LOCAL AREA COMPUTER NETWORKS

Alfred C. Weaver, Ph.D.

Computer Science Report No. RM-85-07
November 18, 1985

LOCAL AREA COMPUTER NETWORKS

Alfred C. Weaver
Department of Computer Science
University of Virginia

(slide title)

Dr. Alfred C. Weaver
Department of Computer Science
Thornton Hall
University of Virginia
Charlottesville, Virginia 22903
(804) 924-1042

OUTLINE

1. INTRODUCTION

- Definitions
- Networks vs. local area networks
- Benefits and dangers
- Examples

2. DESIGN ISSUES

- Topology: arbitrary, star, bus, ring, tree
- Service types: virtual circuits vs. datagrams
- Media: twisted pair, coax, fiber optics, radio
- Protocols: choices and functions
- Suitability: matching LANs to needs

OUTLINE

3. ISO AND IEEE STANDARDS

- The ISO seven layer model
- Physical Layer
- Data link layer
- Network layer
- IEEE committee 802

4. IEEE 802.3 CONTENTION BUS

- Ethernet
- Physical and data link specifications
- Collisions and retransmissions
- Performance

OUTLINE

5. IEEE 802.4 TOKEN BUS

- Bus architecture
- Physical and data link specifications
- Token passing
- Frame formats
- Logical ring membership
- Access classes (priorities)
- Performance

OUTLINE

6. IEEE 802.5 TOKEN RING

- Ring architecture
- Similarities and differences from the token bus
- Token passing
- Frame formats
- Priorities
- Acknowledgements

7. GENERAL MOTORS MAP PROTOCOL

- Summary of usage
- Dependence on 802.4 token bus

8. SUMMARY

SECTION 1

INTRODUCTION

- Definitions
- Networks vs. local area networks
- Benefits and dangers
- Examples

COMPUTER NETWORK

"A COMMUNICATIONS MECHANISM USED TO INTERCONNECT GEOGRAPHICALLY SEPARATED BUT AUTONOMOUS COMPUTERS."

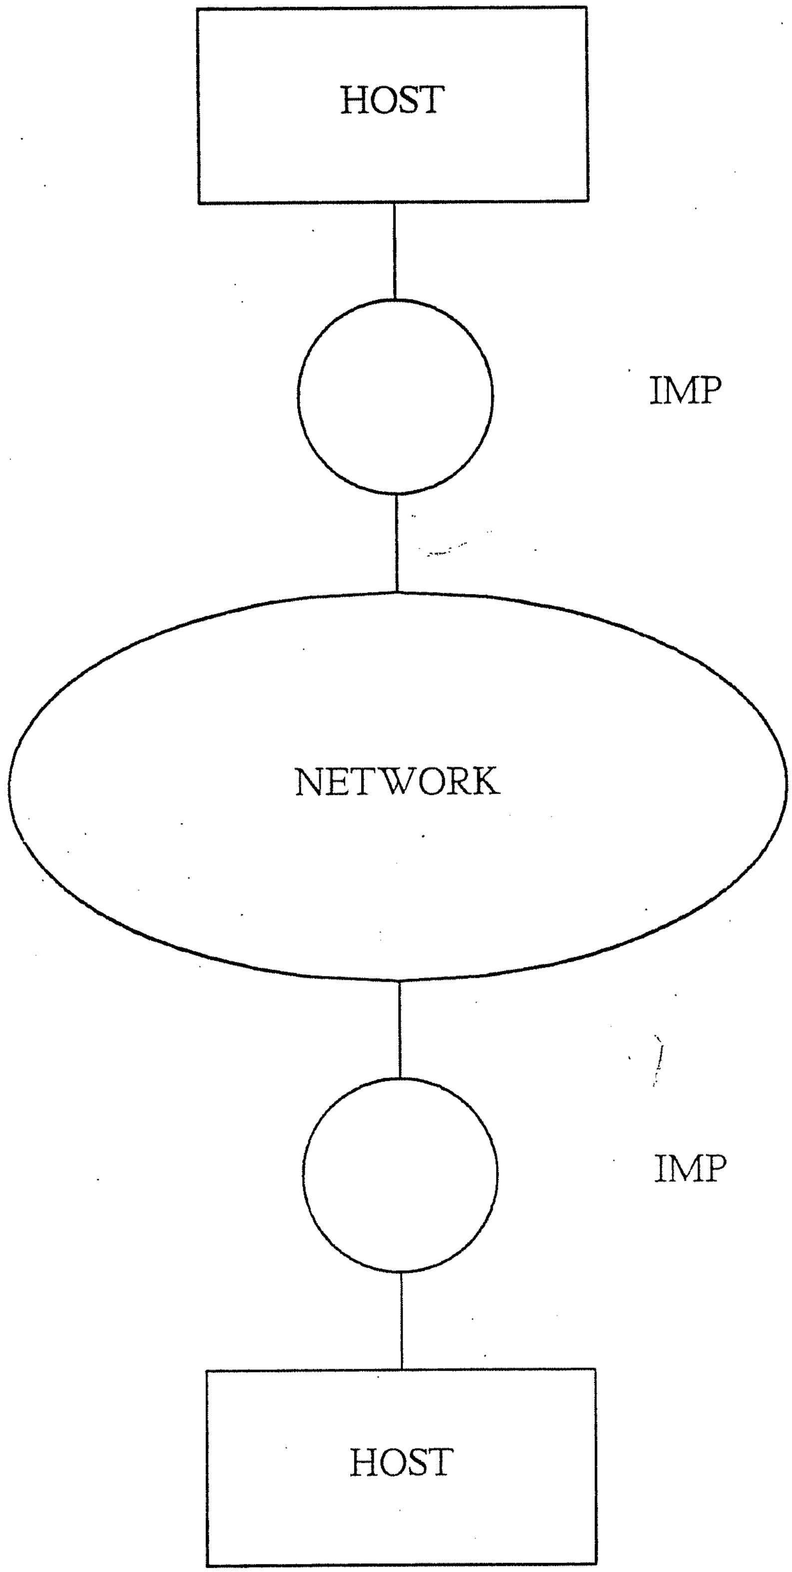

(slide host-imp)

COMPUTER NETWORK ATTRIBUTES

- GEOGRAPHICALLY DISTRIBUTED
- NETWORK-WIDE OPERATING SYSTEM
- MIXED DATA RATES
- TIME-VARYING TOPOLOGY

ATTRIBUTES

INTERPROCESSOR DISTANCE	PROCESSORS IN SAME	TYPE
0.1 m	CIRCUIT BOARD	DATA FLOW MACHINE
1 m	SYSTEM	MULTIPROCESSOR
10 m	ROOM	LOCAL AREA
100 m	BUILDING	NETWORK
1 km	CAMPUS	
10 km	CITY	LONG-HAUL
100 km	COUNTRY	NETWORK
1,000 km	CONTINENT	INTERCONNECTED
10,000 km	PLANET	NETWORKS

LOCAL AREA NETWORKS

"A COMMUNICATIONS MECHANISM USED TO INTERCONNECT GEOGRAPHICALLY ADJACENT COMPUTERS, PERIPHERALS, AND CONTROL DEVICES."

LAN ATTRIBUTES

- SHORT DISTANCES
- HIGH DATA RATES
- SINGLE ENTITY OWNERSHIP
- LOW ERROR RATE

LAN BENEFITS

- INCREMENTAL UPGRADE
- RELIABILITY
- AVAILABILITY
- SURVIVABILITY
- RESOURCE SHARING
- MULTIVENDOR SUPPORT
- IMPROVED PERFORMANCE
- SINGLE TERMINAL ACCESS
- FLEXIBILITY OF LOCATION
- TASK INTEGRATION

POTENTIAL PITFALLS

- INTEROPERABILITY IS NOT ABSOLUTELY GUARANTEED
- SECURITY IS NOT INHERENT
- PRIVACY IS NOT SUPPORTED

COMPUTER ROOM

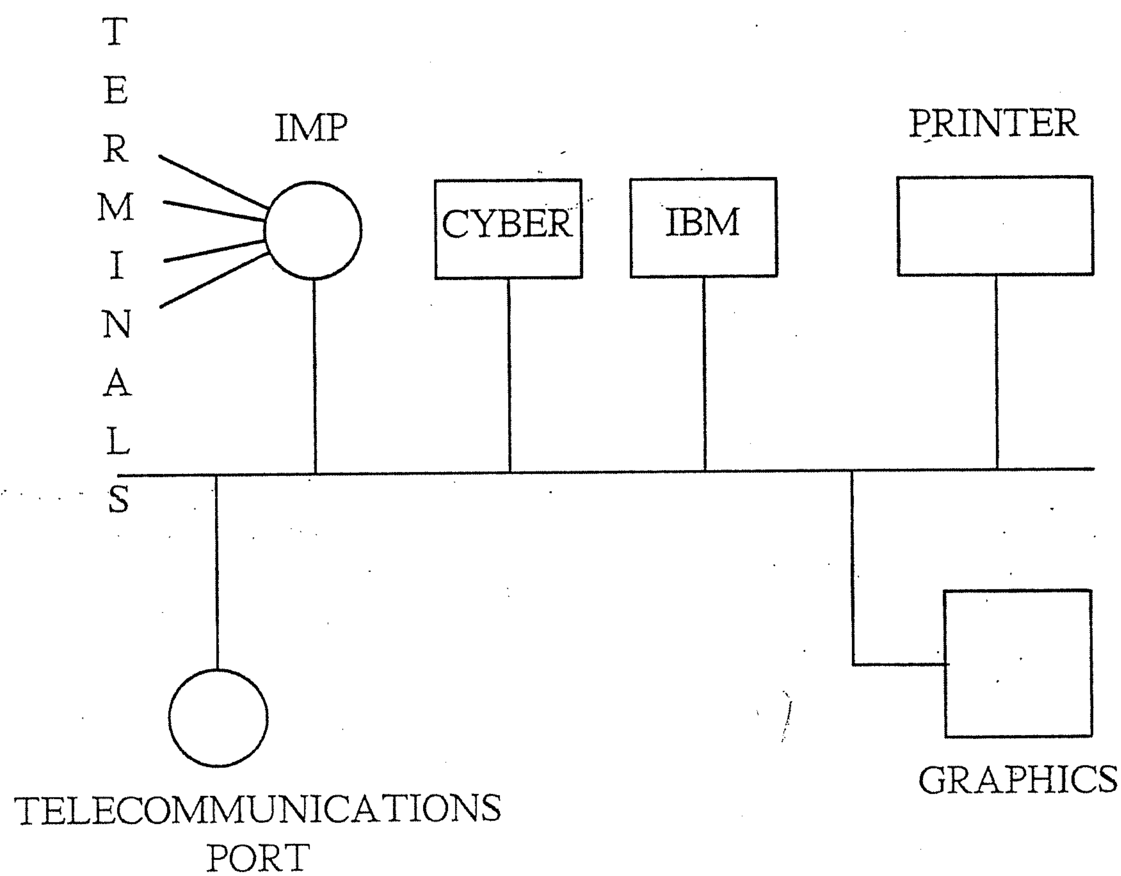

OFFICE AUTOMATION

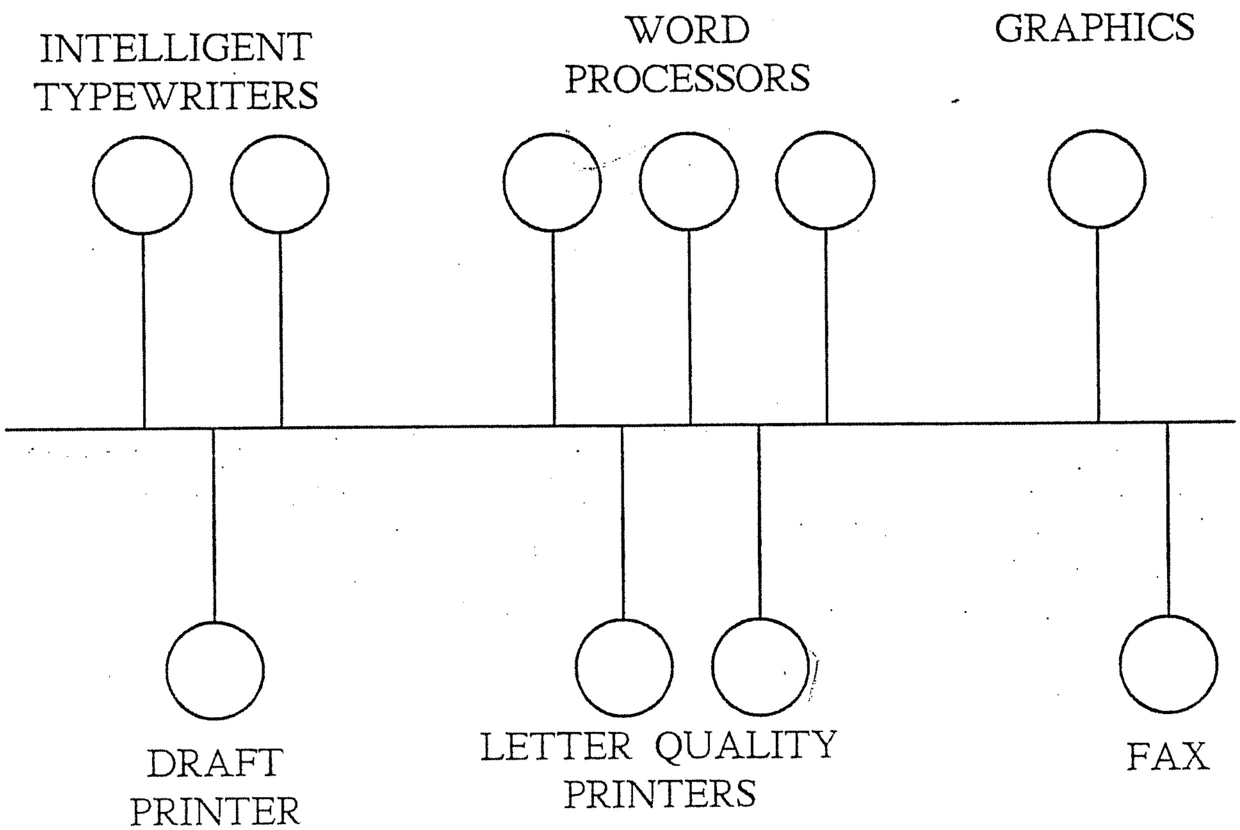

ENVIRONMENTAL CONTROL

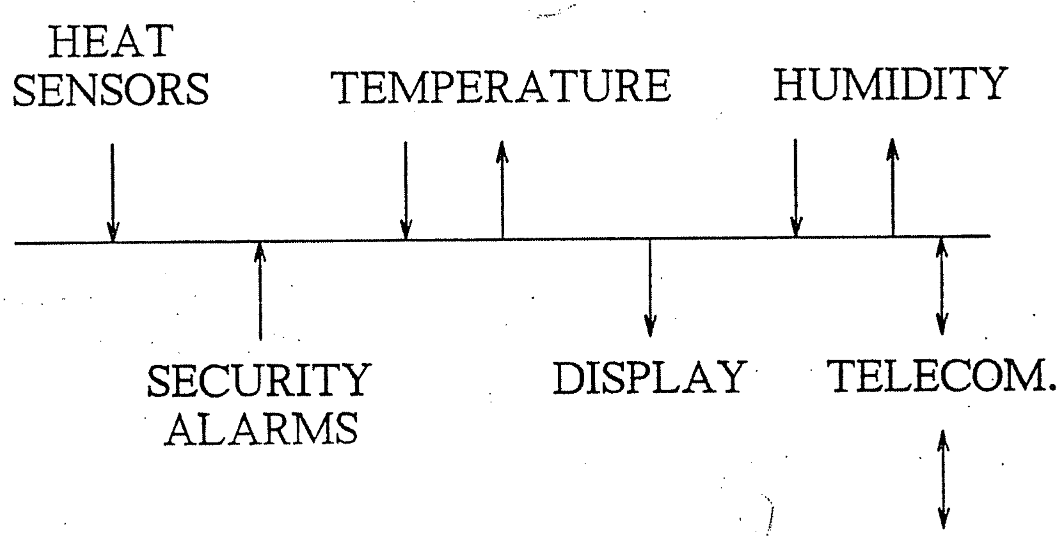

(slide environmental)

PERSONAL COMPUTERS

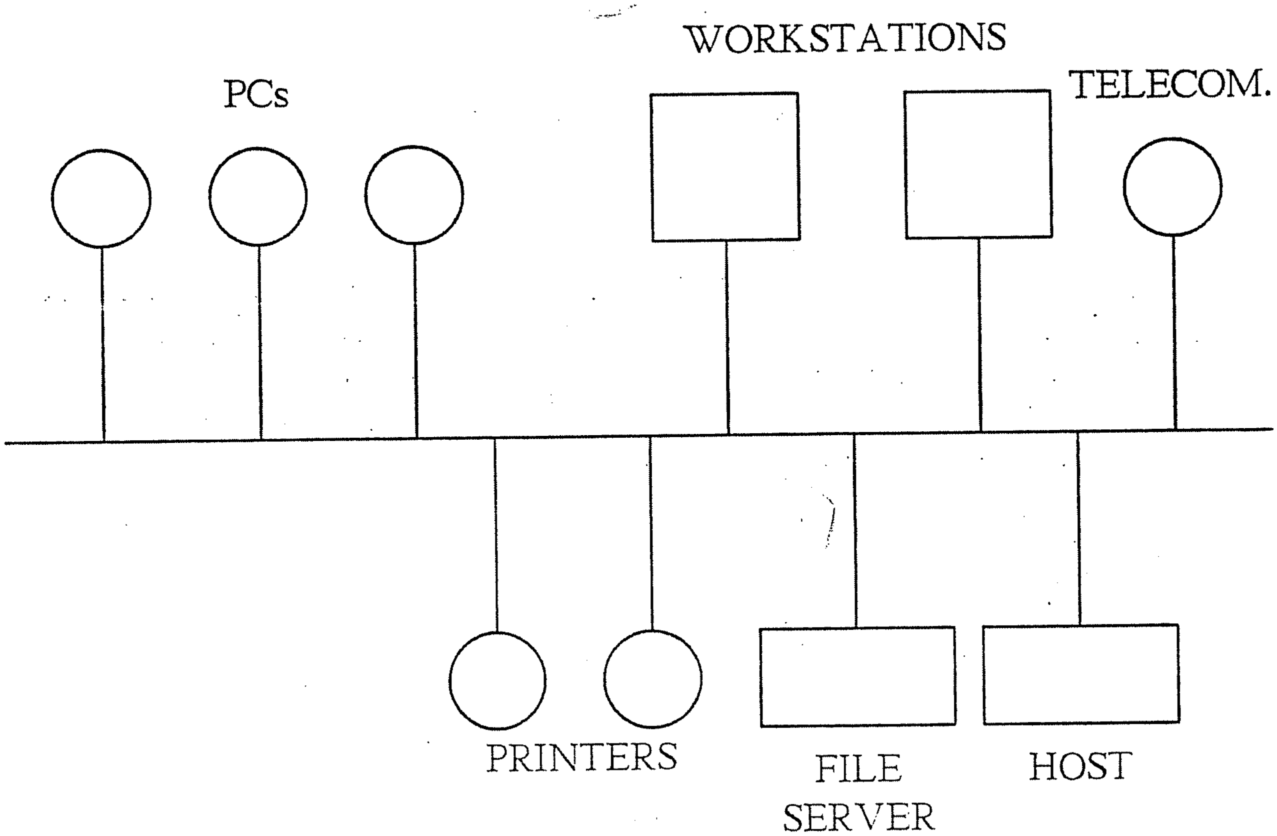

COMMUNICATIONS

- TELEPHONE (REAL-TIME VOICE)
- TV (REAL-TIME VIDEO)
- FAX (IMAGES)
- GRAPHICS
- TEXT
- COMPUTER DATA

SECTION 2

DESIGN ISSUES

- Topology: arbitrary, star, bus, ring, tree
- Service types: virtual circuits vs. datagrams
- Media: twisted pair, coax, fiber optics, radio
- Protocols: choices and functions
- Suitability: matching LANs to needs

TOPOLOGY
ARBITRARILY CONNECTED

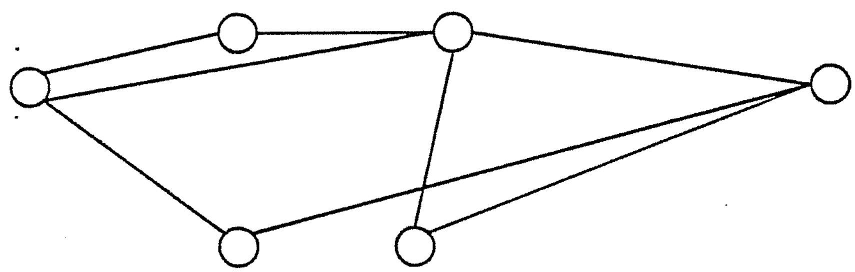

ADVANTAGES

- BANDWIDTH MATCHES NEEDS
- REDUNDANT PATHS

DISADVANTAGES

- WASTED BANDWIDTH
- ROUTING
- WIRING COSTS

STAR

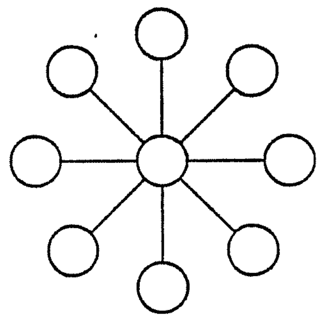

Advantages

- Simple access protocol
- Known transmission distance
- Two hop routing

Disadvantages

- Single point of failure
- All wiring to the "center"
- Sharing typically accomplished via time-slice

RING

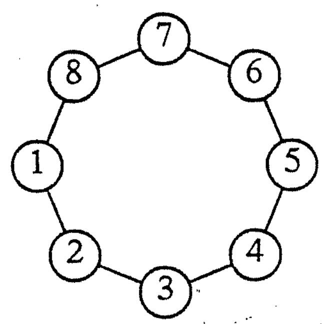

Advantages

- Simplicity
- All wiring is point-to-point; permits mixed media
- Message regenerated at each node

