UI>  Add_Name_Query
          .                .                .
          .                .                .
          .                .                .
```

Station Activation Traces - Trace 2

```
**************************************************************************
*                                                                        *
*    A PC LAN Program user is successful in initializing the             *
*    Token-Ring adapter but is UNSUCCESSFUL starting the LAN Program     *
*                                                                        *
*    The user's adapter address is 400000000002                          *
*                                                                        *
*    FOR CLARITY, IRRELEVANT FRAMES WERE FILTERED OUT OF THIS TRACE      *
*                                                                        *
**************************************************************************
          .                .                .
          .                .                .
          .                .                .
 286 C000FFFFFFFF RS 400035690200 RS <Ring Purge>
 287 400000000002 RS 400000000002 RS <Duplicate Address Test>
          .                .                .
          .                .                .

          .                .                .
 294 C00000000002 RPS 400000000002 RS <Request Parms>
          .                .                .
          .                .                .
          .                .                .
 317 C00000000080 F0 C00000000002 F0 <UI>  Add_Name_Query
 318 400000000002 F0 10005A107CFB F0 <UI>  Add_Name_Response
          .                .                .
          .                .                .
          .                .                .
```

Trace Activity #1

1. In TRACE 2, is 400000000002 a unique address on the ring? How do you know from this trace?

   ____

2. What is the likely cause of the Ring Purge (TRACE 1, record 311 and TRACE 2, record 286)?

   ____

3. What is the address of the active monitor at the time these traces were taken? How do you know?

   ____

4. What kind of address is the destination address in TRACE 2, record 317? What station(s) will copy this frame?

   ____

5. Why isn't the source address in TRACE 2, record 317 400000000002?

   ____

6. What is the significant difference between TRACE 1 and TRACE 2?

   ____

7. Why did the second user fail?

   ____

Messaging Traces - Trace 3

```
*************************************************************************
*                                                                       *
*    A PC LAN Program user sends a message to another user by name.    *
*                                                                       *
*    The sending station is NETA at adapter address 400000000001.       *
*    The destination is NETB at adapter address 400000000002.           *
*    There is a PC LAN Program file server at 400000000099.             *
*                                                                       *
*    FOR CLARITY, IRRELEVANT FRAMES WERE FILTERED OUT OF THIS TRACE     *
*                                                                       *
*************************************************************************
          .                .                 .
          .                .                 .
          .                .                 .
  63 C00000000080 F0 C00000000001 F0 <UI>  Name_Query
          .                .                 .
          .                .                 .
  66 C00000000080 F0 C00000000001 F0 <UI>  Name_Query
  67 C00000000080 F0 C00000000001 F0 <UI>  Name_Query
  68 C00000000080 F0 C00000000001 F0 <UI>  Name_Query
  69 C00000000080 F0 C00000000001 F0 <UI>  Name_Query
          .                .                 .
          .                .                 .
  72 C00000000080 F0 C00000000001 F0 <UI>  Name_Query
  73 C00000000080 F0 C00000000001 F0 <UI>  Name_Query
  74 400000000001 F0 400000000002 F0 <UI>  Name_Recognized
  75 400000000002 F0 400000000001 F0 <SABME>
  76 400000000001 F0 400000000002 F0 <UA>
  77 400000000002 F0 400000000001 F0 <RR> 0
  78 400000000001 F0 400000000002 F0 <RR> 0
  79 400000000002 F0 400000000001 F0 <I> 0 0 Session_Initialize
  80 400000000001 F0 400000000002 F0 <RR> 1
  81 400000000001 F0 400000000002 F0 <I> 0 1 Session_Confirm
  82 400000000002 F0 400000000001 F0 <RR> 1
  83 400000000002 F0 400000000001 F0 <I> 1 1 Data_Only_Last Send Message
  84 400000000001 F0 400000000002 F0 <RR> 2
  85 400000000001 F0 400000000002 F0 <I> 1 2 Data_Ack
  86 400000000002 F0 400000000001 F0 <RR> 2
  87 400000000001 F0 400000000099 F0 <RR> 8
  88 400000000099 F0 400000000001 F0 <RR> 5
  89 400000000001 F0 400000000002 F0 <I> 2 2 Data_Only_Last Send Message
  90 400000000002 F0 400000000001 F0 <RR> 3
  91 400000000002 F0 400000000001 F0 <I> 2 3 Data_Ack
  92 400000000001 F0 400000000002 F0 <RR> 3
  93 400000000002 F0 400000000001 F0 <I> 3 3 Session_Terminate
  94 400000000001 F0 400000000002 F0 <RR> 4
  95 400000000001 F0 400000000002 F0 <I> 3 4 Session_Terminate
  96 400000000002 F0 400000000001 F0 <DISC>
  97 400000000001 F0 400000000002 F0 <UA>
  98 400000000001 F0 400000000002 F0 <DM>
          .                .                 .