Disadvantages

- Failure of one link may break the ring
- Must be wired to a physical neighbor
- More complex management since message goes completely around the ring

BUS

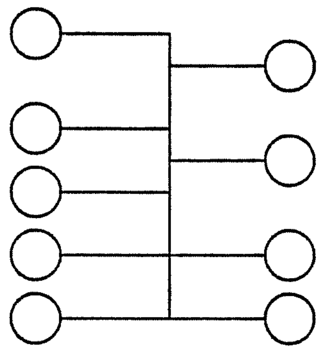

Advantages

- Simplicity of installation
- Potential for maximum channel utilization
- Stations sense common network state
- Twisted-pair and coax media well understood

Disadvantages

- Needs complex access control
- Limitations on bus length

TOPOLOGY

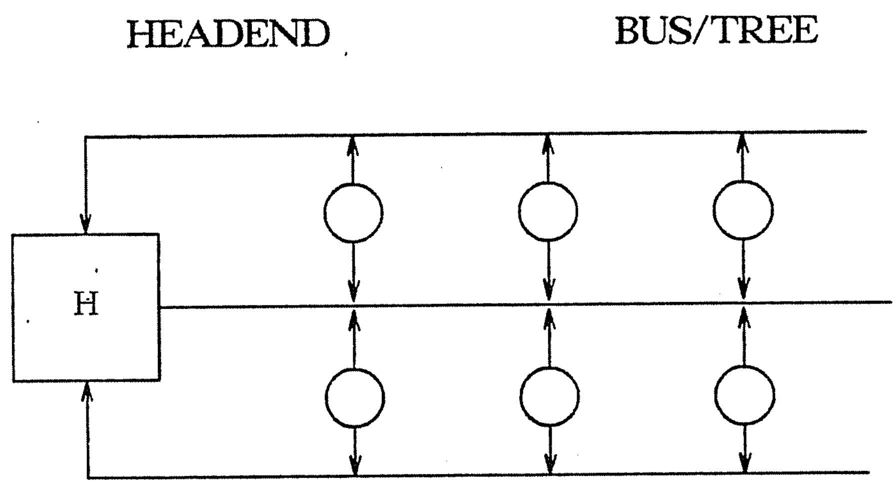

ADVANTAGES

- HIGH SPEED, MULTI-CHANNEL
- CATV TECHNOLOGY
- INTEGRATES DATA, VOICE VIDEO

DISADVANTAGES

- CABLE LENGTH
- COST

SERVICE TYPE

- CIRCUIT SWITCHING
- MESSAGE SWITCHING
- PACKET SWITCHING
  - VIRTUAL CIRCUITS
  - DATAGRAMS

CIRCUIT SWITCHING

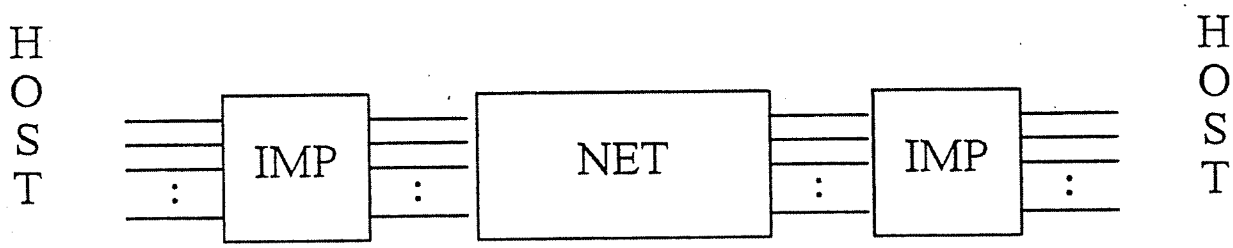

ADVANTAGES

- END-TO-END "COPPER" PATH
- ALL OVERHEAD IN SETUP
- SEQUENTIAL

DISADVANTAGES

- REQUIRES SETUP
- STATIC PATH
- NO SHARING

MESSAGE SWITCHING

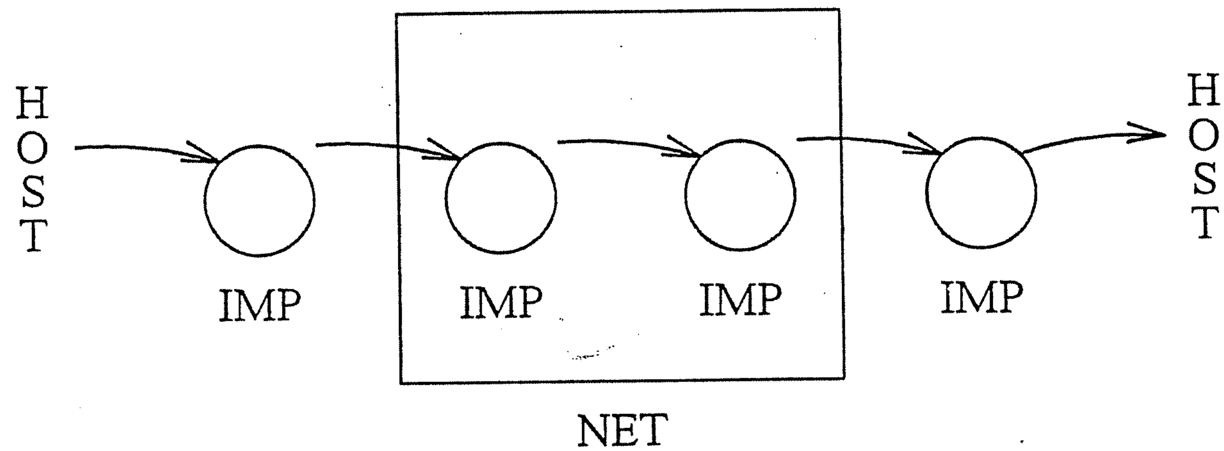

- STORE-AND-FORWARD NETWORK
- MESSAGE SAVED AT EACH IMP
- MESSAGE ARRIVES ALL AT ONCE
- MAXIMUM DELAY
- MAXIMUM MEMORY/STORAGE

PACKET SWITCHING

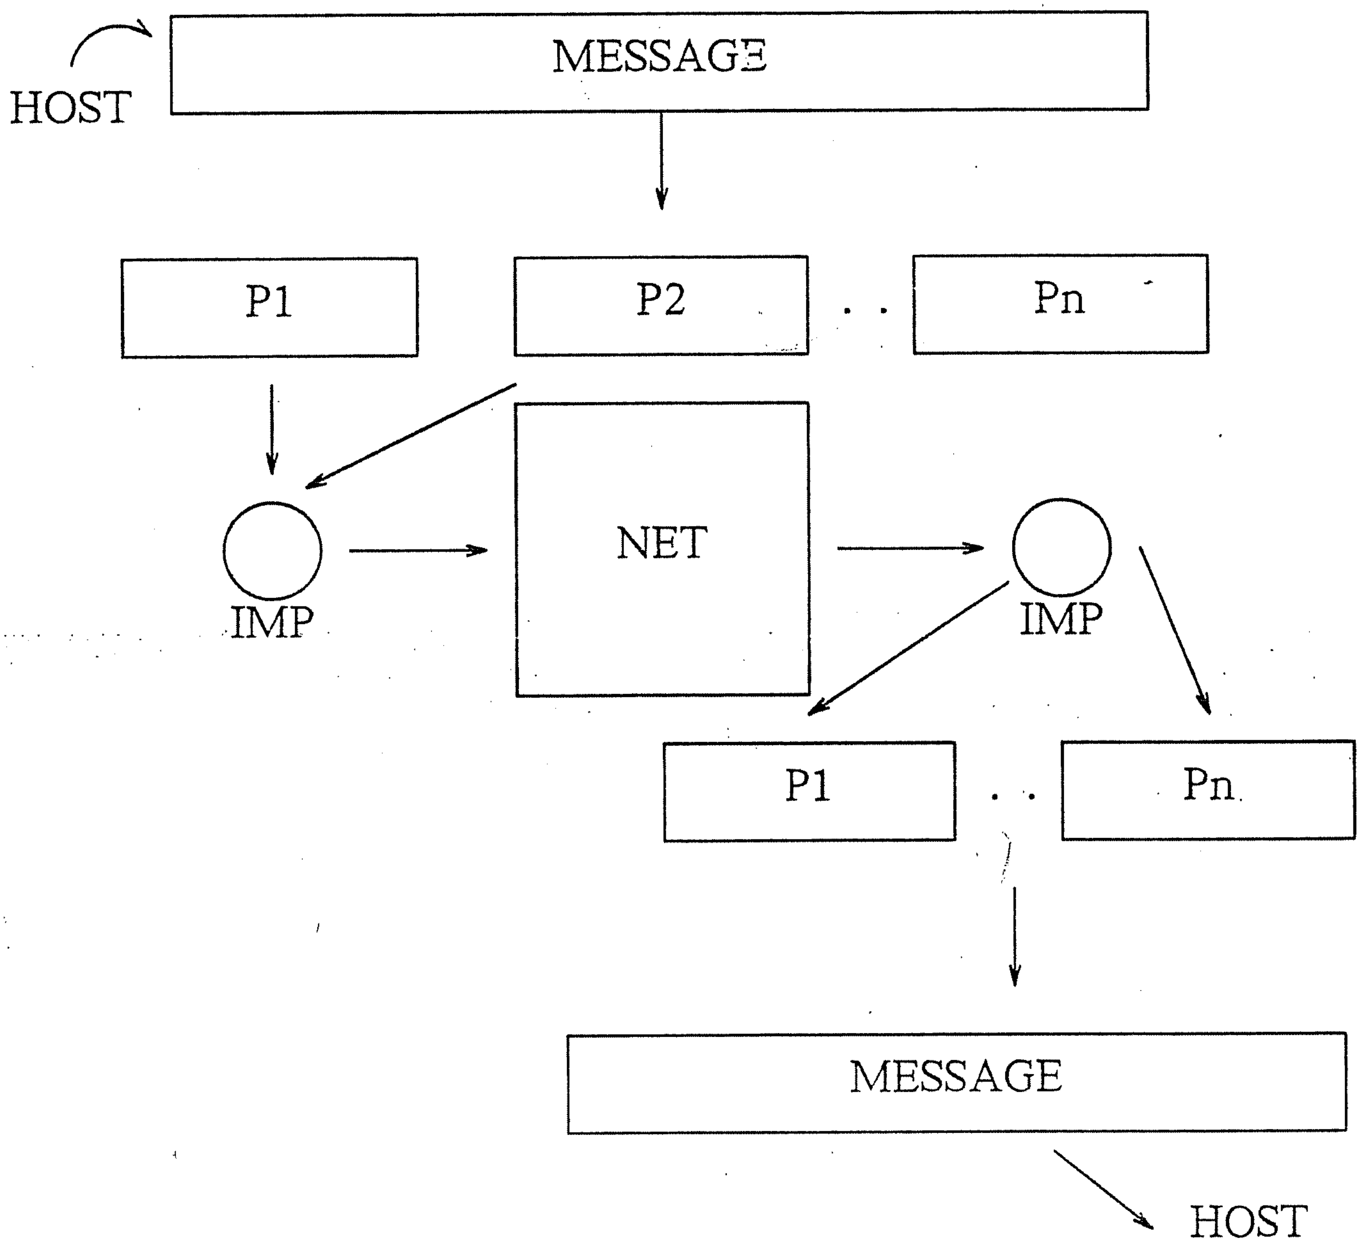

PACKET SWITCHING

VIRTUAL CIRCUITS	DATAGRAMS
ROUTE DECIDED ONCE PER SESSION	ROUTE DECIDED PER PACKET
NO ADDRESSES ON PACKETS, JUST VC#	ADDRESS ON EACH PACKET
SEQUENTIAL	NOT SEQUENTIAL
ROUTE FIXED	ROUTE VARIABLE
NO SHARING OF CIRCUIT	SHARED

VIRTUAL CIRCUITS

MSG → PACKETS

C_1	CONNECT
C_2	ACCEPT
P_1	$DATA_1$
⋮	
P_n	$DATA_n$
C_3	DISCONNECT
C_4	CLEAR

DATAGRAMS

MSG → PACKETS

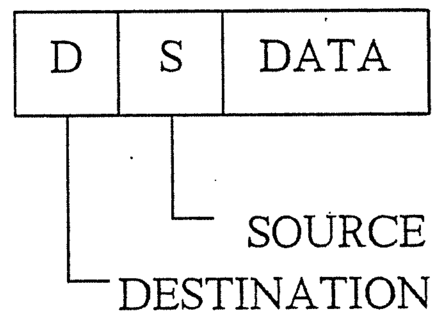

MEDIA
TWISTED PAIR

Two insulated wires

Spiral wound

Bundled and shielded

Usable for analog and digital

Needs analog amplifiers every 5-6 km

Needs digital repeaters every 2-3 km

Total circuit limited to ~15 km

Good for point-to-point and multi-point

Low cost

MEDIA
COAX CABLE

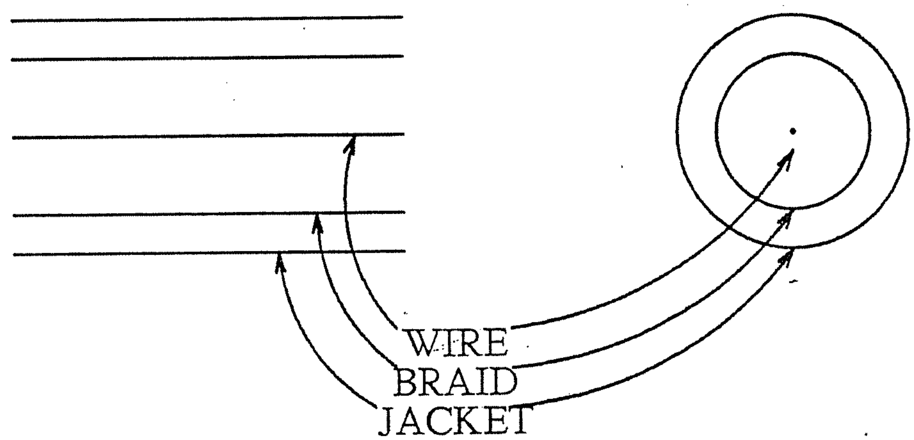

CATV 75 Ω
ANALOG (300 - 400 MHz)
DIGITAL ONLY (50 Mbps)
1000's DEVICES

BASEBAND 50 Ω
DIGITAL ONLY (50 Mbps)
100's DEVICES

EASY TO INSTALL
CATV WIDELY UNDERSTOOD
GOOD INSTALLATION TOOLS

DISTANCE: FEW KM
VERY GOOD NOISE IMMUNITY
MEDIUM COST

MEDIA
FIBER OPTICS

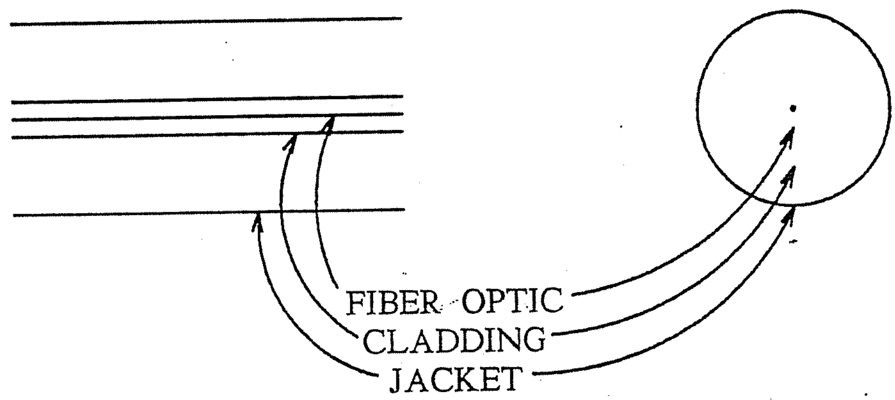

LIGHT SOURCE
LED
LASER

LIGHT DETECTOR
PIN PHOTODIODE
APD (AVALANCHE PHOTODIODE)

BEST FOR POINT-TO-POINT
BASEBAND EASILY 10 Mbps
6-8 KM WITHOUT REPEATERS
NOISE IMMUNE

COST: MORE THAN COAX

FUTURE: BROADBAND

MEDIA
LINE-OF-SIGHT

MICROWAVES

- $10^9 - 10^{10}$ Hz
- SECURITY PROBLEM
- FCC LICENSING

LASER

- $10^{14} - 10^{15}$ Hz
- VERY DIRECTIONAL
- SENSITIVE TO ENVIRONMENT
- FDA LICENSING

INFRARED

- $10^{11} - 10^{14}$ Hz
- HIGHLY DIRECTIONAL
- EXCELLENT SECURITY
- NO LICENSING

COMMUNICATIONS PROTOCOLS

NEEDED TO

- ACCESS CHANNEL
- SHARE CHANNEL
- ENFORCE FAIRNESS
- ENABLE PRIORITY
- FORMAT MESSAGES
- ACKNOWLEDGE RECEIPT
- DETECT/CORRECT ERRORS
- ACHIEVE STANDARDS

COMMUNICATIONS PROTOCOLS

ISSUES

- PHYSICAL COMPATIBILITY
- PACKETIZING
- SEQUENCING
- FLOW CONTROL
- CONGESTION CONTROL
- ROUTING
- ACCESS MECHANISM
- SHARING
- PERFORMANCE OPTIMIZATION
- STANDARDS

PROTOCOLS

Fully scheduled

- Time division multiple access
- Polling

Contention

- Radio channels (Aloha)
- CSMA/CD (Ethernet)

Token passing

- Token passing bus
- Token passing ring

FULLY SCHEDULED

Time division multiple access uses fixed scheduling

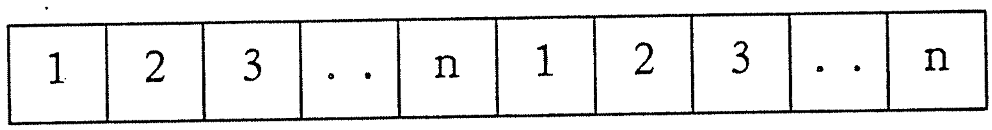

Master/slave (polling) can do adaptive scheduling

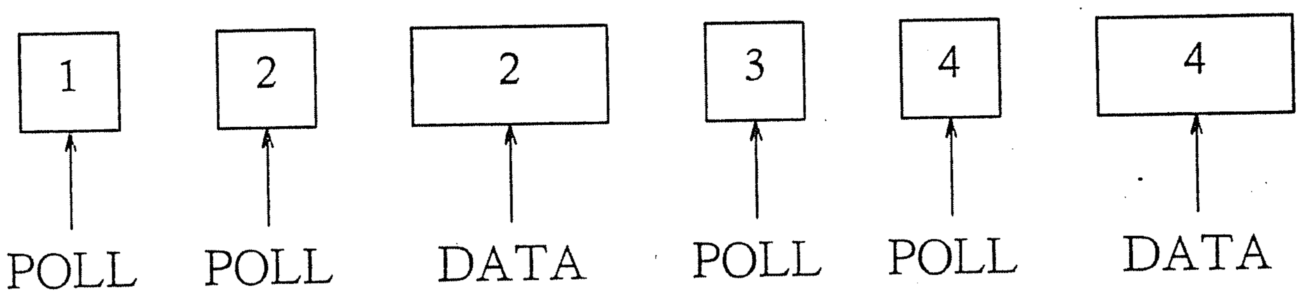

CONTENTION

"Schedules" on a packet-by-packet basis

Each node arbitrates for access independently

No guarantee of fairness

No concept of priority

Restricted to baseband implementation

Performance limited by bus length

Message delay near zero at low offered load

Message delay is exponential with offered load

Message delay highly variable

Only 20-40% of bus capacity effectively utilized

TOKEN PASSING

Token confers the right to transmit

Token holder uses bus for a (bounded) period of time, rather than for a single packet

Fairness can be assured or ignored

Priorities easily accommodated

Performance affected by bus length

Message delay is exponential with offered load

Message delay is non-zero even at low offered loads

Message delay increases less rapidly than CSMA/CD

Perhaps 60-80% of bus capacity utilized

LAN SUITABILITY

Star

- Centralized control
- Deterministic traffic patterns
- Circuit-oriented communication
- Continuous (or at least frequent) transmission
- Popular with military

Ring

- Need to know your physical neighbor
- Wiring distance should be short
- Popular with vendors of engineering workstations

LAN SUITABILITY

Contention Bus

- Low average offered loads
- Traffic pattern unpredictable and bursty
- No need for fairness or priority
- Intended for office automation

Token Bus

- Handles deterministic and non-deterministic loads
- Can accommodate real-time demands
- Implements priority
- Intended for general purpose applications

SECTION 3

ISO AND IEEE STANDARDS

- The ISO seven layer model
- Physical Layer
- Data link layer
- Network layer
- IEEE committee 802

ISO MODEL FOR OPEN SYSTEMS INTERCONNECTION

- LAYERED ARCHITECTURE
- SEVEN LAYERS
- PEER PROTOCOLS
- DETAIL HIDING
- ONLY HOPE FOR COMPATIBILITY
- X.25 IS LOWEST THREE LEVELS

ISO OSI MODEL

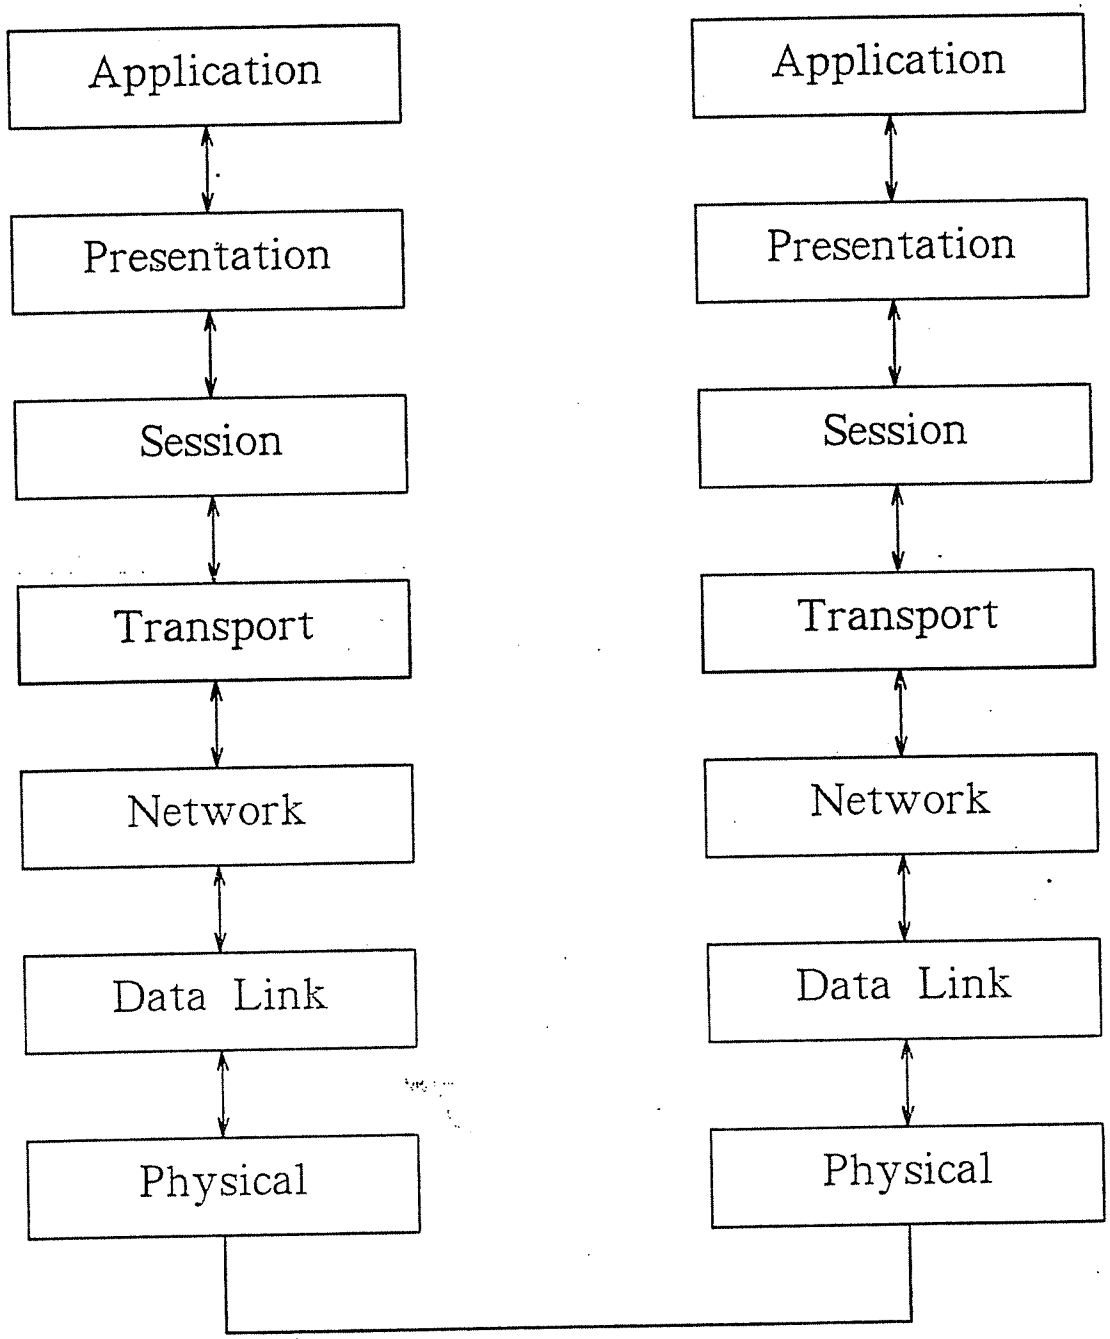

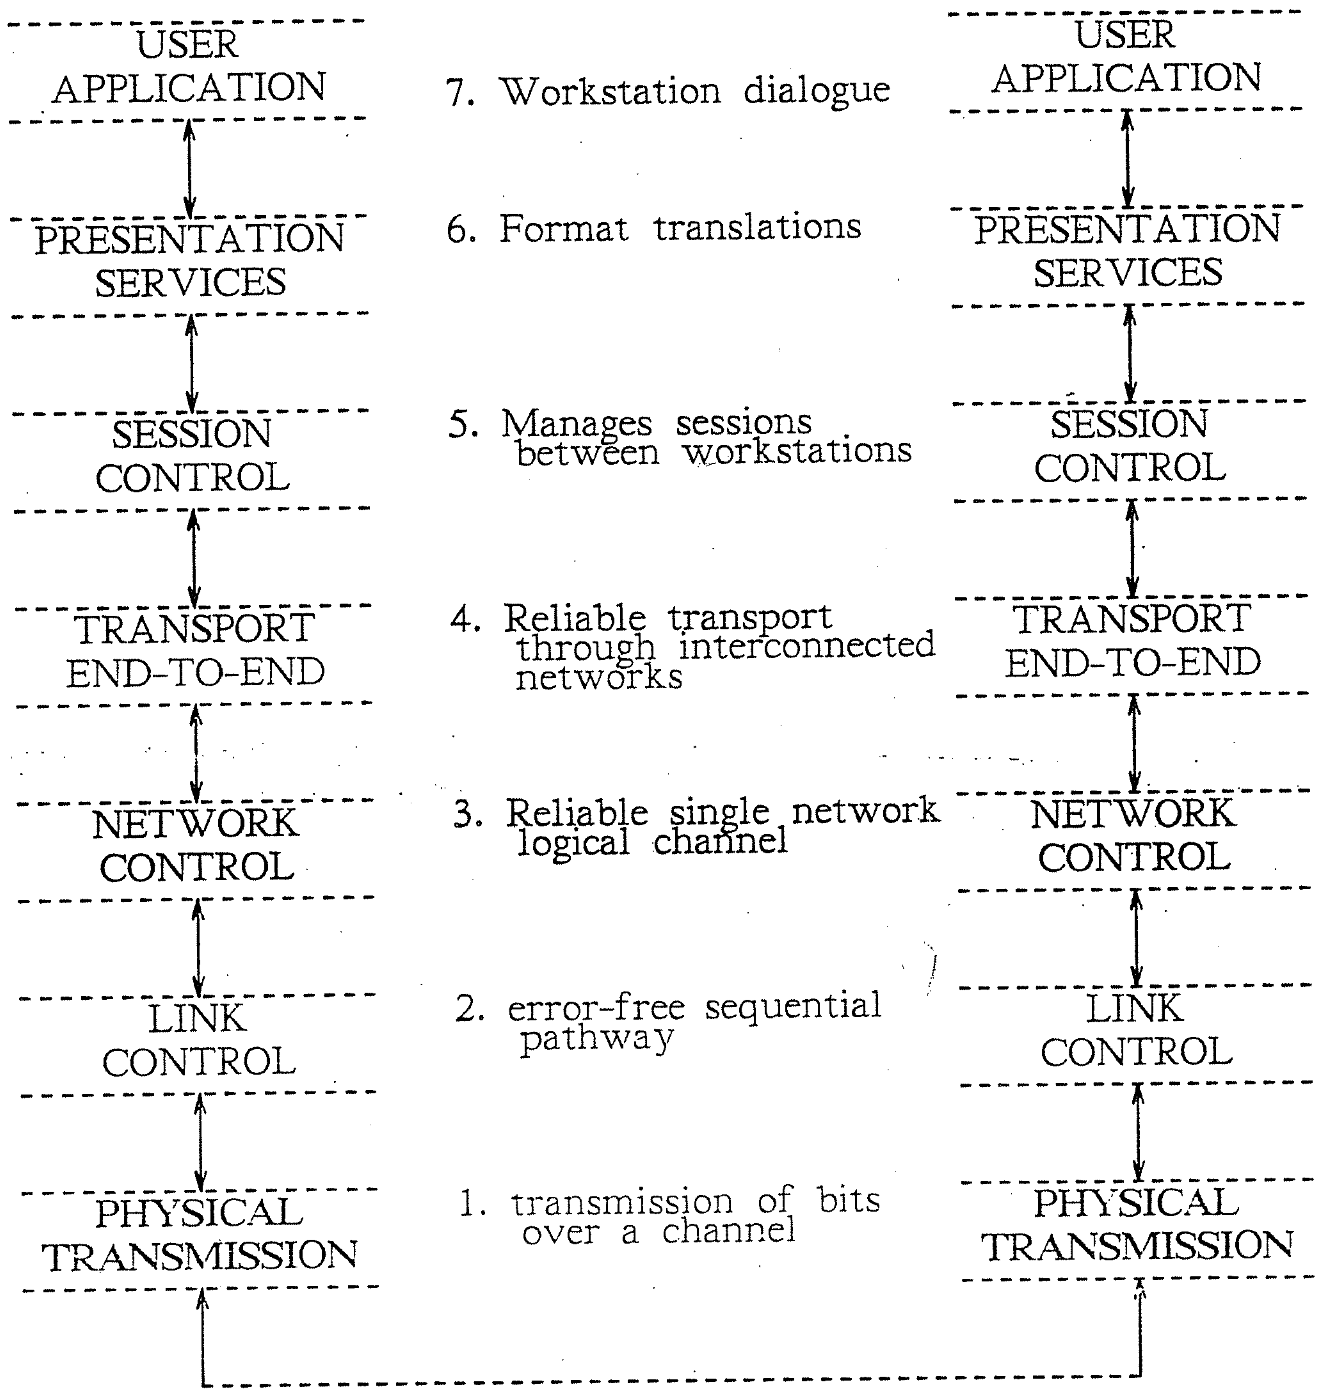
USER APPLICATION
7. Workstation dialogue
USER APPLICATION
PRESENTATION SERVICES
6. Format translations
PRESENTATION SERVICES
SESSION CONTROL
5. Manages sessions between workstations
SESSION CONTROL
TRANSPORT END-TO-END
4. Reliable transport through interconnected networks
TRANSPORT END-TO-END
NETWORK CONTROL
3. Reliable single network logical channel
NETWORK CONTROL
LINK CONTROL
2. error-free sequential pathway
LINK CONTROL
PHYSICAL TRANSMISSION
1. transmission of bits over a channel
PHYSICAL TRANSMISSION