          .                .                 .
          .                .                 .
```

Messaging Traces - Trace 4

```
**************************************************************************
*                                                                        *
*    A PC LAN Program user sends a message to ALL computers by sending   *
*    a message to *. The sending computer is NETA at 400000000001.       *
*                                                                        *
*    FOR CLARITY, IRRELEVANT FRAMES WERE FILTERED OUT OF THIS TRACE      *
*                                                                        *
**************************************************************************
          .                .                .
          .                .                .
          .                .                .
  35 C00000000080 F0 C00000000001 F0 <UI>  Datagram_Broadcast
          .                .                .
          .                .                .
          .                .                .
```

Trace Activity #2

8. In TRACE 3, what is the purpose of the multiple UI frames, records 63 through 73? Are they connection oriented or connectionless?

9. What important information is carried by the source address field in TRACE 3, record 74?

10. If NETB were not on the ring, where would the message traffic stop?

11. What is happening in TRACE 3 between records 75 through 78?

12. What is the purpose of record 84 in TRACE 3?

13. From the user's perspective, what additional information is provided by sending the message to NETB rather than all computers (*).

14. Which of these traces shows an example of a logical connection?

UNIT C: Appendix C

Appendix C - Interconnection of Multiple LAN Types

Introduction

Interconnection of dissimilar LANs (i.e. Token-Ring and CSMA/CD LANs) require changes to MAC layer fields, and in some cases, the LLC fields. There are at least three reasons for this situation:

1. MAC protocols differ slightly between the two LANs
2. Some CSMA/CD LANs do not implement the 802.2 LLC protocols (i.e. Ethernet V2.0)
3. Routing schemes may be different between the two LANs (i.e. source routing vs. learning bridges)

Although bridges can still be used to interconnect the various types of LANs, they need to be more sophisticated and require additional processing to perform the connection tasks.

IBM has one product which can perform these tasks, called the 8209 LAN Bridge. Let's take a look at how it works.

Note: Although this is an Architecture course and not a product course, understanding how this product works will should help you reinforce your understanding of the architectures involved.

Frame Conversions

Ref.: *8209 Local Area Network Bridge Customer Information* manual (SA21-9994)

In each of the following configurations, the highlights of the conversion process will be noted. By now, you should be familiar enough with the terminology and acronyms to understand the concepts of what is occurring.

Prior to looking at the various conversions, one important point needs to be discussed, and that is the fact that the 802.5 and 802.3 (and Ethernet) address fields are not quite the same; they are bit reversed at the byte level. When going from one type of LAN to another, the address bit mapping must be reversed.

Note: All references here to Ethernet apply only to Ethernet (DIX) Version 2.