NETWORK INTERFACE

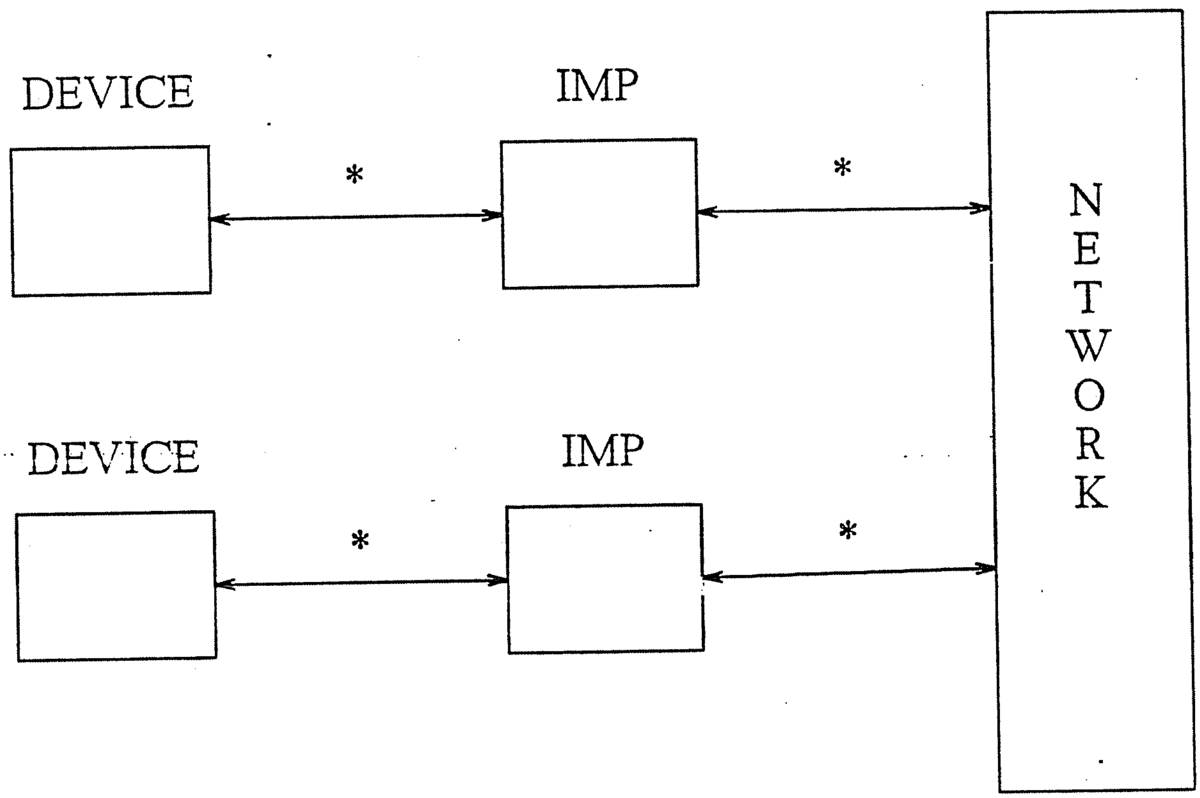

* ANALOG OR DIGITAL

PHYSICAL LAYER

- CHOOSE ANALOG OR DIGITAL NETWORK
- CHOOSE VIRTUAL CIRCUITS OR DATAGRAMS

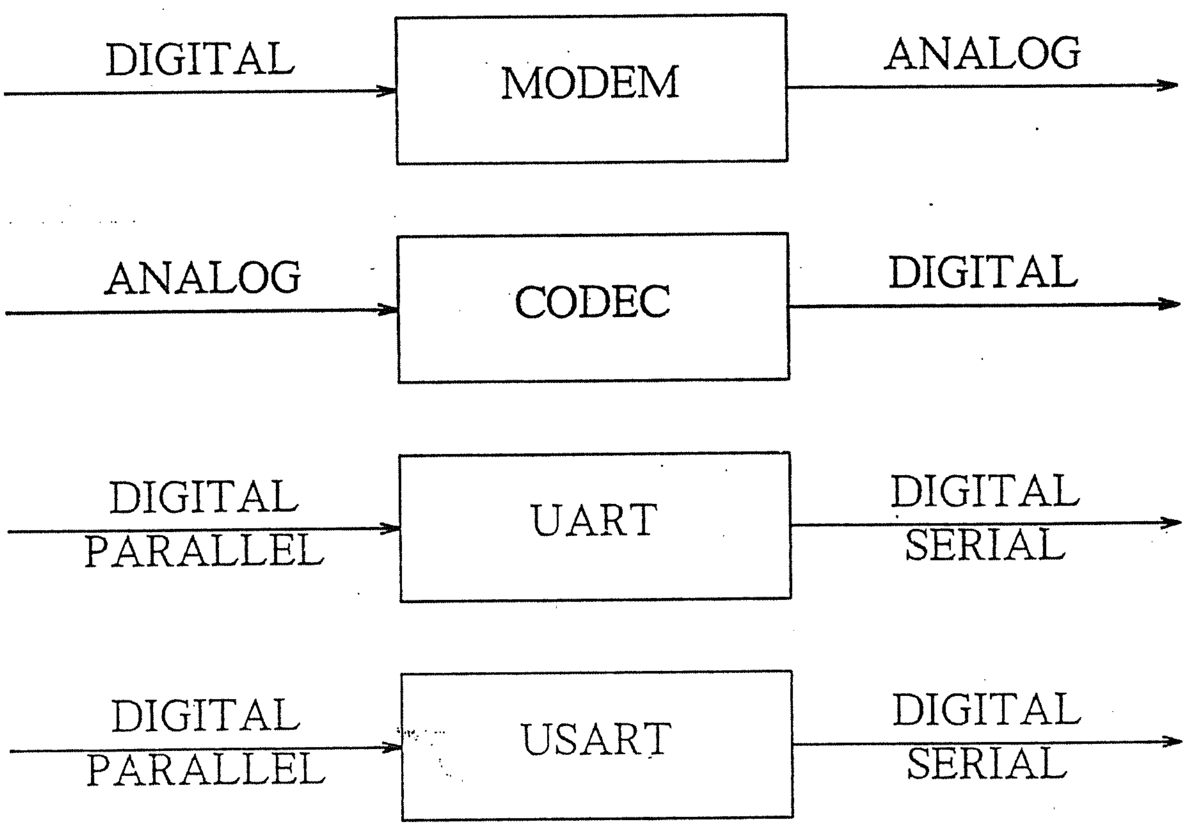

TDMA

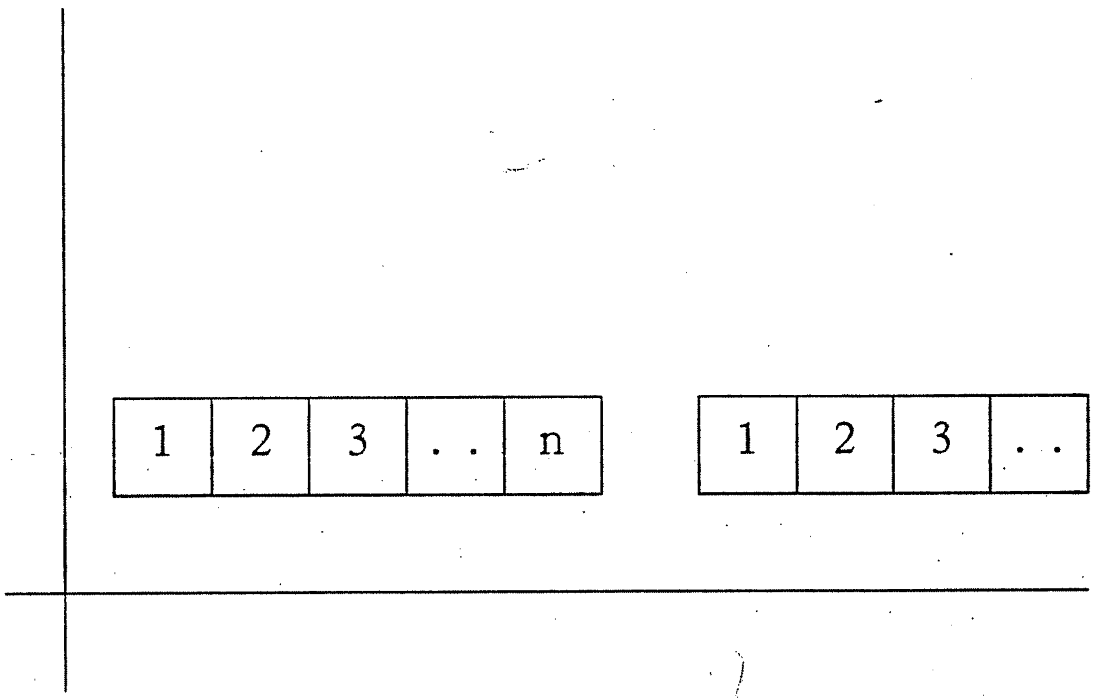

TIME DIVISION

MULTIPLE ACCESS

FDMA

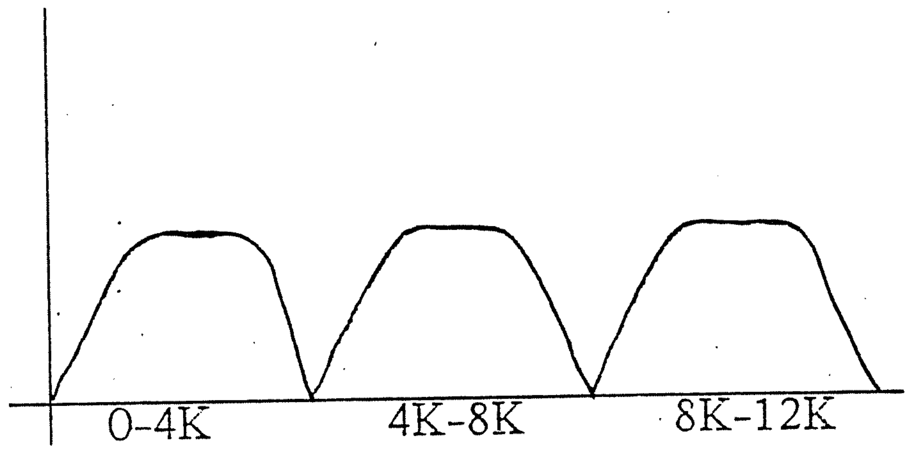

FREQUENCY DIVISION
MULTIPLE ACCESS

POLLING

BUS CONTROLLER

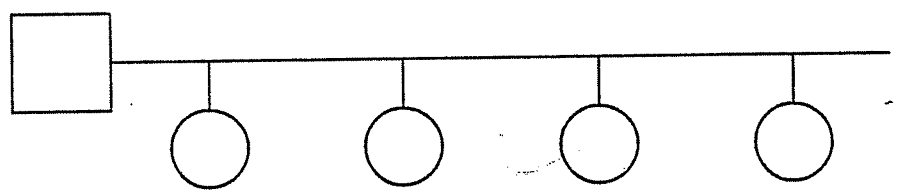

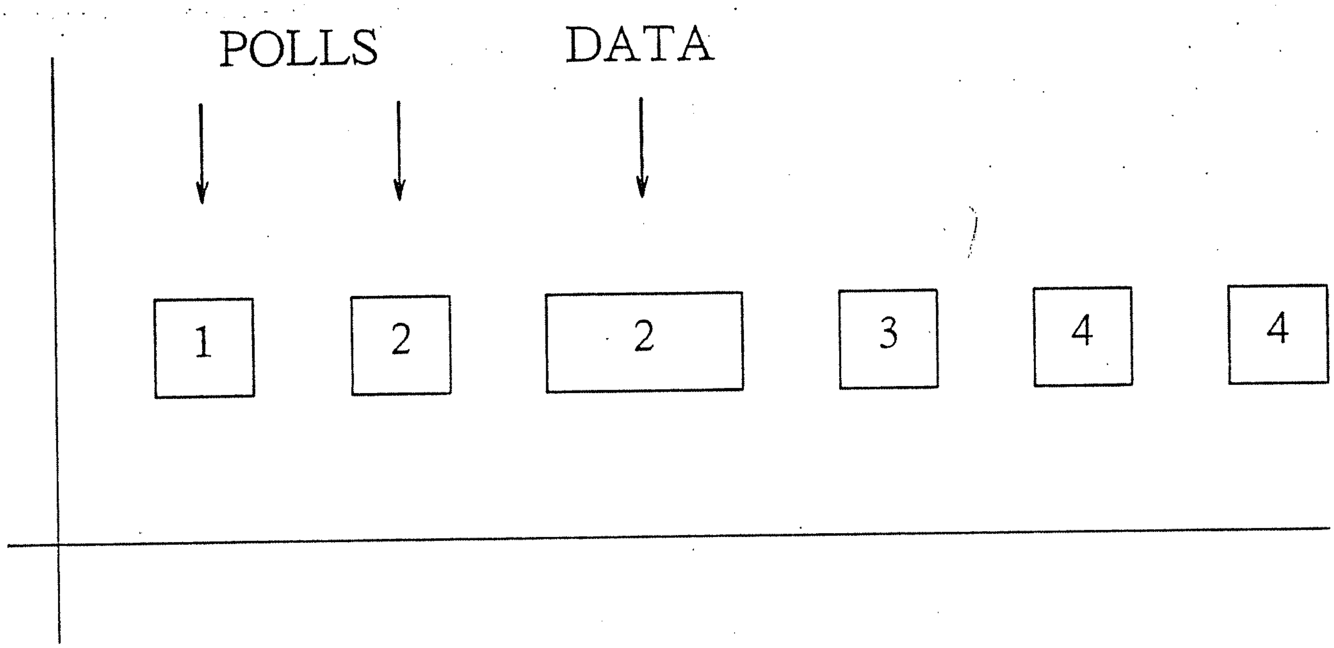

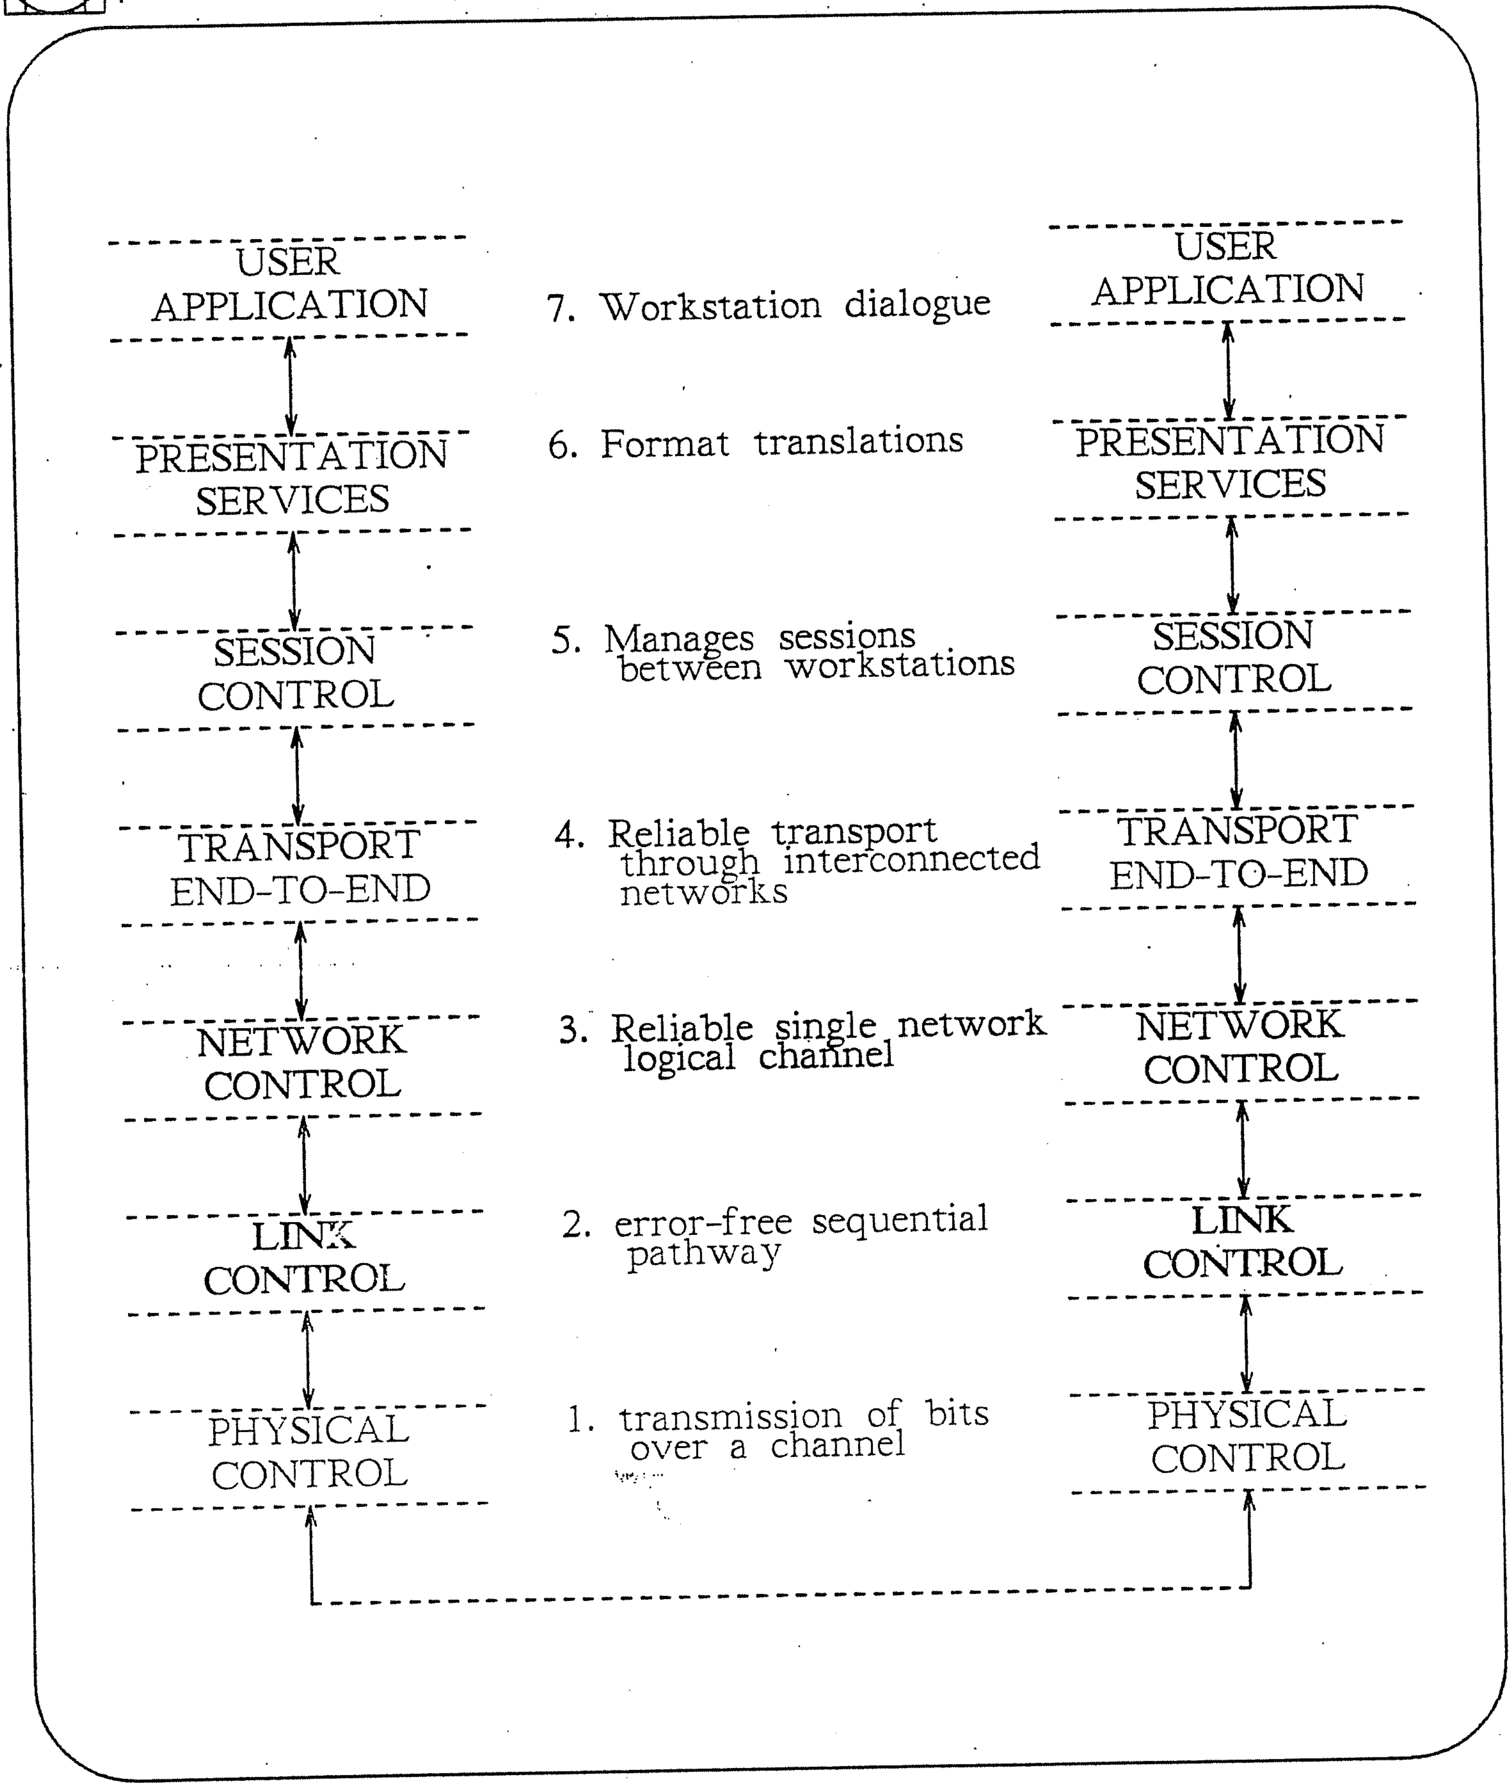
USER APPLICATION
7. Workstation dialogue
USER APPLICATION
PRESENTATION SERVICES
6. Format translations
PRESENTATION SERVICES
SESSION CONTROL
5. Manages sessions between workstations
SESSION CONTROL
TRANSPORT END-TO-END
4. Reliable transport through interconnected networks
TRANSPORT END-TO-END
NETWORK CONTROL
3. Reliable single network logical channel
NETWORK CONTROL
LINK CONTROL
2. error-free sequential pathway
LINK CONTROL
PHYSICAL CONTROL
1. transmission of bits over a channel
PHYSICAL CONTROL

ERRORS

1001110100111100
DATA BLOCK

PROBABILITY OF BIT IN ERROR IS "e"

PROBABILITY OF BLOCK CORRECT IS $(1-e)^N$

PROBABILITY OF BLOCK IN ERROR IS $1-(1-e)^N$

FOR $e \ll 1$, $(1-(1-e)^N) \approx e \cdot N$

NOT REALISTIC !

ERRORS OCCUR IN BURSTS

HAMMING CODES

MESSAGE	REDUNDANT
m	r

$n = m + r$ CODEWORD

2^m LEGAL MESSAGES

$< 2^n$ LEGAL CODEWORDS

HAMMING DISTANCE BETWEEN CODEWORDS IS NUMBER OF BIT POSITIONS WHERE THEY DIFFER

HAMMING DISTANCE OF CODE IS MINIMUM OVER ALL LEGAL CODEWORDS

DETECT d ERRORS NEED DISTANCE $d + 1$

CORRECT d ERRORS NEED DISTANCE $2d + 1$

CYCLIC REDUNDACY CODES

M(X) = ORIGINAL BINARY MESSAGE

G(X) = BINARY GENERATOR POLYNOMIAL, DEGREE r

T(X) = BINARY TRANSMITTED MESSAGE

METHOD:

1) APPEND r ZEROES TO M(X)
2) DIVIDE $M(X) \cdot 2^r$ BY G(X)
3) SAVE REMAINDER R(X)
4) $T(X) = 2^r \cdot M(X) - R(X)$
5) G(X) NOW DIVIDES T(X) EVENLY
6) TRANSMIT T(X)
7) RECEIVER VERIFIES

$$\frac{T(X)}{G(X)} = 0$$

RECOMMENDATIONS

CRC-16: $x^{16} + x^{15} + x^2 + 1$
CRC-CCITT: $x^{16} + x^{12} + x^5 + 1$

This CRC scheme detects:

- all single bit errors
- all double bit errors
- all odd-number-of-bit errors
- all busrts of length $\leqslant 16$
- $>99.99\%$ of all bursts of length >16
- does not detect errors involving a factor of $G(x)$, but these are extremely rare

RETRANSMISSION

How best to handle errors on networks?

Could use error correcting code (like Hamming), but at great cost in terms of number of bits needed per packet

Error correcting codes are necessary for applications involving long propagation delays, like satellite links and spacecraft control

For LANs, we take advantage of the short propagation time and simply retransmit a failed message

Transmitter sets a timer when message is sent; if no acknowledgement received by timeout, message is resent

Retransmission implies acknowledgement and sequencing

In normal operation, the error rate on LANs is so low that automatic retransmission is not provided at the data link layer; instead it is a user-selectable option at layer four or five

BROADCAST ROUTING

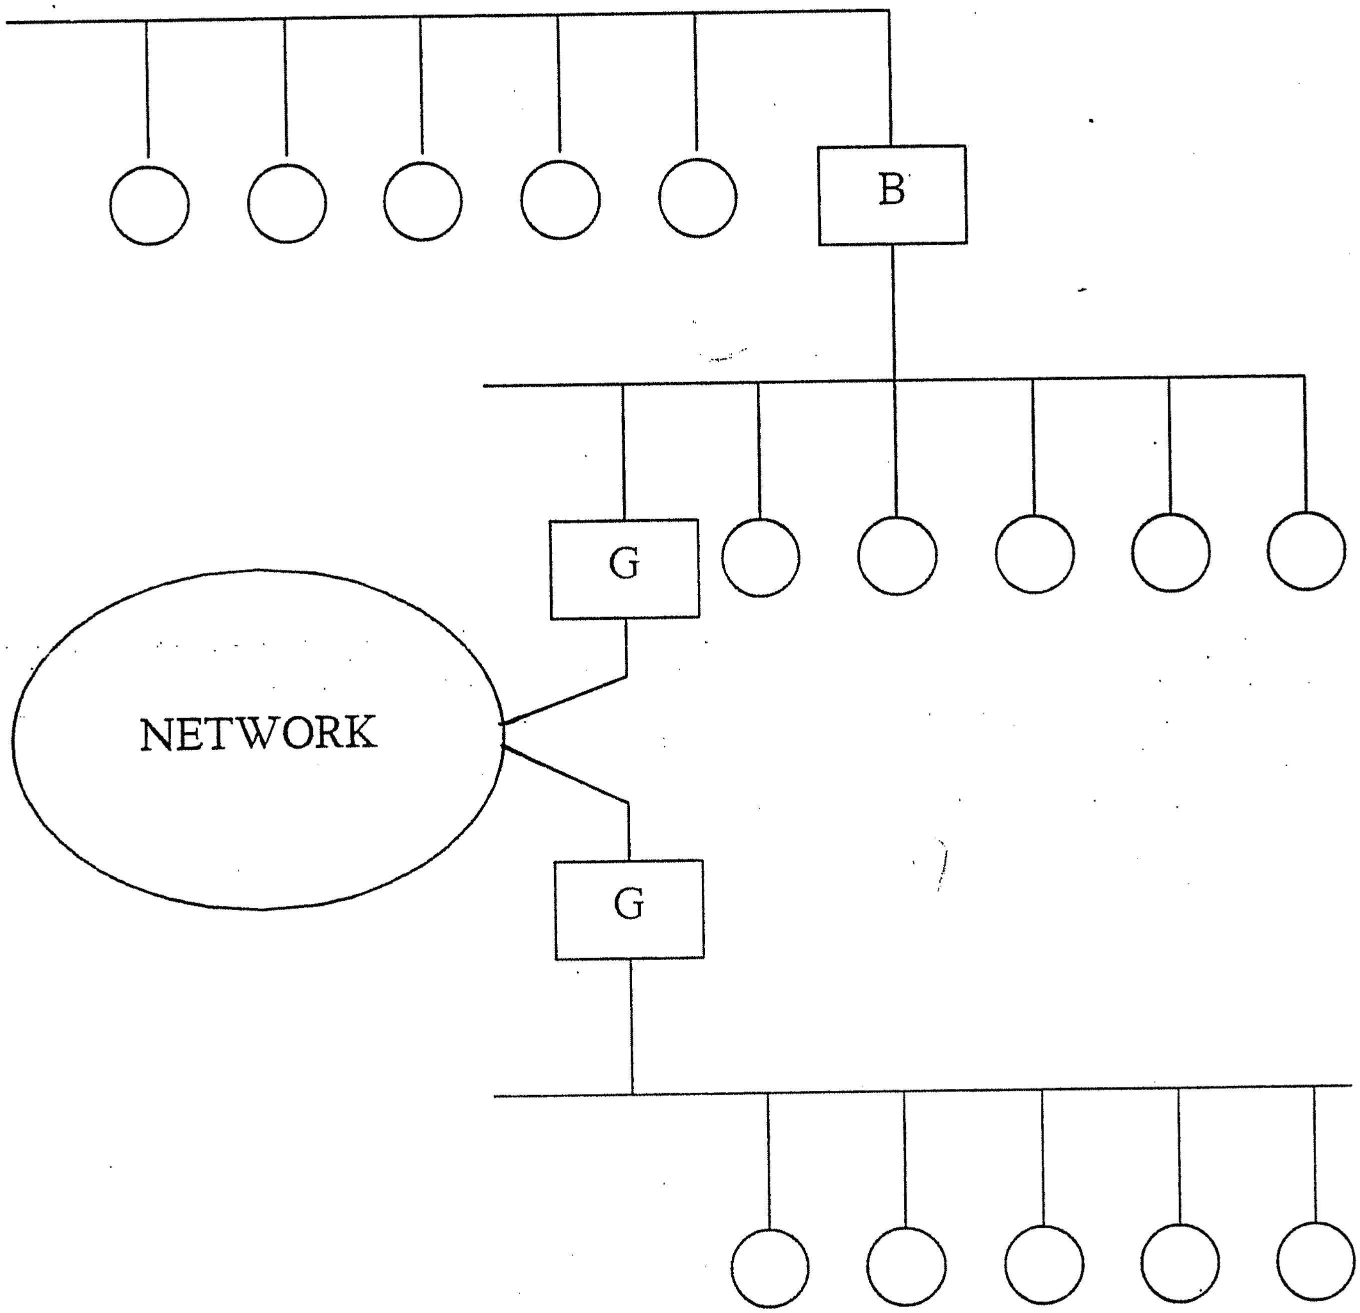

(slide broadcast)

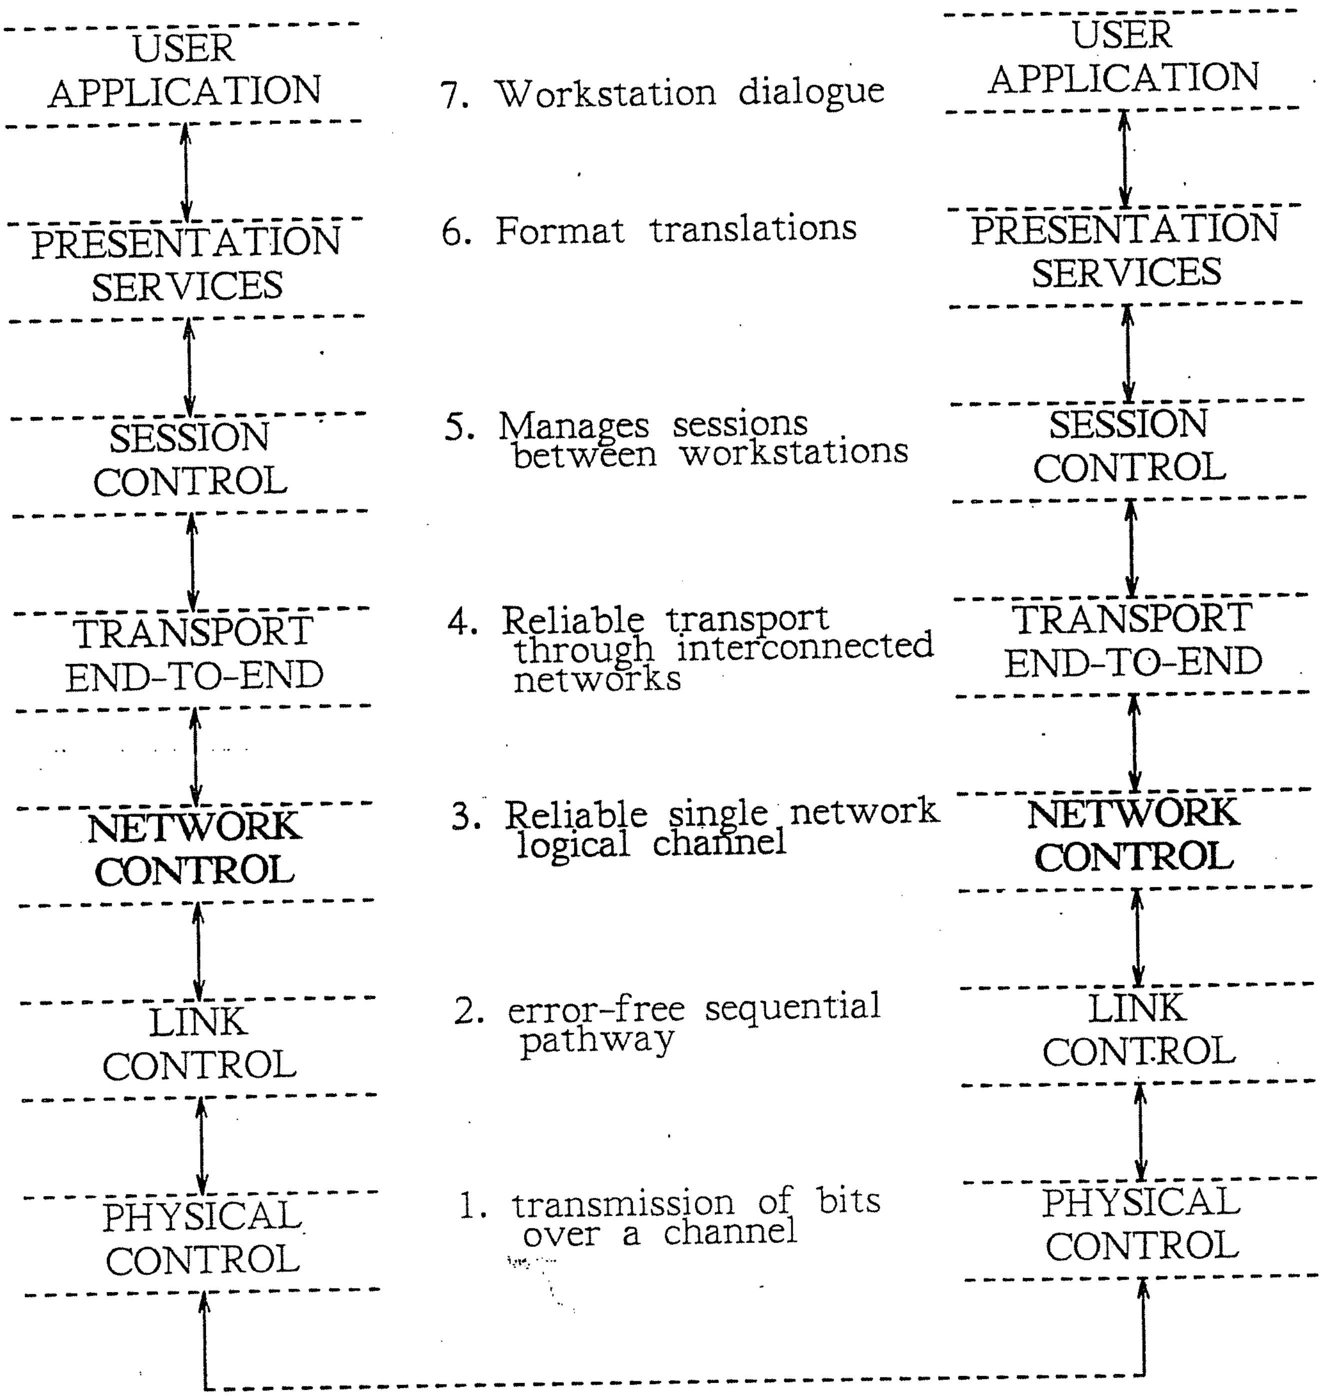
USER APPLICATION
7. Workstation dialogue
USER APPLICATION
PRESENTATION SERVICES
6. Format translations
PRESENTATION SERVICES
SESSION CONTROL
5. Manages sessions between workstations
SESSION CONTROL
TRANSPORT END-TO-END
4. Reliable transport through interconnected networks
TRANSPORT END-TO-END
NETWORK CONTROL
3. Reliable single network logical channel
NETWORK CONTROL
LINK CONTROL
2. error-free sequential pathway
LINK CONTROL
PHYSICAL CONTROL
1. transmission of bits over a channel
PHYSICAL CONTROL

IEEE COMMITTEE 802

Charged to develop a local area network standard

Committee membership predominately industrial

Recognized that application areas had different needs

- office automation
- factory automation
- real-time control systems

Impossible to agree on a single standard

- application areas are truly diverse
- mixed corporate membership inhibited agreement

802 SUBCOMMITTEES

802.1 -- Relationship to ISO OSI model

802.2 -- Logical Link Control of Data Link Layer

802.3 -- Contention Bus

802.4 -- Token Bus

802.5 -- Token Ring

802.6 -- Metropolitan Networks

LAYERS 1, 2, and 3

Layer 3 -- Network

- Performs routing in an arbitrary network
- Usually null in a broadcast protocol

Layer 2 -- Data Link

- Logical Link Control -- framing
- Medium Access Control -- controls transmission

Layer 1 -- Physical

- Wiring considerations
- Modulation and propagation

802.x RELATIONSHIPS

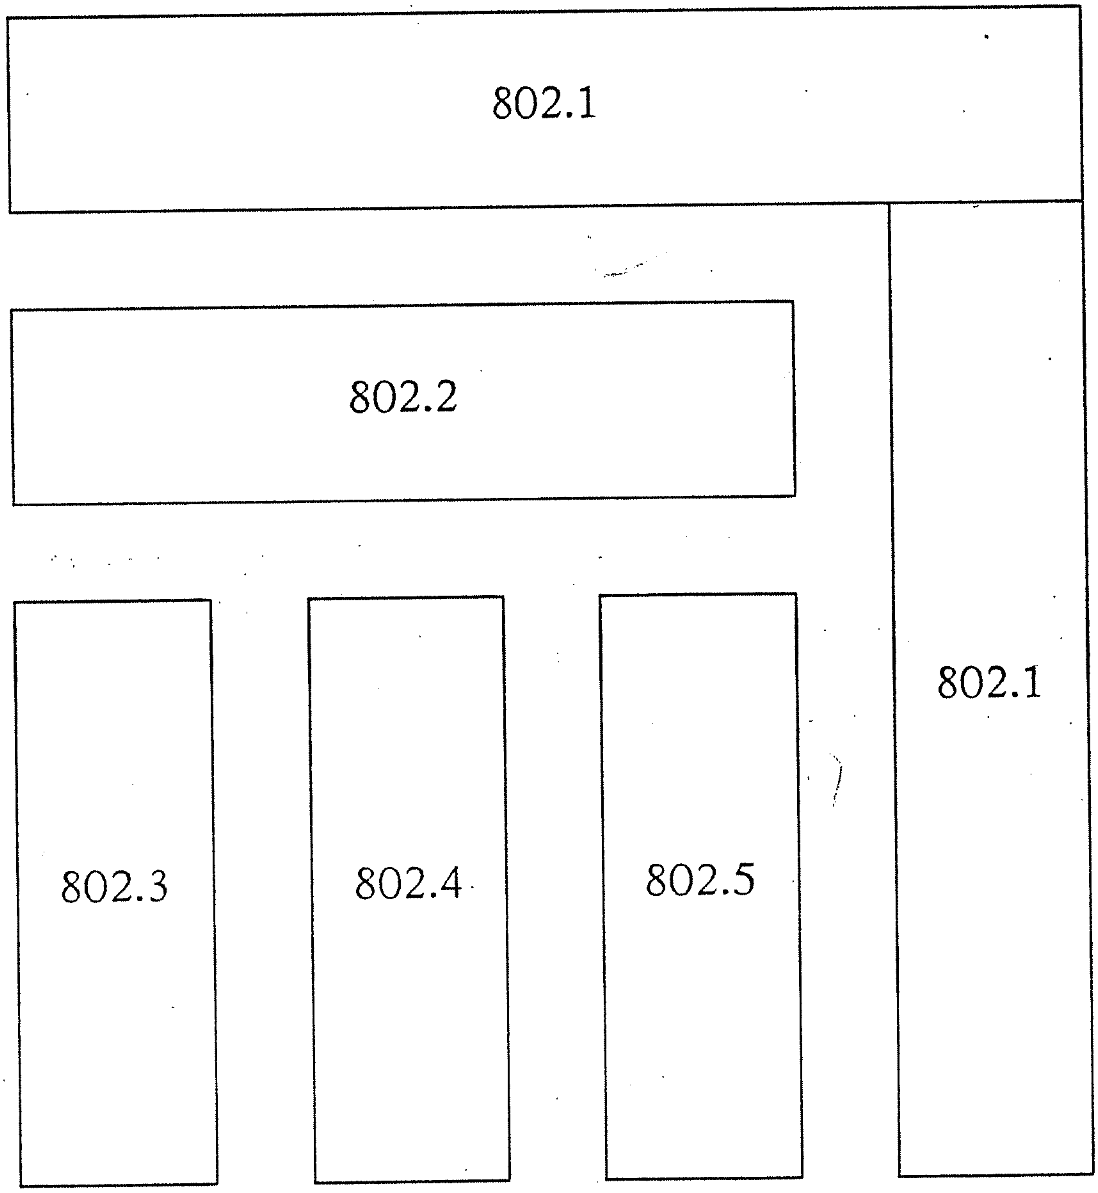

SECTION 4

IEEE 802.3 CONTENTION BUS

- Ethernet
- Physical and data link specifications
- Collisions and retransmissions
- Performance

ETHERNET GOALS

Simplicity

- Features which would complicate the design without substantially contributing to the meeting of other goals have been excluded.

Low Cost

- Since technological improvements will continue to reduce the overall costs of stations wishing to connect to Ethernet, the cost of the connection itself should be minimized

Compatibility

- All implementations of Ethernet should be capable of exchanging information at the data link level. For this reason the specification avoids optional features, to eliminate the possibility of incompatible variants of Ethernet.

ETHERNET GOALS

Addressing flexibility

- The addressing mechanisms should provide the capability to target frames to a single station, a group of stations, or to all stations on the network.

Fairness

- All stations should have equal access to the network when averaged over time.

Progress

- No single station operaing in accordance with the protocol should be able to prevent the progress of other stations.

High speed

- The network should operate efficiently at a data rate of 10 Mbps.

ETHERNET GOALS

Low delay

- At any given level of offered traffic, the network should introduce as little delay as possible in the transfer of a frame.

Stability

- The network should be stable under all load conditions, in the sense that the delivered traffic should be a monotonically non-decreasing function of the total offered traffic.

Maintainability

- The Ethernet design should allow for network maintenance, operation, and planning.

Layered architecture

- The Ethernet design should be specified in layered terms to separate the logical aspects of the data link protocol from the physical details of the communications medium.

ETHERNET NON-GOALS

Full duplex operation

- At any given instant, the Ethernet can transfer data from one source station to one or more destination stations. Bi-directional communication is provided by rapid exchange of frames, rather than full duplex operation.

Error control

- Error handling at the data link level is limited to the detection of bit errors in the physical channel, and the detection and recovery from collisions. Provision of a complete error control facility to handle detected errors is relegated to higher layers of the network architecture.

Security

- The data link protocol does not employ encryption or other mechanisms to provide security. Higher layers of the network architecture may provide such facilities as appropriate.

ETHERNET NON-GOALS

Speed flexibility

- This specification defines a physical channel operating at a single fixed data rate of 10 Mbps.

Priority

- The data link protocol provides no support of priority station operation.

Hostile user

- There is no attempt to protect the network from a malicious user at the data link layer.