- 802.3 (or Ethernet) address = X '08 00 5A B2 3D 69'
- 802.5 address = X '10 00 5A 4D BC 96'
- Address conversion procedure:
  - Write the 802.3 address as separate bytes
  - Write each byte in binary (bit) format
  - Reverse the order of the bits for each byte
  - Rewrite the numbers in byte format
  - For example:
    - 1st 802.3 byte = X'01'

- Bit format = B'0000 0001'
- Reverse bit format = B'1000 0000'
- 802.5 byte format = X'80'

802.5 to Ethernet (with imbedded TCP/IP)

- MAC header prior to the DA is changed from 802.5 to Ethernet
- DA and SA left intact (though the bit order is reversed)
- 802.5 RI, DSAP, SSAP, Control and SNAP protocol ID fields dropped
- Ethernet Type field and all information remains intact
- FCS is recalculated
- 802.5 MAC trailer dropped

Ethernet to 802.5 (with imbedded TCP/IP)

- MAC header prior to the DA is changed from Ethernet to 802.5
- DA and SA left intact (though the bit order is reversed)
- 802.5 RI, DSAP, SSAP, Control and SNAP protocol ID fields are added
  - RI field is retrieved from a table in the 8209
  - Remaining fields will be given the following values:
    X 'AA AA 03 00 00 00'
- Ethernet Type field and all information remains intact
- FCS is recalculated
- 802.5 MAC trailer is added

802.5 to 802.3

- MAC header prior to the DA is changed from 802.5 to 802.3
- DA and SA left intact (though the bit order is reversed)
- 802.5 RI field is dropped
- A length field is calculated and inserted
- 802.2 fields and all information remains intact
- Pad is inserted if needed
- FCS is recalculated
- 802.5 MAC trailer dropped

802.3 to 802.5

- MAC header prior to the DA is changed from 802.3 to 802.5
- DA and SA left intact (though the bit order is reversed)
- Length field is dropped
- RI field is retrieved from a table in the 8209
- 802.2 fields and all information remains intact

- Pad is dropped if present
- FCS is recalculated
- 802.5 MAC trailer is added

802.5 to Ethernet (with imbedded LLC based protocols)

- MAC header prior to the DA is changed from 802.5 to Ethernet
- DA and SA left intact (though the bit order is reversed)
- 802.5 RI field is dropped
- An Ethernet Type field (80D5) is added along with a length field and pad
- All LLC fields and the information (complete L-PDU) are left intact
- A pad is added if needed
- FCS is recalculated
- 802.5 MAC trailer dropped

Ethernet to 802.5 (with imbedded LLC based protocols)

- MAC header prior to the DA is changed from Ethernet to 802.5
- DA and SA left intact (though the bit order is reversed)
- Ethernet Type field (80D5), length field and pad are all dropped
- 802.5 RI field retrieved from a table in the 8209
- All LLC fields and the information (complete L-PDU) are left intact
- Pad is dropped if present
- FCS is recalculated
- 802.5 MAC trailer added

ARP and RARP Conversions

- ARP and RARP are used to locate a physical address based on a known IP address; therefore, the ARP and RARP packets will have physical addresses contained in them
- Like the physical addresses in the DA and SA fields, the ARP and RARP physical address fields will have to be converted using the same conversion procedure (discussed above) as the packet crosses the bridge

Exercise 3

In this exercise, you are given the IEEE 802.5 Token-Ring MAC frame. The MAC frame is to be transferred across an IBM 8209 Bridge to an Ethernet LAN; the MAC frame contains TCP/IP protocols. Your assignment is to identify which fields will be dumped (if any), which fields will be added (if any) and which fields will be altered (if any).

802.5 MAC Frame

SD	AC	FC	DA	SA	RI	DSAP	SSAP	CONT	P_ID	TYPE	INFO	FCS	ED	FS

Directions:

1. Which fields will be dumped?

2. Which fields will be altered?

3. Which fields will be added?

4. In the box below, complete the frame format as it would appear on the Ethernet LAN; use the 802.5 MAC Frame as an example.

Ethernet Frame

UNIT D: Appendix D

Overview of OEM Architectures

What's in this appendix?

The intent of this appendix is to provide you with an overview of the architectures implemented by major OEM vendors in the LAN marketplace. It is our intent to expand this material as it becomes abailable to us, and that eventually this will be a separate follow-on education offering in the LAN curriculum.

Each vendor's offering will be shown by its relationship to the OSI seven layer model. As in the case of the IBM hardware and software, the comparison is NOT meant that the products are OSI compliant.

The vendor LAN implementations that will be reviewed are:

- General Motors: MAP
- Boeing: TOP
- Xerox: XNS
- 3Com: 3+
- Novell: NetWare
- Apple: AppleTalk
- AT&T: StarLan

OEM ARCHITECTURES

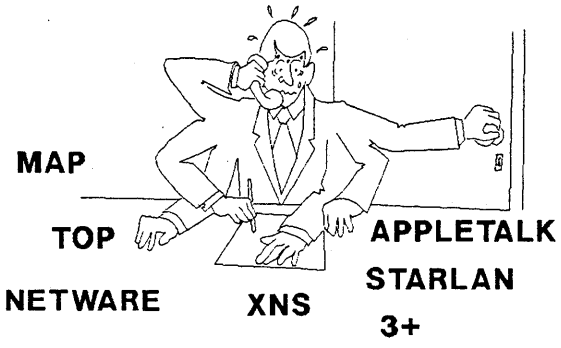

LANA705

Introduction

Since there are minimal standards in layers 3-7 of the OSI model, just about every networking company has come up with its own architecture for these layers. In this appendix, we will be looking at some of the major network players and how they have implemented these upper layer protocols.

A major reference for this appendix is the Hewlett Packard chart which shows most of the architectures in a concise format. The one architecture not on the chart that will be discussed is Apple's AppleTalk network.

For additional details about these architectures, refer to James Martin's *Local Area Networks - Architectures and Implementations.*

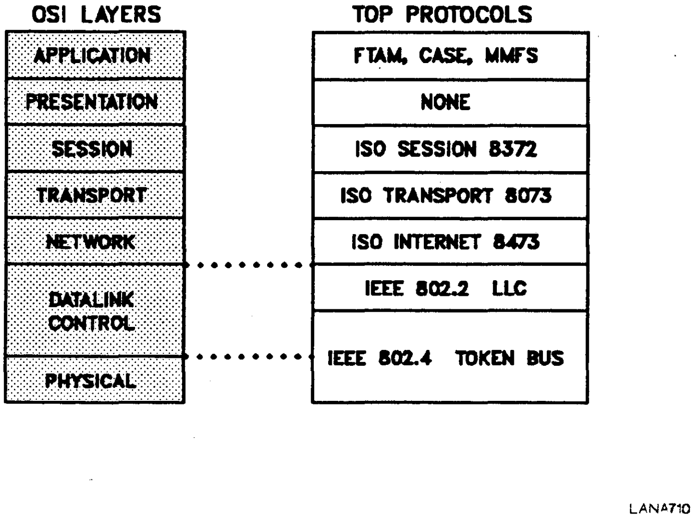

General Motors: MAP

Manufacturing Automation Protocol (MAP) was created by General Motors, along with others, to address the area of factory automation. It closely adheres to the functionality defined in the OSI seven-layer model. Each layer is summarized below:

- LAYER 1 - Physical layer
  - Two implementations: broadband at 10Mbps or baseband at 5Mbps
  - Data Encoding Scheme: Phase-coherent Frequency Modulation
- LAYER 2 - MAC SUBLAYER
  - IEEE 802.4 Token Bus
  - Only 6 byte addresses
- LAYER 2 - LLC SUBLAYER
  - Uses connectionless services (Type 1)
  - Optional use of acknowledged connectionless services (Type 3)
- LAYER 3 - Network Layer
  - Relays and routes messages between network nodes
  - Message segmentation and reassembly if message exceeds data unit size
  - For WAN services, X.25 is specified
  - Address format:
    - Authority and Format ID
    - Network ID
    - Primary Subnet ID

    - End System Address
  - Optional features
    - Security checking
    - source routing
    - priorities
    - route recording
    - quality of service maintenance
- LAYER 4 - Transport Layer
  - Provides services for data transfer and connection management
  - OSI's Class 4 service is used. Its functions include:
    - Get and release transport connections and transmit data over them
    - Flow control
    - Multiplex several transport connections
    - Error detection and recovery using:
      - Checksums (single packet error detection)
      - Sequence Numbers (multi-packet flow control)
- LAYER 5 - Session Layer
  - Implements full-duplex communications
  - Connection establishment and termination
  - Data transfer
- LAYER 6 - Presentation Layer
  - No functions implemented in this layer
- LAYER 7 - Application Layer
  - File Transfer, Access, and Management (FTAM)
  - Common Application Service Elements (CASE)
  - Manufacturing Message Format Standard (MMFS)
  - Directory Services

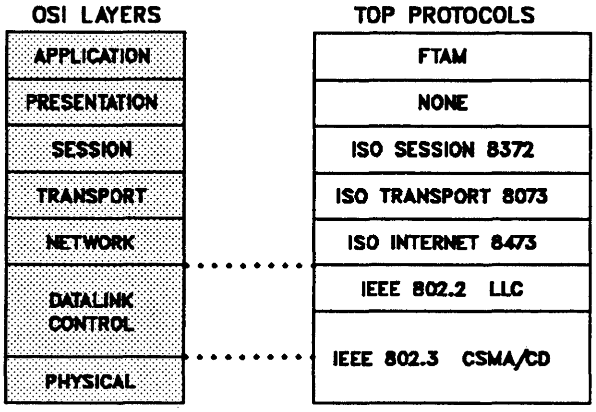

LANA715

Boeing: TOP

Technical and Office Protocols (TOP) was sponsored by the Boeing Company to address engineering, manufacturing, and general office applications. It has many similarities to MAP. In fact, from the Logical Link Control sublayer of the Data Link Control layer to the Presentation layer (Layer 6), the protocols are identical.