ETHERNET

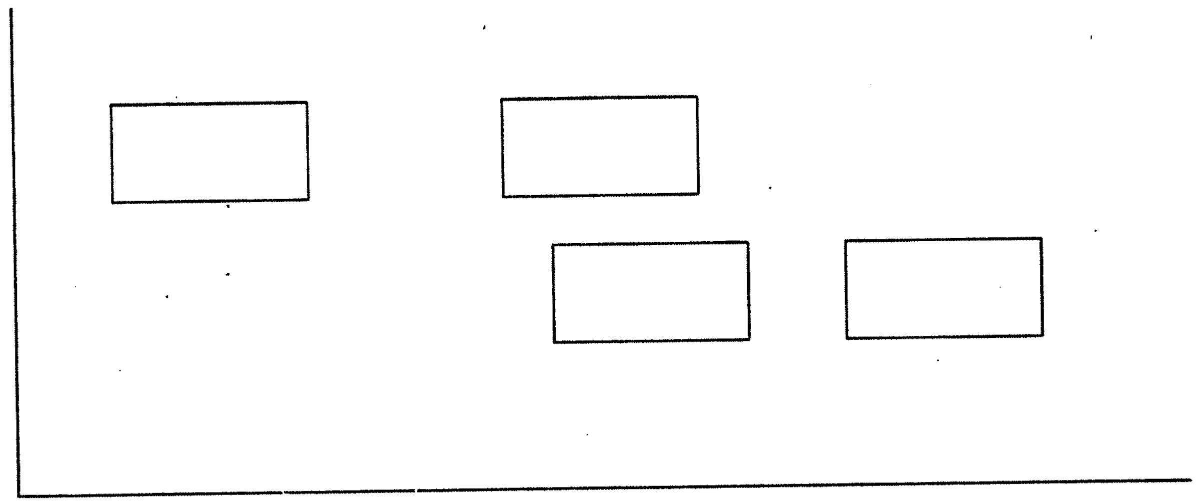

Form of CSMA/CD

Uses 50 ohm baseband coax cable

Maximum bus capacity of 10 Mbps

Maximum segment distance of 500 meters

Maximum end-to-end separation of 2,500 meters

Maximum of 1,000 stations

Uses coax taps

ETHERNET ARCHITECTURE

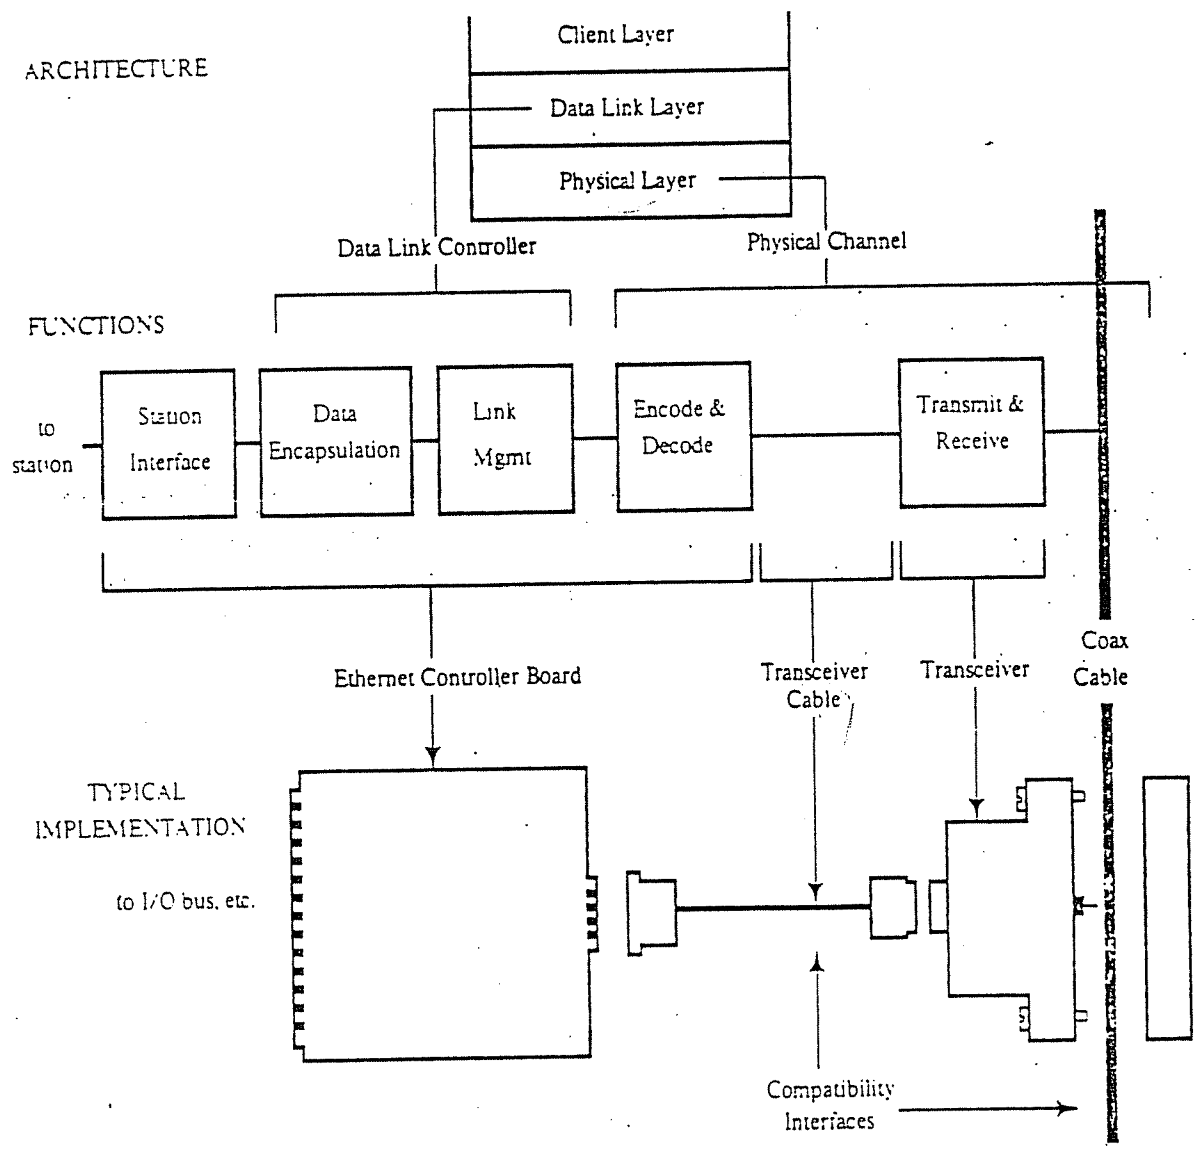

ETHERNET SEGMENTS

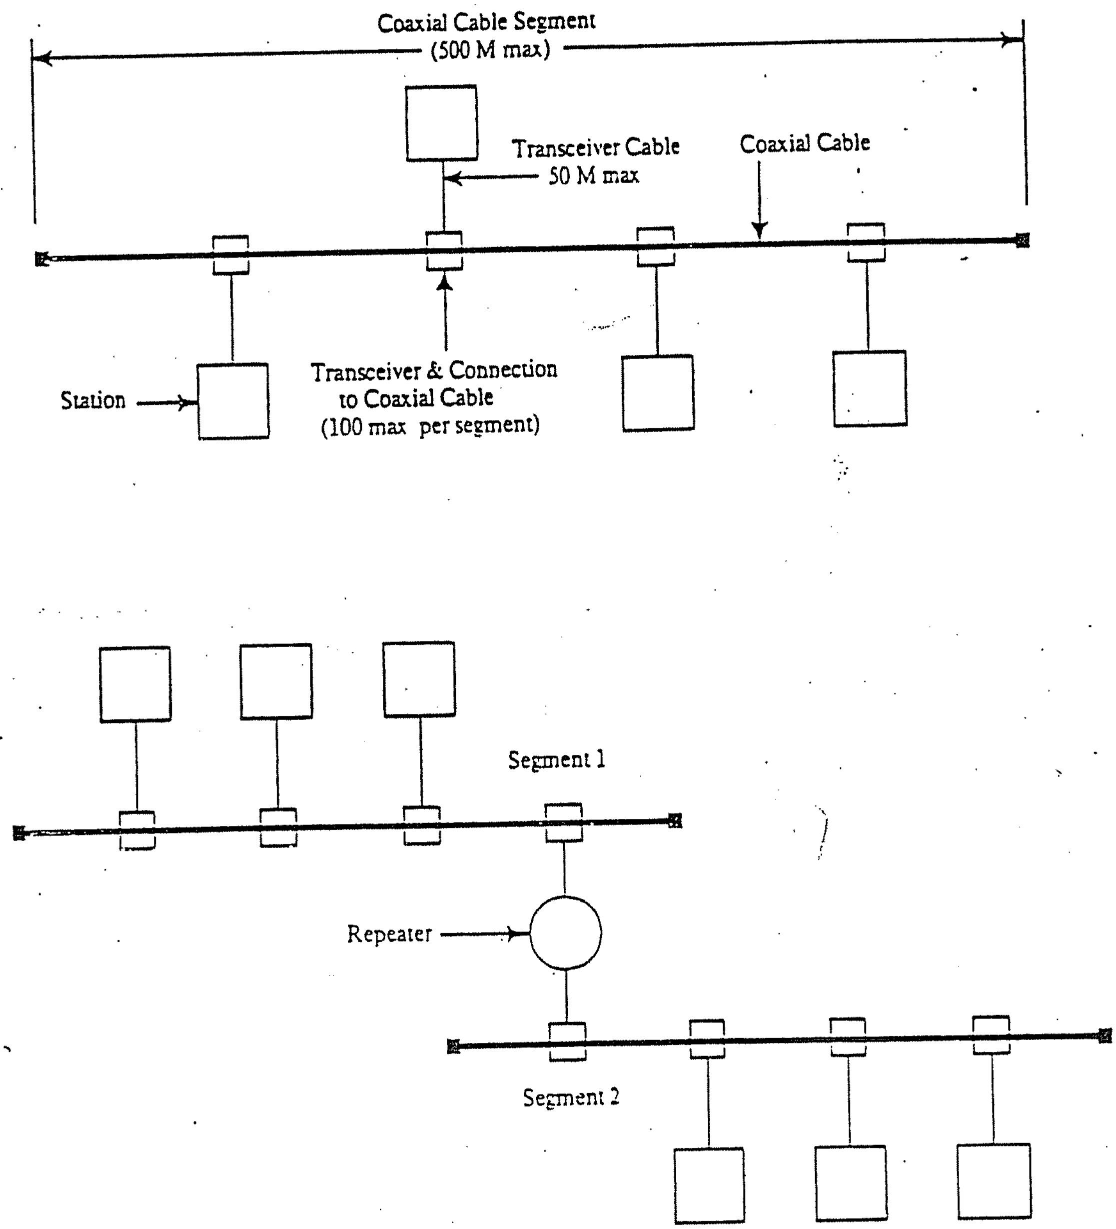

ETHERNET SEGMENTS

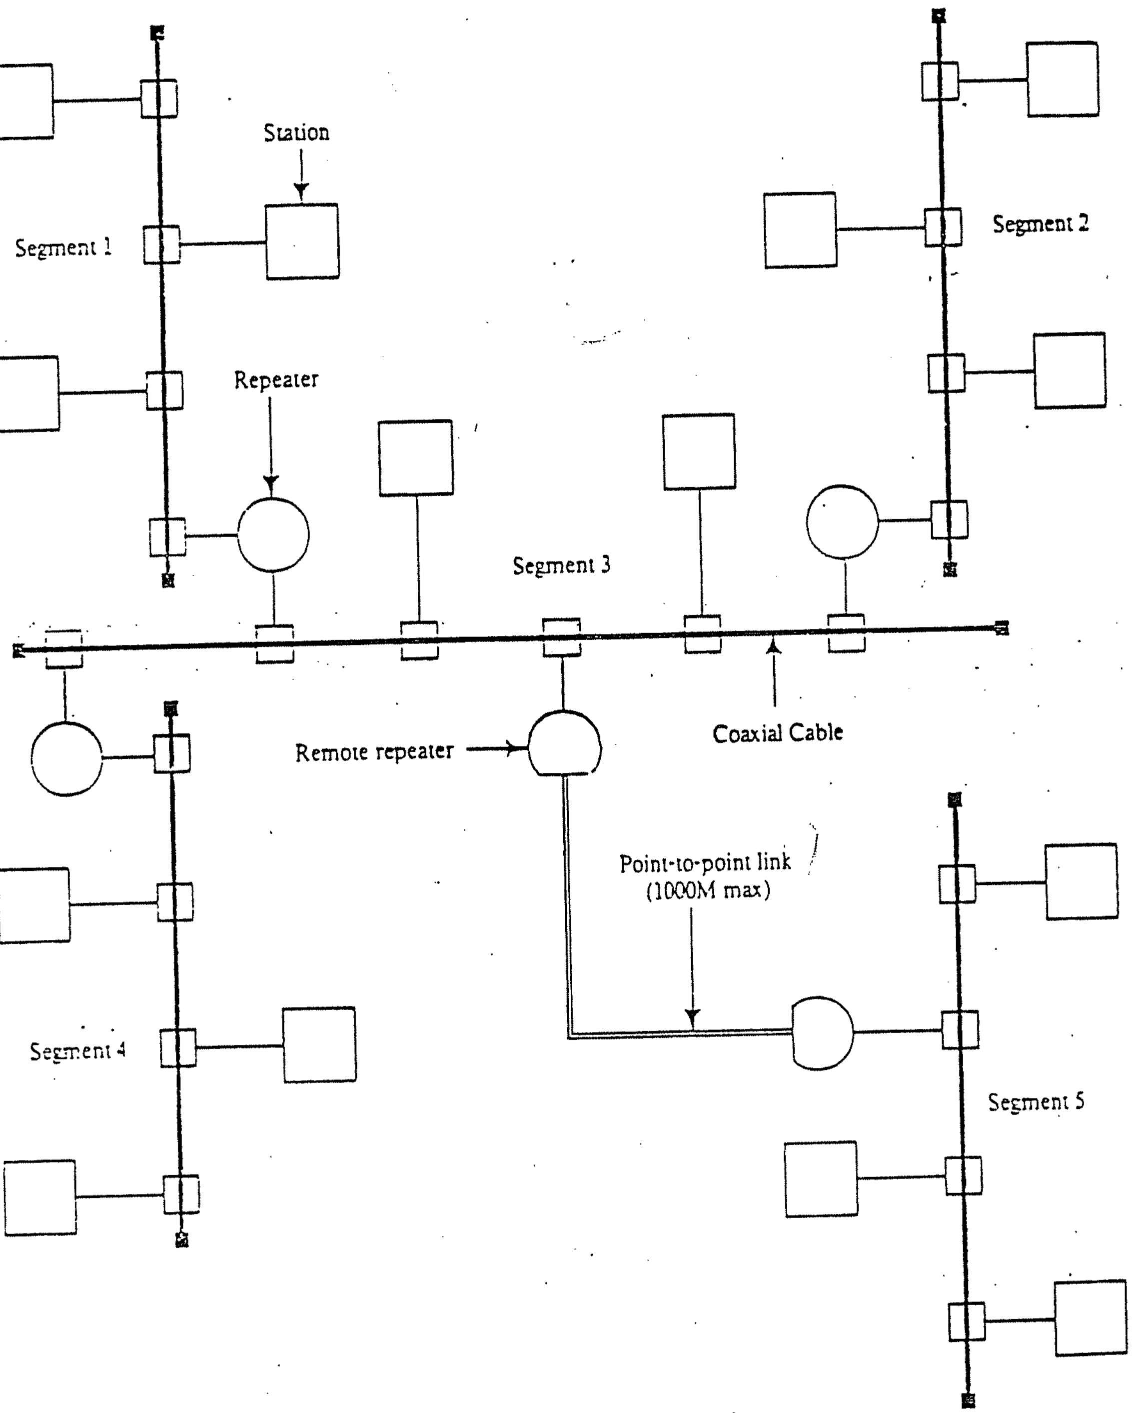

ETHERNET PACKET FORMAT

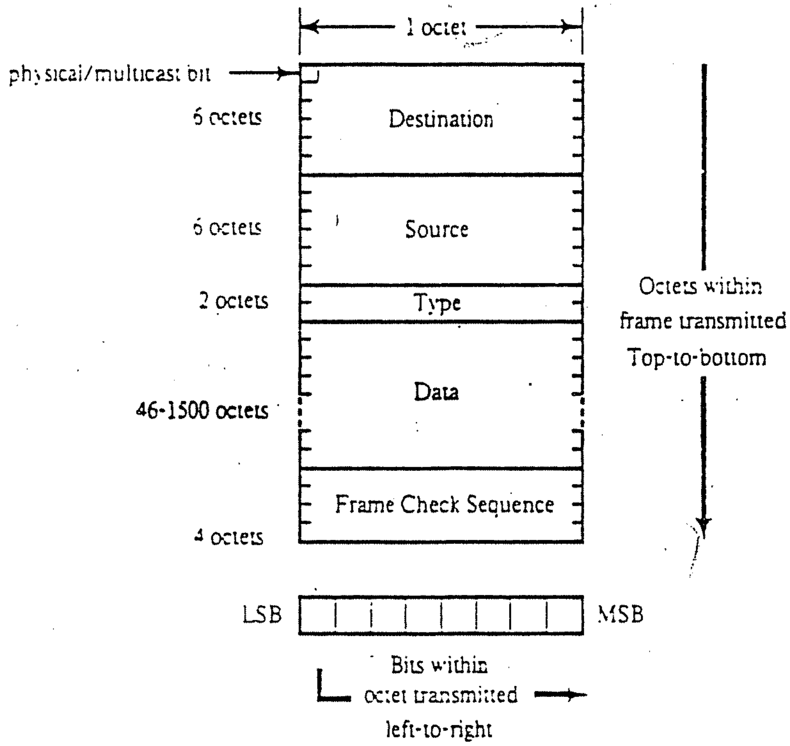

DATA LINK LAYER

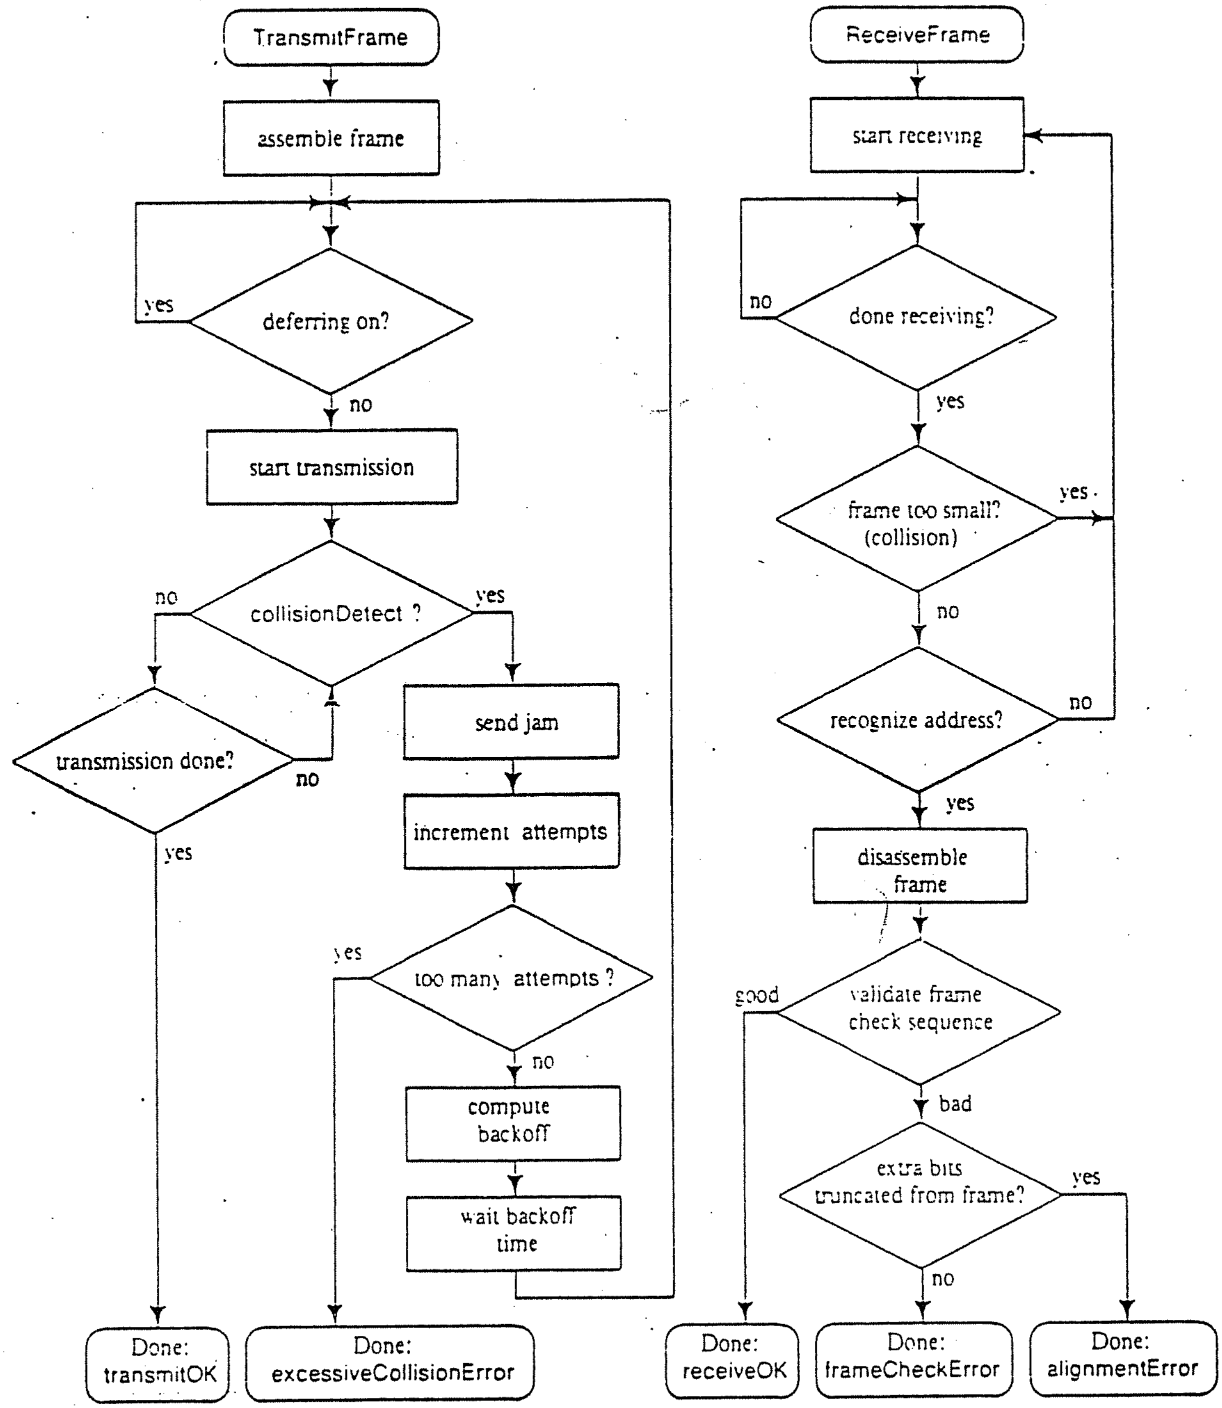

FrameTransmitter process
(Invoking Data Link TransmitFrame operation)

FrameReceiver process
(Invoking Data Link ReceiveFrame operation)

DATA LINK LAYER

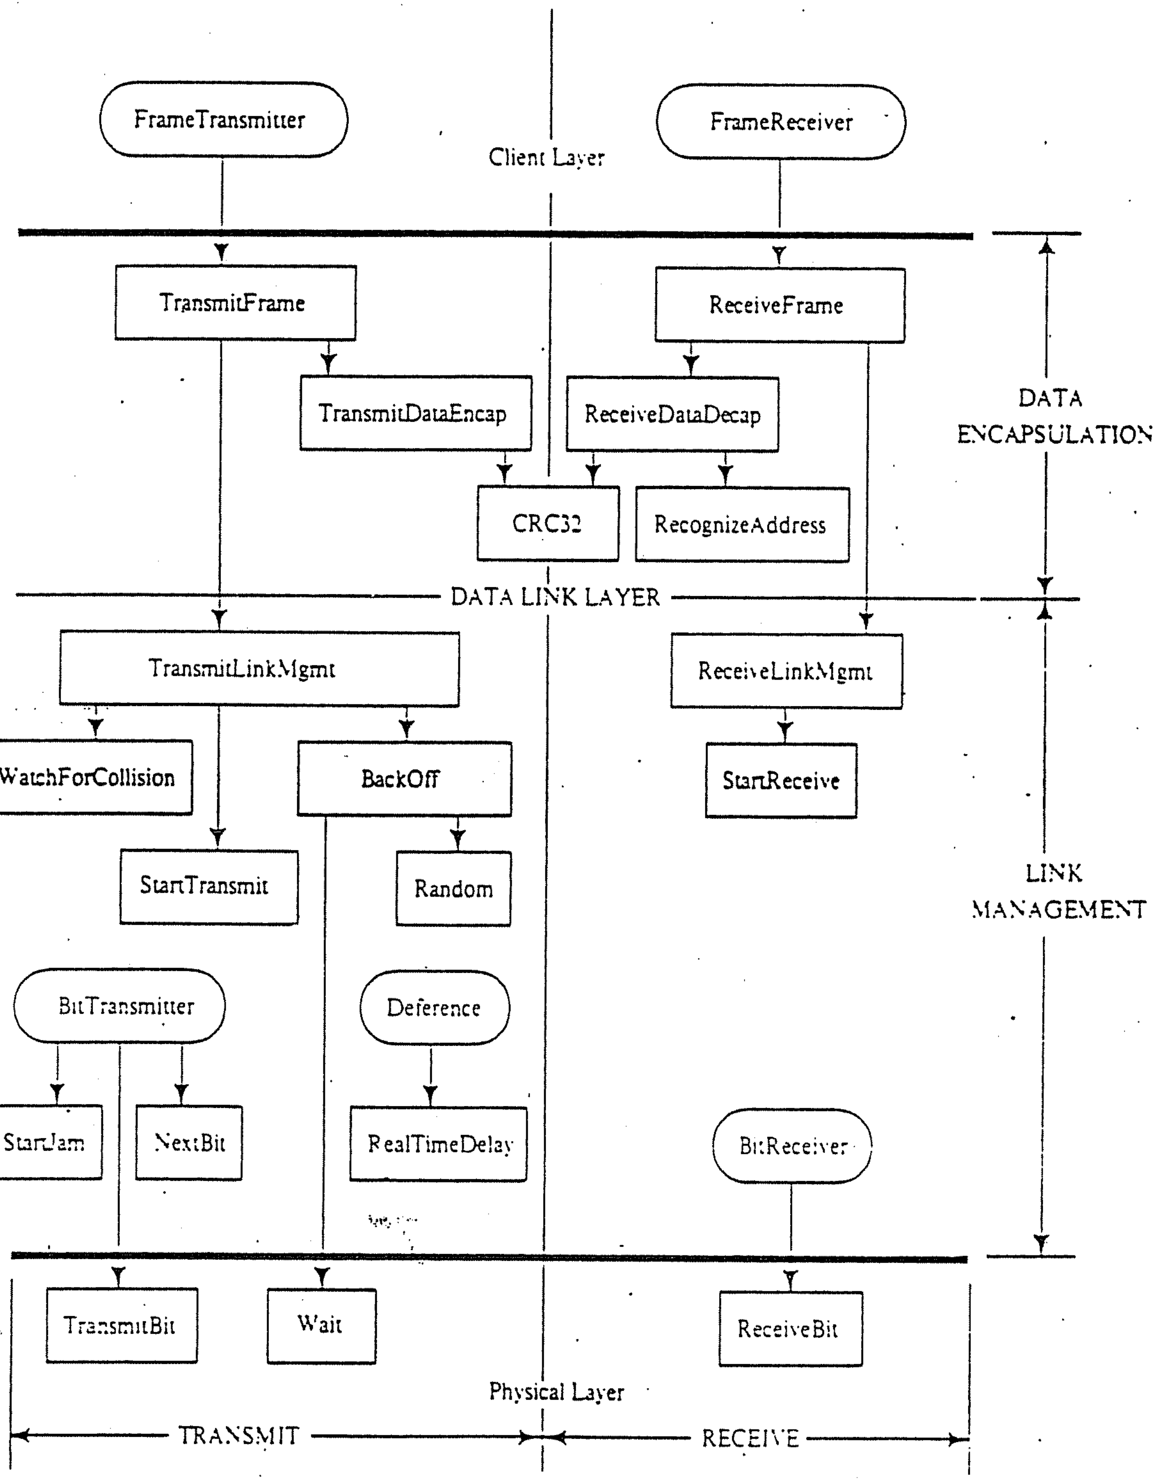

BACKOFF

Ethernet: CSMA/CD, 1-persistent, binary exponential backoff

At the i^{th} collision:

- j := min (i, 10)
- k := random (0 .. 2^j)
- wait k slot times = $512 \cdot k$ bit times
- sense channel
- transmit when idle