Where they differ is in the physical layer, MAC sublayer, and application layer. These differences are as follows:

- LAYER 1 - Physical Layer
  - Baseband transmission
  - Coaxial Cable
  - 10 Mbps over 500 meter segments (10Base5)
- LAYER 2 - MAC Sublayer
  - IEEE 802.3 CSMA/CD
  - 6 byte addressing
- LAYER 7 - Application Layer
  - Specifies only File Transfer, Access, and Management (FTAM)
  - CASE, X.400 messaging, and other protocols under consideration

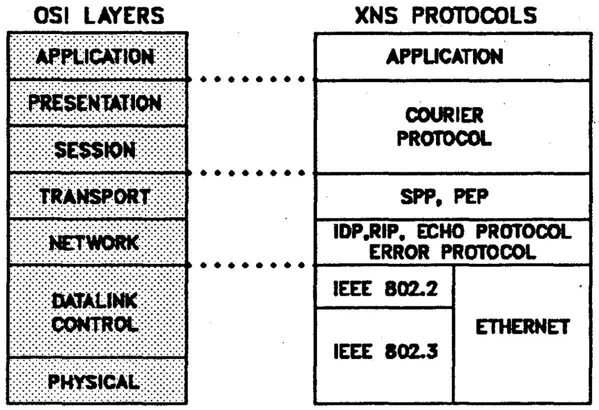

LANA720

Xerox: XNS

The Xerox Network Systems (XNS) architecture was created by Xerox Corporation for internal use to interconnect their products and was later made generally available. Currently, part of the Network layer, the Internetwork Datagram Protocol, is used by both Xerox and Novell. Layer 1 through 7 summaries follow.

- LAYER 1 - Physical Layer
  - Bus Topology using telephone twisted-pair, coaxial cable, or fiber.
- LAYER 2 - Data Link Control Layer
  - Ethernet or IEEE 802.3 CSMA/CD
  - If IEEE 802.3, then IEEE 802.2 Logical Link Control also
- LAYER 3 - Network Layer
  - Internetwork Datagram Protocol - connectionless datagram service.
    - 4-byte Network Address
    - 6-byte Station Address (Xerox calls each station a Host)
    - Does not provide reliable data transfer
  - Error Protocol - Standardizes error reporting
  - Echo Protocol - Verifies existence and correct operation of a station and the route to it.
  - Routing Information Protocol (RIP) - Provides exchange of routing information between network nodes
- LAYER 4 - Transport Layer

  - Sequence Packet Protocol (SPP) - Multi-packet reliable data transfer
  - Packet Exchange Protocol (PEP) - Single packet with response for reliable data transfer

- LAYERS 5 and 6 - Session and Presentation Layers

  XNS Courier Protocol resides in these layers.

  - Translates request from the application layer into packets of information suitable for sending across the network.