WORST CASE DELAY

Retry mechanism: transmit, collide, jam, backoff

Worst case delay

$$= \sum_{i=1}^{10} 2^i \cdot 512 + 5 \cdot 2^{10} + 15 \cdot 48$$

$$= 512(2^{11} - 2) + 5 \cdot 1024 \cdot 512 + 720$$

$$= 3{,}669{,}712 \quad \textit{bittimes}$$

$$\approx 0.37 \quad \textit{seconds}$$

ETHERNET THROUGHPUT

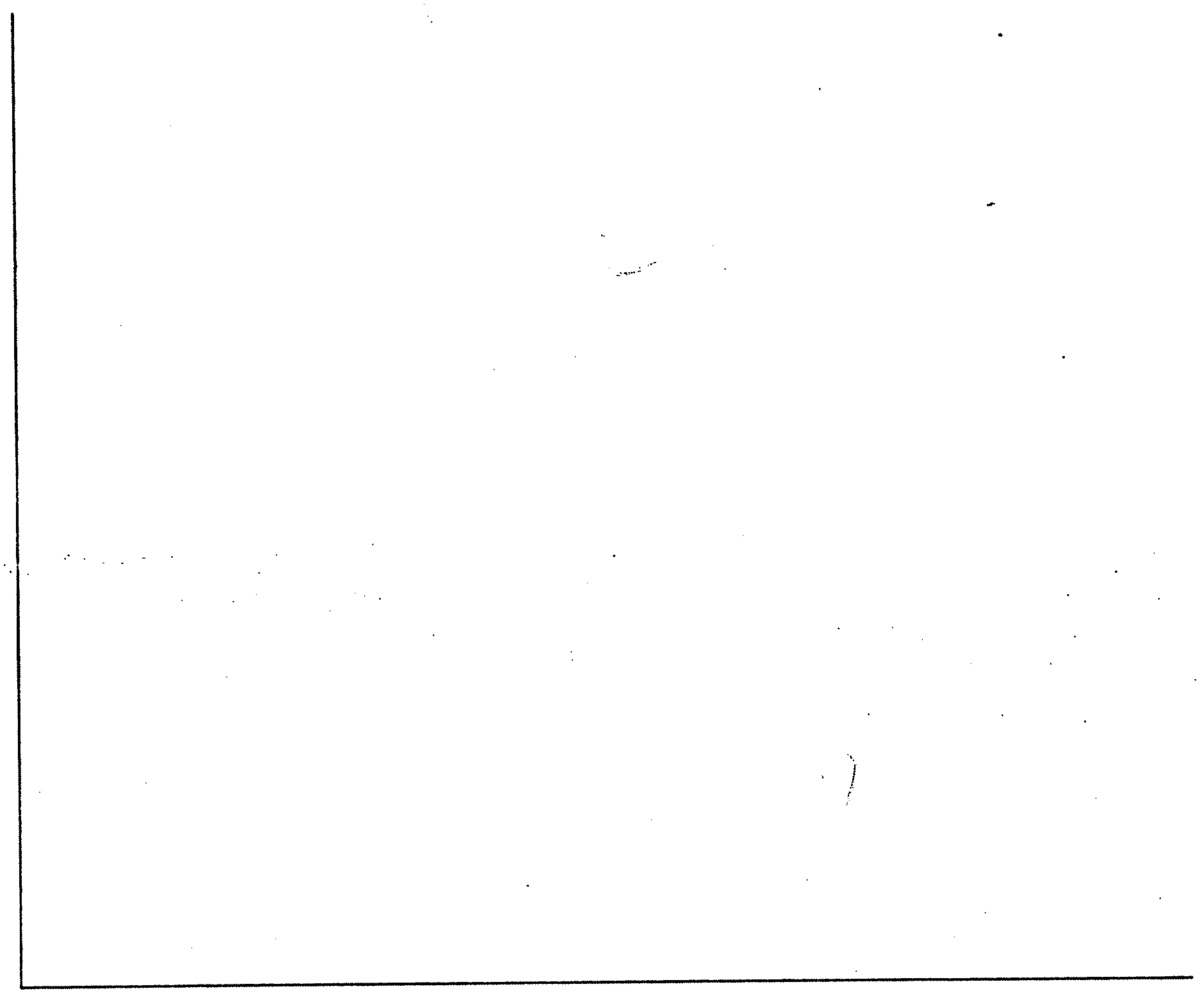

ETHERNET DELAY

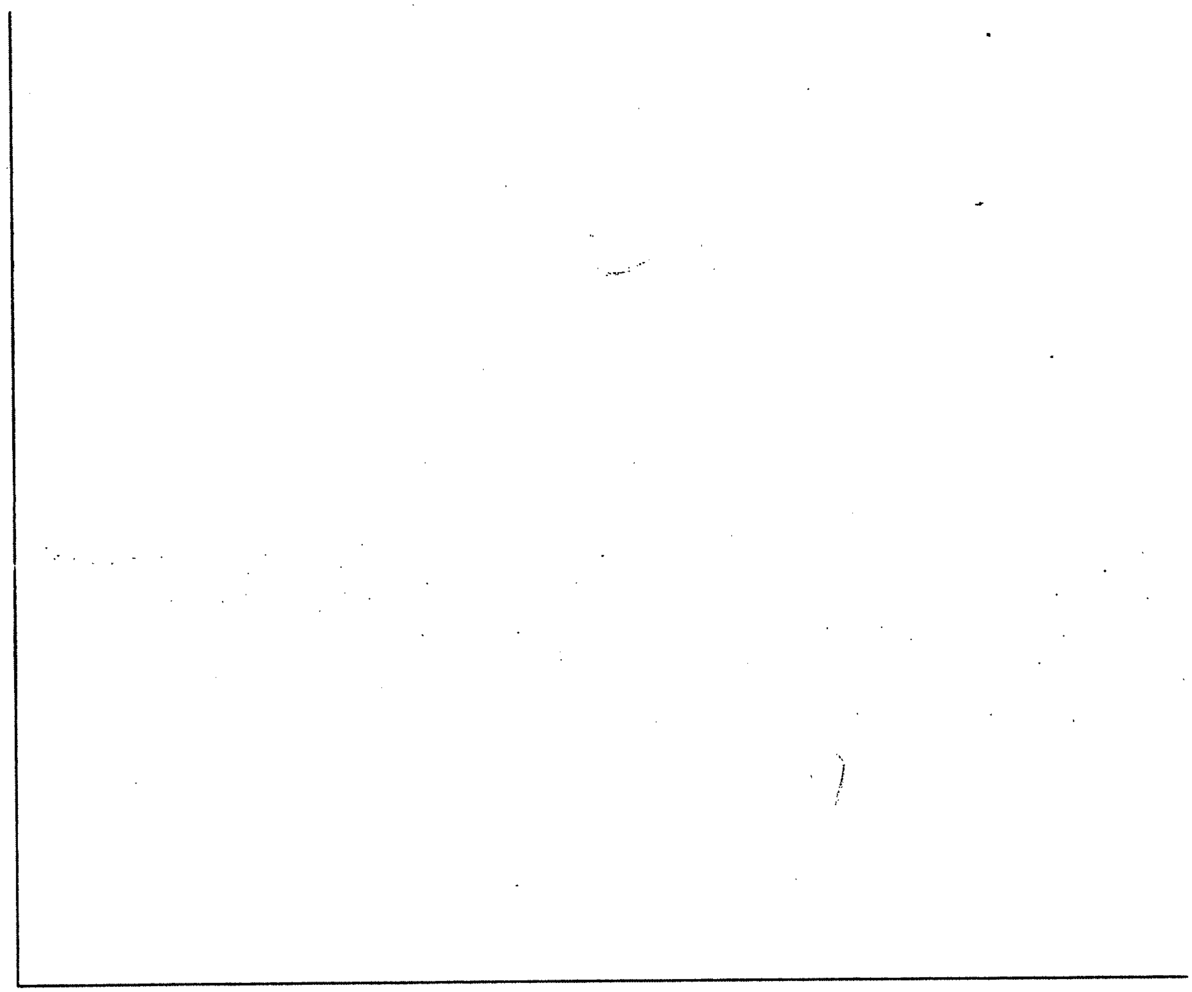

ETHERNET SENSITIVITY

THROUGHPUT sensitive to

- propagation delay
- packet size

DELAY sensitive to

- offered load
- throughput

ETHERNET SUMMARY

Designed for

- office automation environment
- bursty, asynchronous loads
- low average offered loads

Throughput

- is linear with offered load at low loads
- is stable and becomes asymptotic
- has maximum determined by propagation time and packet size

Delay

- is highly variable
- has large (0.37 sec) worst-case delay

ETHERNET SUMMARY

Accepted as IEEE and ISO standard

Supported jointly by Xerox, Digital Equipment, and Intel

Now has many manufacturers producing Ethernet-compatible equipment

Sure to be around for another 10 years

SECTION 5

IEEE 802.4 TOKEN BUS

- Bus architecture
- Physical and data link specifications
- Token passing
- Frame formats
- Logical ring membership
- Access classes (priorities)
- Performance

THE IEEE 802.4 TOKEN BUS

The 802.4 standard defines

- electrical and physical characteristics of the transmission medium
- electrical signaling method used
- frame formats transmitted
- actions of a station upon receipt of a data frame
- services provided by the Medium Access Control sublayer of the Data Link Layer

RELATIONSHIP OF 802.4 TO ISO OSI

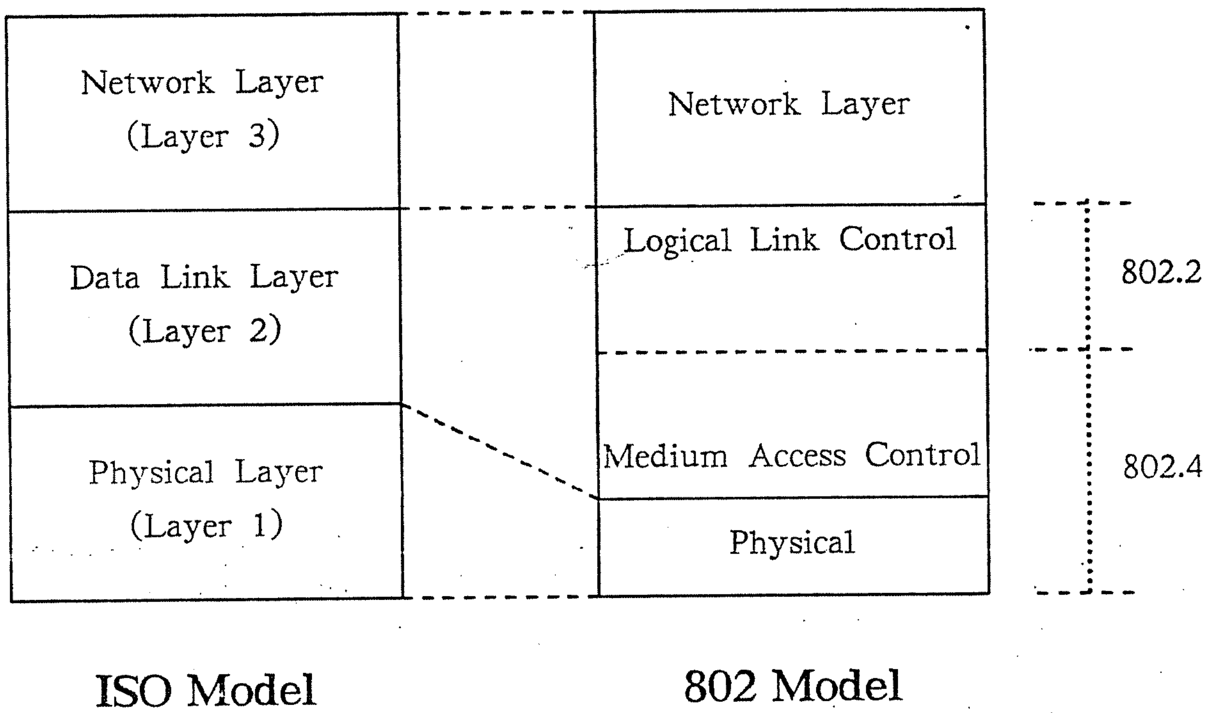

(slide ieee vs iso)

802.4 SUMMARY

A *token* controls access to the physical medium

Token holder is momentarily the network master

Implements four priorities, or *access_classes*

Station must pass the token to a known successor within a bounded time

Orderly progression of the token from station to station forms a logical ring on a physical bus

Station's interface is a Medium Access Controller (MAC)

MAC implements the protocol, including

- token recognition, passing, and regeneration after loss
- message encapsulation and framing
- service of the four priorities
- error control and recovery

TOKEN

Standard allows a network-wide choice of either 16 or 48 bit addresses

Token is an explicit message of at least 96 or 160 bits depending upon address size (token frame can carry data)

Token consists of:

- preamble (1 or more octets)
- start delimiter (1 octet)
- control information (1 octet)
- destination address (2 or 6 octets)
- source address (2 or 6 octets)
- optional data (0 or more octets)
- frame check sequence (4 octets)
- end delimiter (1 octet)

MESSAGE FRAME

Message frame contains

- preamble (1 or more octets)
- start delimiter (1 octet)
- control information (1 octet)
- destination address (2 or 6 octets)
- source address (2 or 6 octets)
- data (0 or more octets up to maximum frame size of 8191 octets)
- frame check sequence (4 octets)
- end delimiter (1 octet)

TOKEN PASSING

At startup, each station is assigned a unique logical address

During startup, stations add themselves to the logical ring one at a time, thereby learning their successor

A station may transmit only while it holds the token

When all data has been transmitted or certain timers expire, token must be passed to successor

A station will periodically query the network to determine whether additional stations wish to join the ring

Special cases to recover from loss of token, failure of successor, etc.

LOGICAL RING MEMBERSHIP

At startup, each station is assigned a unique logical address

The station's *inter_solicit_count* is initialized to zero

The *Max_Inter_Solicit_Count* is randomized in its low order two bits every 50 milliseconds or after each use

The *inter_solicit_count* is decremented once per token receipt, just prior to passing the token

When the *inter_solicit_count* reaches zero, and if the *ring_maintenance_timer* has not expired, a *response_window* is opened

A *solicit_successor* message is sent

The response, if any, is resolved

If no response is received, *inter_solicit_count* is reloaded with *Max_Inter_Solicit_Count*

RESOLVING SOLICIT SUCCESSOR

Only three possible responses: none, one, or many

- No response
  - — pass token
- One response
  - — responding station sends *Set_Successor* frame to token holder
  - — token holder changes his "next station" variable to be address of new station
  - — pass token

RESOLVING SOLICIT SUCCESSOR

- Many responses (collision)
  - — open four demand windows and send *Solicit_Successor*
  - — other stations respond in window 0,1,2,3 depending upon value of the one's complement of the first two bits of station address
  - — contending stations who hear a valid frame drop out of contention
  - — process continues, using successively lower ordered pairs of address bits, as long as there are multiple responses
  - — if no resolution on lowest pair of address bits, there is a duplicate address error
  - — try once more; competing stations generate random address
  - — stations which lose report duplicate address to network manager

LEAVING THE LOGICAL RING

Station has two variables, *in_ring_desired* and *any_send_pending*

Variable *any_send_pending* is true whenever there are any messages in any *access_class*

When ready to leave, sets *in_ring_desired* to false

When both flags are false, waits for token, then sends this station's predecessor a *Set_Successor* message with this station's successor as its value

OTHER ERRORS

Standard specifies recovery actions for

- lost or multiple tokens
- token pass failure
- deaf stations
- duplicate address
- stuck transmitter

ACCESS CLASSES

Four priorities or *access_classes*

- Synchronous
- Urgent Asynchronous
- Normal Asynchronous
- Time Available

Only *Synchronous* is guaranteed a level of service

Other classes receive "best effort"

An 802.4 station must implement either the *Synchronous* class alone or else all four classes simultaneously

Note that priority applies to a message, not a station

NETWORK PARAMETERS

Let the *Synchronous*, *Urgent Asynchronous*, *Normal Asynchronous*, and *Time Available access_classes* be abbreviated by *S*, *UA*, *NA*, and *TA*, respectively

High Priority Token Hold Time (HPTHT) -- the maximum amount of time a station may serve its *Synchronous* class

Urgent Asynchronous Target Rotation Time (TRT_{UA}) -- goal token rotation time for class UA

Normal Asynchronous Target Rotation Time (TRT_{NA}) -- goal token rotation time for class NA

Time Available Target Rotation Time (TRT_{TA}) -- goal token rotation time for class TA

STATION TIMERS

token_hold_timer (tht) -- when it expires, station may complete sending the packet in progress, but must then sequence to the next lower priority *access_class* or, if serving class TA, must pass the token to the station's successor

$token\_rotation\_timer_{UA}$ (trt_{UA})

$token\_rotation\_timer_{NA}$ (trt_{NA})

$token\_rotation\_timer_{TA}$ (trt_{TA})

- used to monitor the token cycle time at *access_classes* UA, NA, and TA

$token\_rotation\_timer_{RM}$ (trt_{RM})

- ring maintenance timer

SERVICE DISCIPLINE

At station startup, *trt*'s are initialized to zero

Token arrives at a station; its *tht* is set to the *HPTHT* and begins timing

Station serves its *Synchronous* class until its queue is empty or the *tht* expires

If station is not implementing priority, it passes the token

Station copies the residue of its trt_{ua} into the *tht*

Station reloads the trt_{ua} with TRT_{UA}

Station serves its *Urgent Asynchronous* class until its queue is empty or the *tht* expires

SERVICE DISCIPLINE

Station copies the residue of its trt_{na} into the *tht*

Station reloads the trt_{na} with TRT_{NA}

Station serves its *Normal Asynchronous* class until its queue is empty or the *tht* expires

Station copies the residue of its trt_{ta} into the *tht*

Station reloads the trt_{ta} with TRT_{TA}

Station serves its *Time Available* class until its queue is empty or the *tht* expires

Station passes the token to its successor

PERFORMANCE ANALYSIS

802.4 TOKEN BUS

ESTABLISH BASE CONFIGURATION

- 64 stations, always members
- Synchronous traffic only
- Infinite *High_Priority_Token_Hold_Time*
- Constant length 160 bit messages (256 bit frames)
- Max_Inter_Solicit_Count = 255
- Bus capacity = 10,000,000 bits per second
- Error free
- 16 bit addresses
- 96 bit tokens

COMPONENTS OF BUS CAPACITY

At any offered load, 100% of bus capacity is divided among

- data messages
- token traffic
- propagation delays

Graph shows distribution for base configuration

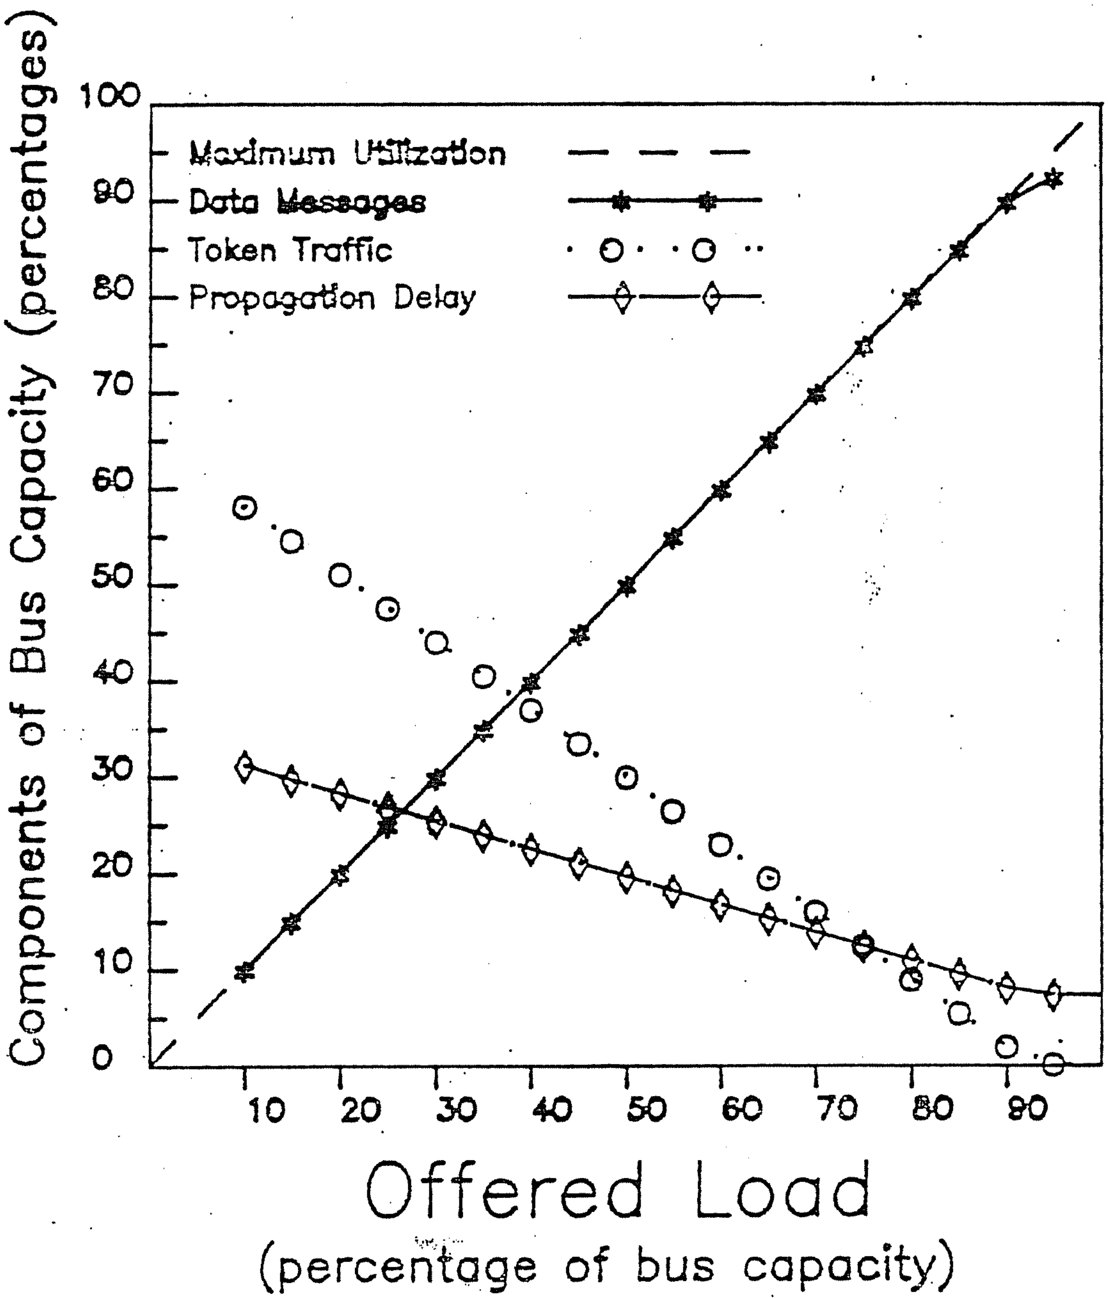

(slide graph1)

AVERAGE MESSAGE DELIVERY TIME

Average message delivery time consists of

- queueing delay
- network access delay
- propagation delay (constant here)

Graph shows components of total observed delay

Note that average message delay is about one-half the token cycle time

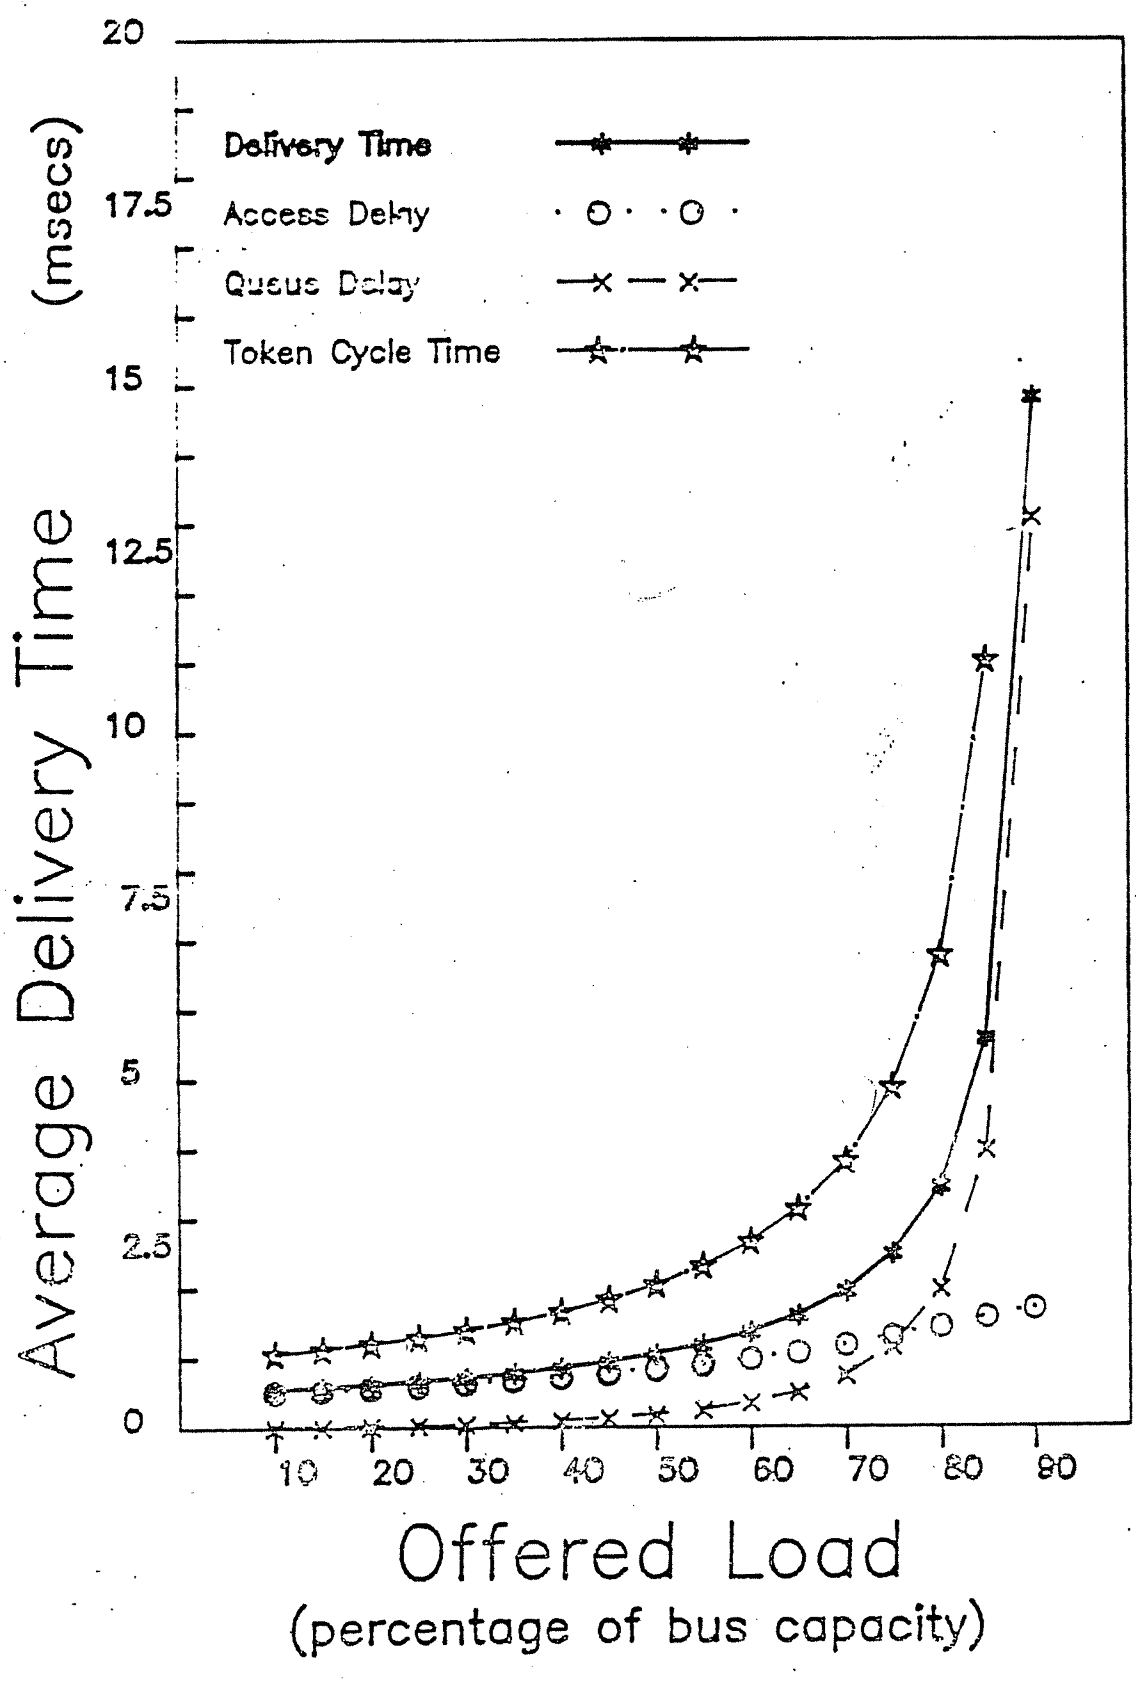

(slide graph2)