- LAYER 7 - Application Layer

  A wide variety of applications use the XNS communication services

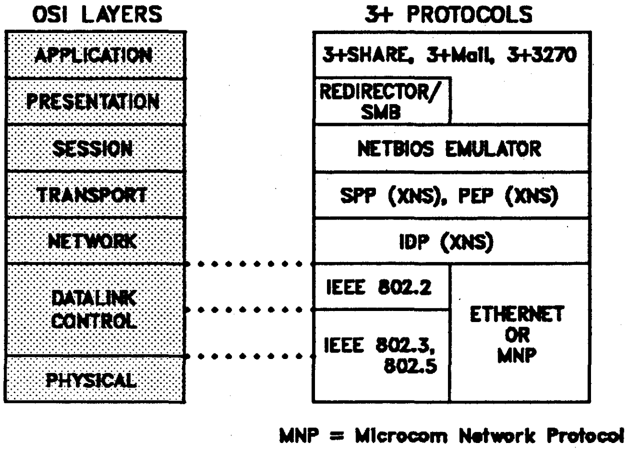

LANA725

3COM: 3+

3Com Corporation is a major LAN vendor (3Com stands for Computers, Communication, and Compatibility). 3COM offers hardware and software products for many LAN architectures, including products compatible with IBM's token ring. Functionality supported by each layer of architecture is detailed below.

- LAYER 1 - Physical Layer
  - Since 3Com-3+ supports both the 802.2 and Ethernet interfaces at Layer 2, it can run on many physical networks including coaxial cable, Telephone Twisted Pair, fiber and shielded twisted pair.
  - Can use IBM's token ring adapters or its own
- LAYER 2 - Data Link Layer
  - Ethernet
  - IEEE 802.2 Logical Link Control with the following MAC layers:
    - IEEE 802.3 CSMA/CD
    - IEEE 802.5 Token Passing Ring
  - Microcom Network Protocol - Asynchronous communication protocol
- LAYERS 3 and 4 - Network and Transport Layers

  3Com uses a protocol package that implements both layers 3 and 4. It is called MS-DOS Internal Network Drive Protocol (MINDP) and implements XNS protocols as summarized below:

  - Layer 3

    - Internetwork Datagram Protocol
  - Layer 4
    - Sequenced Packet Protocol
    - Packet Exchange Protocol
- LAYER 5 - Session Layer
  - 3Com provides a NETBIOS interface to layer 6.
  - Provides network names and session establishment
- LAYER 6 - Presentation Layer
  - Redirector/SMB (Server Message Block) Protocol
  - Provides compatibility with IBM's PC LAN Program
  - Intercepts DOS I/O requests that need to be redirected to a server and generates SMB control blocks
- LAYER 7 - Application Layer
  - 3+Share - Network Operating System for the IBM environment
  - 3+Mail - Electronic mail services
  - 3+3270 - Emulation of a 3270 terminal or 3287 printer

In addition to network access, 3Com Corporation provides several network connectivity products

- 3+Remote - Dialup access from remote PC
- 3+NetConnect - Gateway between 3+ Networks
- 3+Route - Connection of remote 3+ Networks

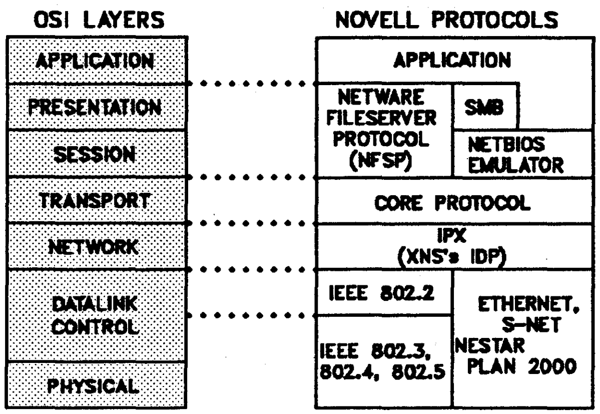

Novell: Netware

Novell's Netware provides a number of services that work on a variety of LANs. These services include:

- file sharing
- printer sharing
- electronic mail
- remote access
- interconnection of NetWare-compatible LANs
- file and record locking
- access security

Netware implements Layers 3-5 primarily with some support software in layers 6 and 7. Layers 1 and 2 can be Ethernet or any other LAN implementations that support IEEE 802.2 Logical Link Control. A summary of the Netware layered architecture is given below:

- LAYERS 1 and 2 - Physical and Data Link Layers
  - CSMA/CD networks like IBM PC Network and AT&T StarLan.
  - Token ring architectures like IBM's Token Ring Network and Proteon's ProNET.
  - Proprietary networks such as Nestar Plan 2000 and Novell's own S-Net.
- LAYER 3 - Network Layer
  - Internet Packet Exchange (IPX) Protocol
    - An implementation of Xerox's XNS Internetwork Datagram Protocol
      - 4-byte Network Address

        - 6-byte Station Address
        - Does not provide reliable data transfer

- LAYER 4 - Transport Layer
    - NetWare Core Protocol (NCP)
        - Common (Core) communications software for both the NETBIOS Emulator and the Netware File Service Protocol.
        - Ensures integrity of data
- LAYERS 5, 6, and 7 - Session, Presentation and Application Layers
    - NETBIOS Emulator
    - Netware File Server Protocol
        - Connection control
        - File maintenance services
        - Directory maintenance
        - File/record locking
        - Directory maintenance
        - User maintenance
        - File server statistics
        - Communications
        - Printing services
        - Application software copy protection

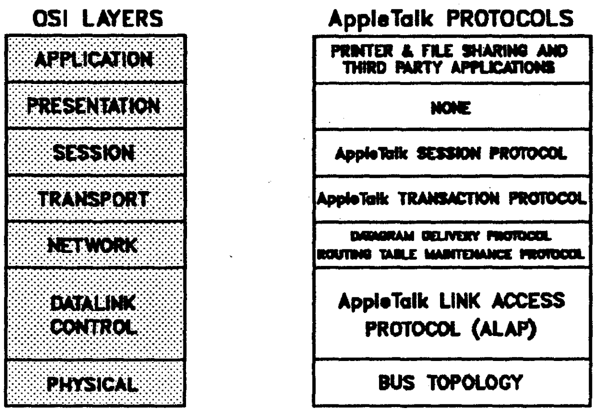

OSI LAYERS	AppleTalk PROTOCOLS
APPLICATION	PRINTER & FILE SHARING AND THIRD PARTY APPLICATIONS
PRESENTATION	NONE
SESSION	AppleTalk SESSION PROTOCOL
TRANSPORT	AppleTalk TRANSACTION PROTOCOL
NETWORK	DATAGRAM DELIVERY PROTOCOL ROUTING TABLE MAINTENANCE PROTOCOL
DATALINK CONTROL	AppleTalk LINK ACCESS PROTOCOL (ALAP)
PHYSICAL	BUS TOPOLOGY

LANA735

Apple Computer: AppleTalk

AppleTalk is implemented and included in both the Macintosh hardware and software as well as in some peripheral devices. All that is required is to cable the devices together. If a server is desired, it comes at an additional cost. Its features include:

- Maximum of 32 devices
- Maximum total cable length of 300 meters
- Data rate of .23Mbps
- Doesn't follow IEEE 802 standards

Third-party products extend AppleTalk to allow:

- Non-Apple devices such as:
  - IBM PCs
  - DEC VAX minicomuters
  - Hewlett-Packard computers
  - Various UNIX computers
- File sharing
- Electronic mail
- Communications Servers
- Print spooling
- Gateways for dial-up access to AppleTalk network
- Interconnection with Ethernet networks

- Extended number of devices
- Longer distances

The architecture consists of:

- LAYER 1 - Physical Layer
  - Bus or tree topology
  - FM-0 encoding
  - Signalling is based on EIA RS-422
- LAYER 2 - Data Link Layer
  - AppleTalk Link Access Protocol (ALAP)
  - Carrier Sense Multiple Access. If no carrier, transmits
    - If collision, higher layer requests retry. Not handled in layer 2
  - Addressing: 8 bit node ID
    - 0-127: user node ID
    - 128-254: server node ID
    - Max data size: 600 bytes
- LAYER 3 - Network Layer
  - Datagram Delivery Protocol
    - Error recovery left to higher layers
    - End devices are called socket clients
    - Logical connection called a socket
    - Network address consists of:
      - 2-byte network address
      - 1-byte node ID
      - 1-byte socket number
  - Routing Table Maintenance Protocol - Used to update complete routing table in routing nodes
- LAYER 4 - Transport Layer
  - AppleTalk Transaction Protocol
    - Provides reliable transmission between socket clients
- LAYER 5 - Session Layer
  - AppleTalk Session Protocol (ASP)
    - Establishes a session between network entities
- LAYER 6 - Presentation Layer
  - No functions at this layer
- LAYER 7 - Application Layer
  - Printer and file sharing

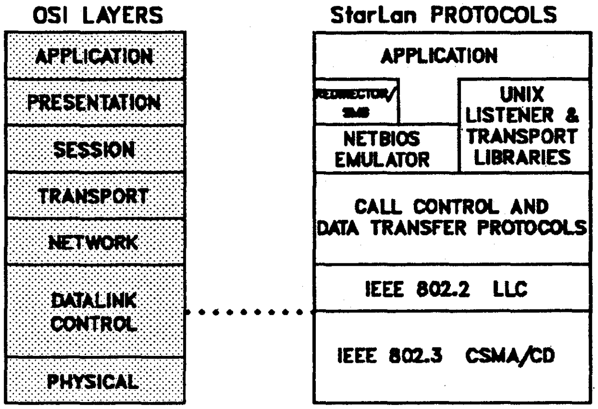

AT&T: StarLan

StarLan is an inexpensive, telephone twisted-pair implementation of an IEEE 802.3 CSMA/CD network. It is known as 1Base5 which stands for 1 Mbps and a segment length of approximately 500 meters.

The layers of the architecture are explained below.

- LAYER 1 - Physical Layer
  - Daisy chain configuration:
    - 10 stations per single length of cable
    - maximum length of cable: 400 feet
  - Star configuration
    - Network Extension Unit connects up to 11 daisy chains
    - Up to 12 Network Extension Units can be interconnected
  - 1 Mbps
- LAYER 2 - Data Link Layer
  - Standard IEEE 802.3 Media Access Control
  - IEEE 802.2 Logical Link Control
- LAYERS 3 and 4 - Network and Transport Layers
  - Call Control and Data Transfer Protocols

AT&T StarLan

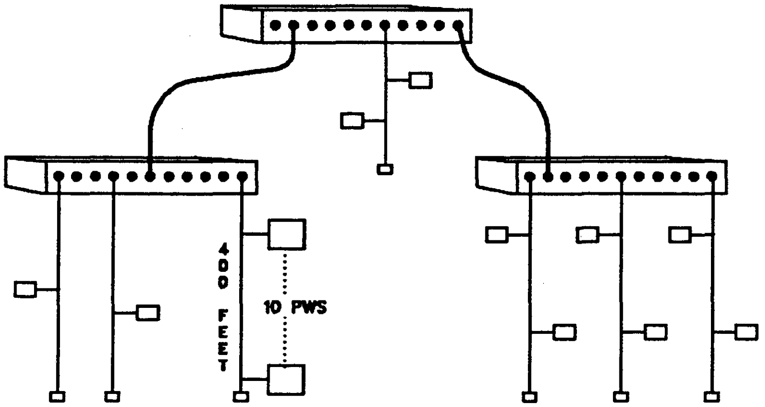

LANA745

- LAYER 5 - Session Layer
  - Two interfaces supported:
    - UNIX Listener and Transport Libraries
    - NETBIOS Emulation. This interface lacks the following:
      - Remote IPL
      - Unlink
      - Two network access units in one network station
      - Send and receive timeout values
      - Direct I/O
      - Setting and getting a post address
- LAYER 6 - Presentation Layer
  - DOS Redirector/Server Message Block
- LAYER 7 - Application Layer
  - File sharing with file and record locking
  - File transfer
  - Printer sharing
  - Electronic mail
  - Asynchronous remote access
  - Command line or menu-driven interface

Session 2 - Overview of OSI Standards and Protocols

Introduction

As vendors work toward the goal of interoperability, we will begin to see products that implement services that are ISO-OSI compliant. In this session, we will take a brief look at the OSI compliant standards on which these products are based. Examples of these standards include:

- FTAM
- X.400
- X.500
- CMIP and CMIS