IMPACT OF PACKET SIZE

When a station sends multiple packets, they are separated by one inter-frame gap (2 microseconds)

When a station with small packets generates as much offered load as a station with large packets, then at large offered loads the station with small packets suffers longer total delays from its higher frame overhead and delays

Average delivery time is plotted against offered load for data lengths of 160..5120 bits

Frame size is 96 bits longer than data length

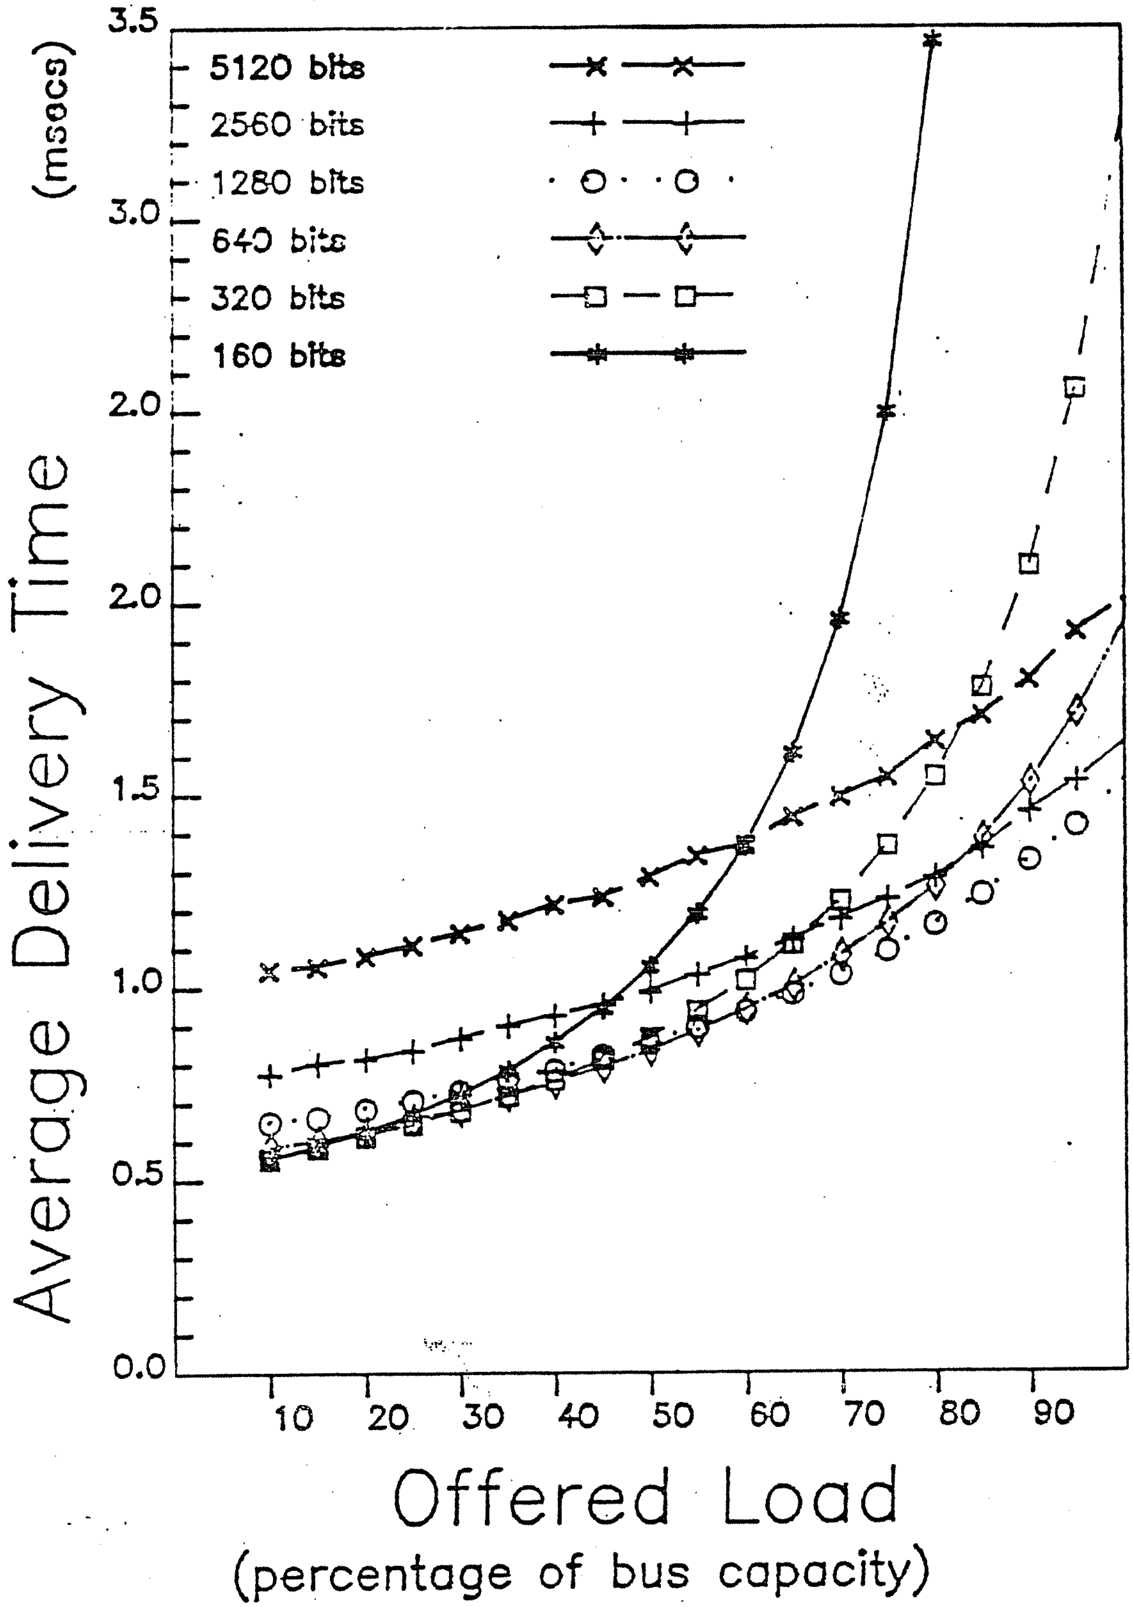

(slide graph3)

NUMBER OF ACTIVE STATIONS

Average message delivery time increases linearly with the number of active stations

Graphs shows 8..256 stations

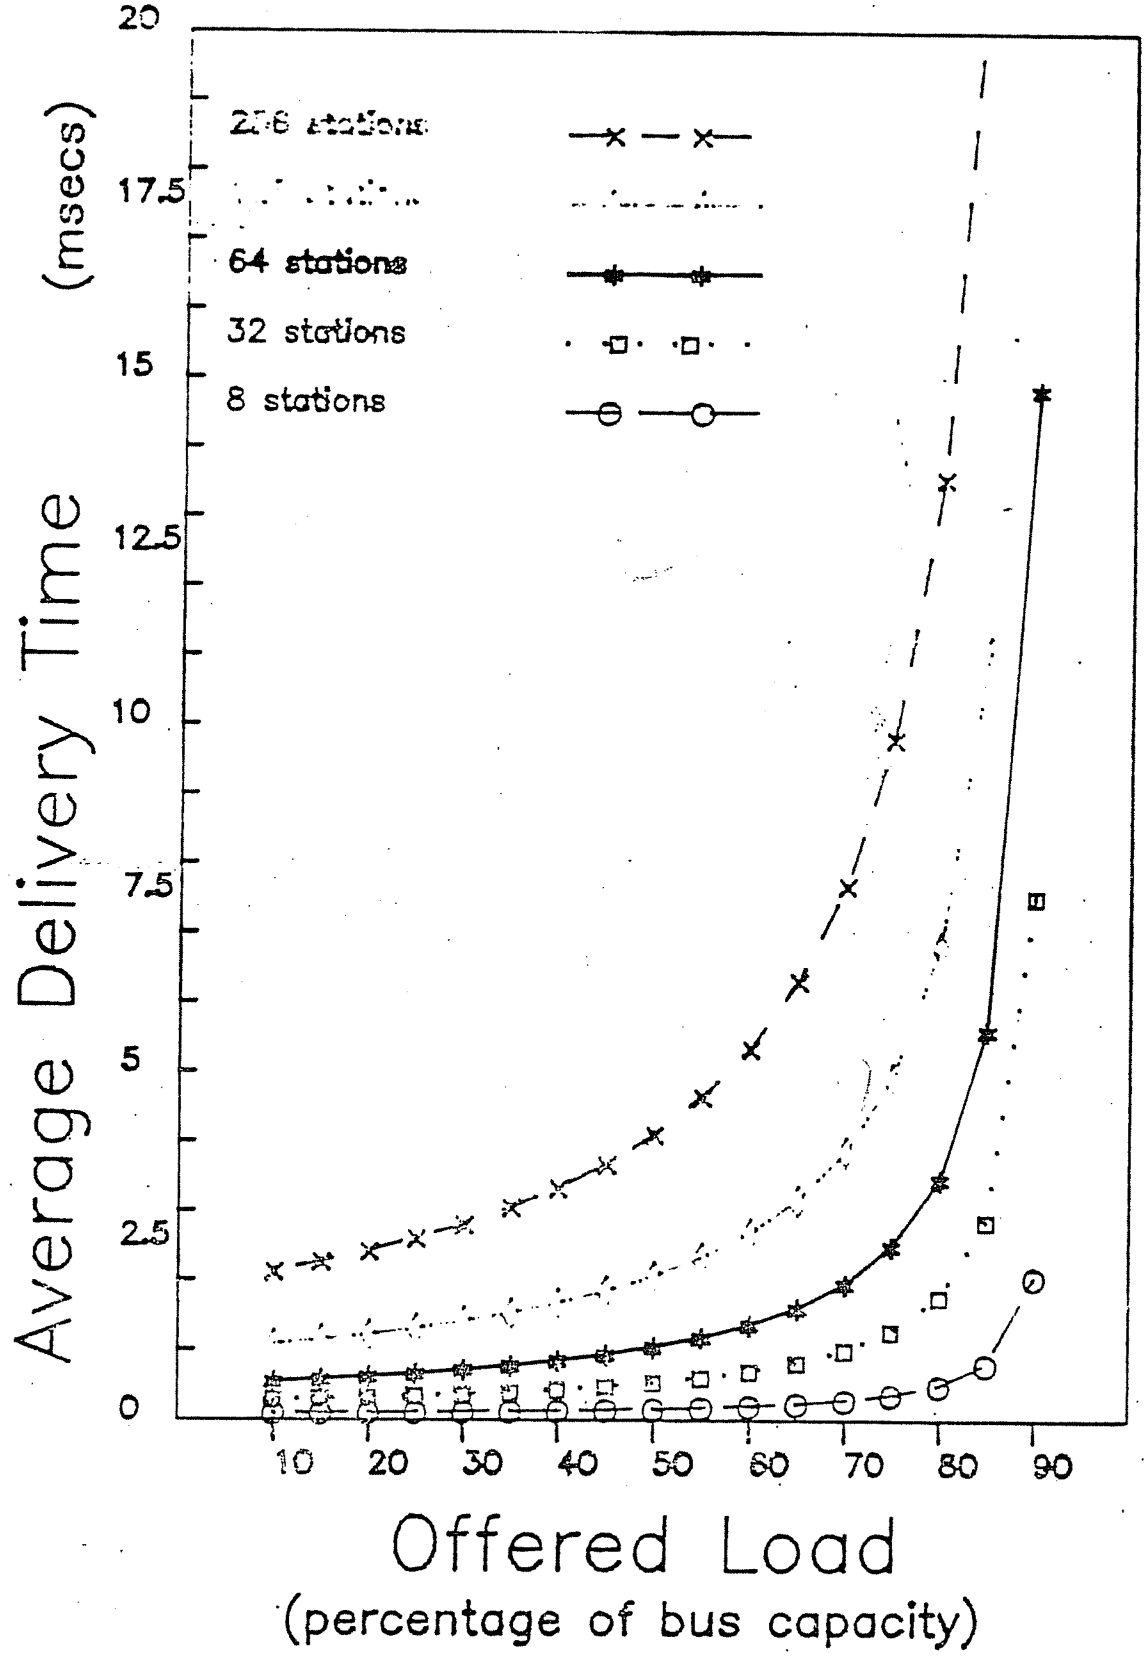

Average Delivery Time (msecs)
20
17.5
15
12.5
10
7.5
5
2.5
0
64 stations
32 stations
8 stations
10
20
30
40
50
60
70
80
90
Offered Load
(percentage of bus capacity)

OBSERVATIONS

For *Synchronous* traffic with *HPTHT* set to infinity:

- Fraction of bus capacity used for data traffic (throughput) increases with offered load because protocol traffic decreases
- Average message delivery time is approximately one-half the token cycle time
- At high offered loads, long messages experience shorter average delivery times than short messages
- Network performance improves as the number of active stations decreases

HIGH PRIORITY TOKEN HOLD TIME

Setting the *HPTHT* has two side-effects:

- As desired, it guarantees a minimum frequency of service to the *Synchronous* class
- The penalty is that it also bounds the maximum data utilization of the bus

As *HPTHT* increases (thus draining the *Synchronous* queues more often and minimizing token traffic), bus utilization for data throughput increases

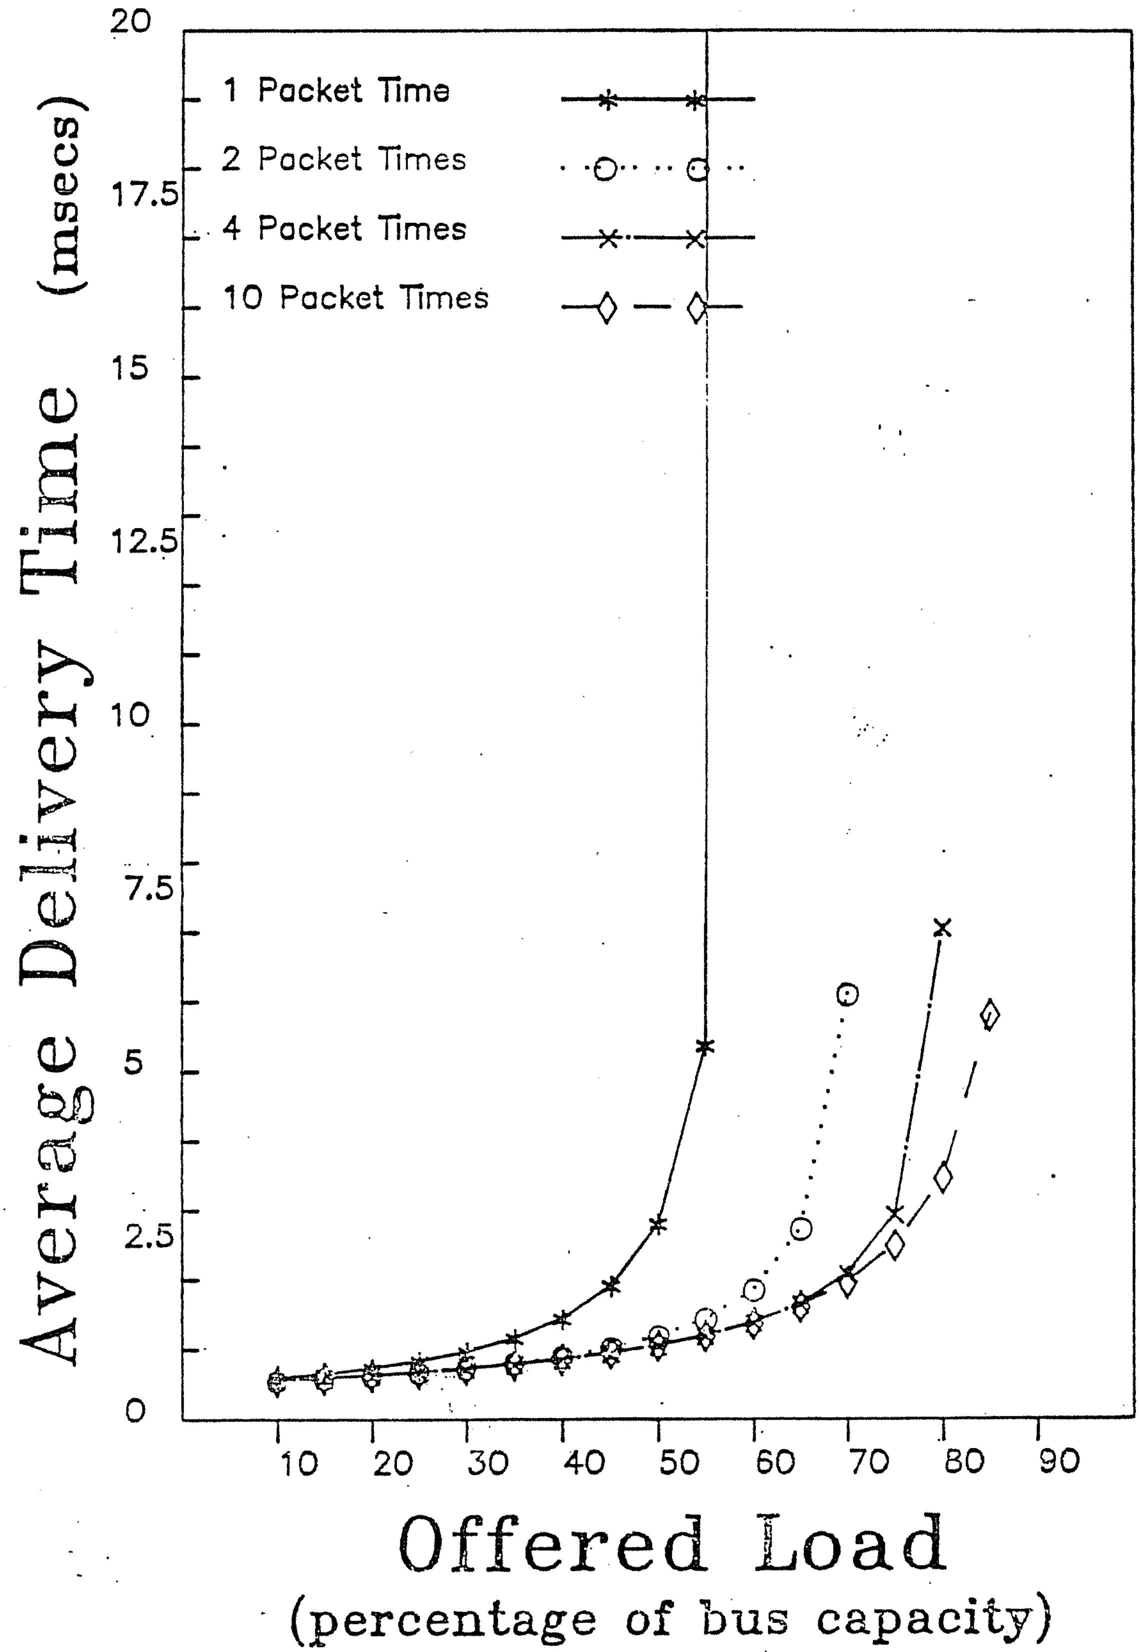

(slide graph5)

NON-SYNCHRONOUS CLASSES

When token cycle time increases to equal the *Target_Rotation_Time* at an access_class, service to that access_class ceases

Graph plots delay of *Synchronous* and *Urgent_Asynchronous* classes

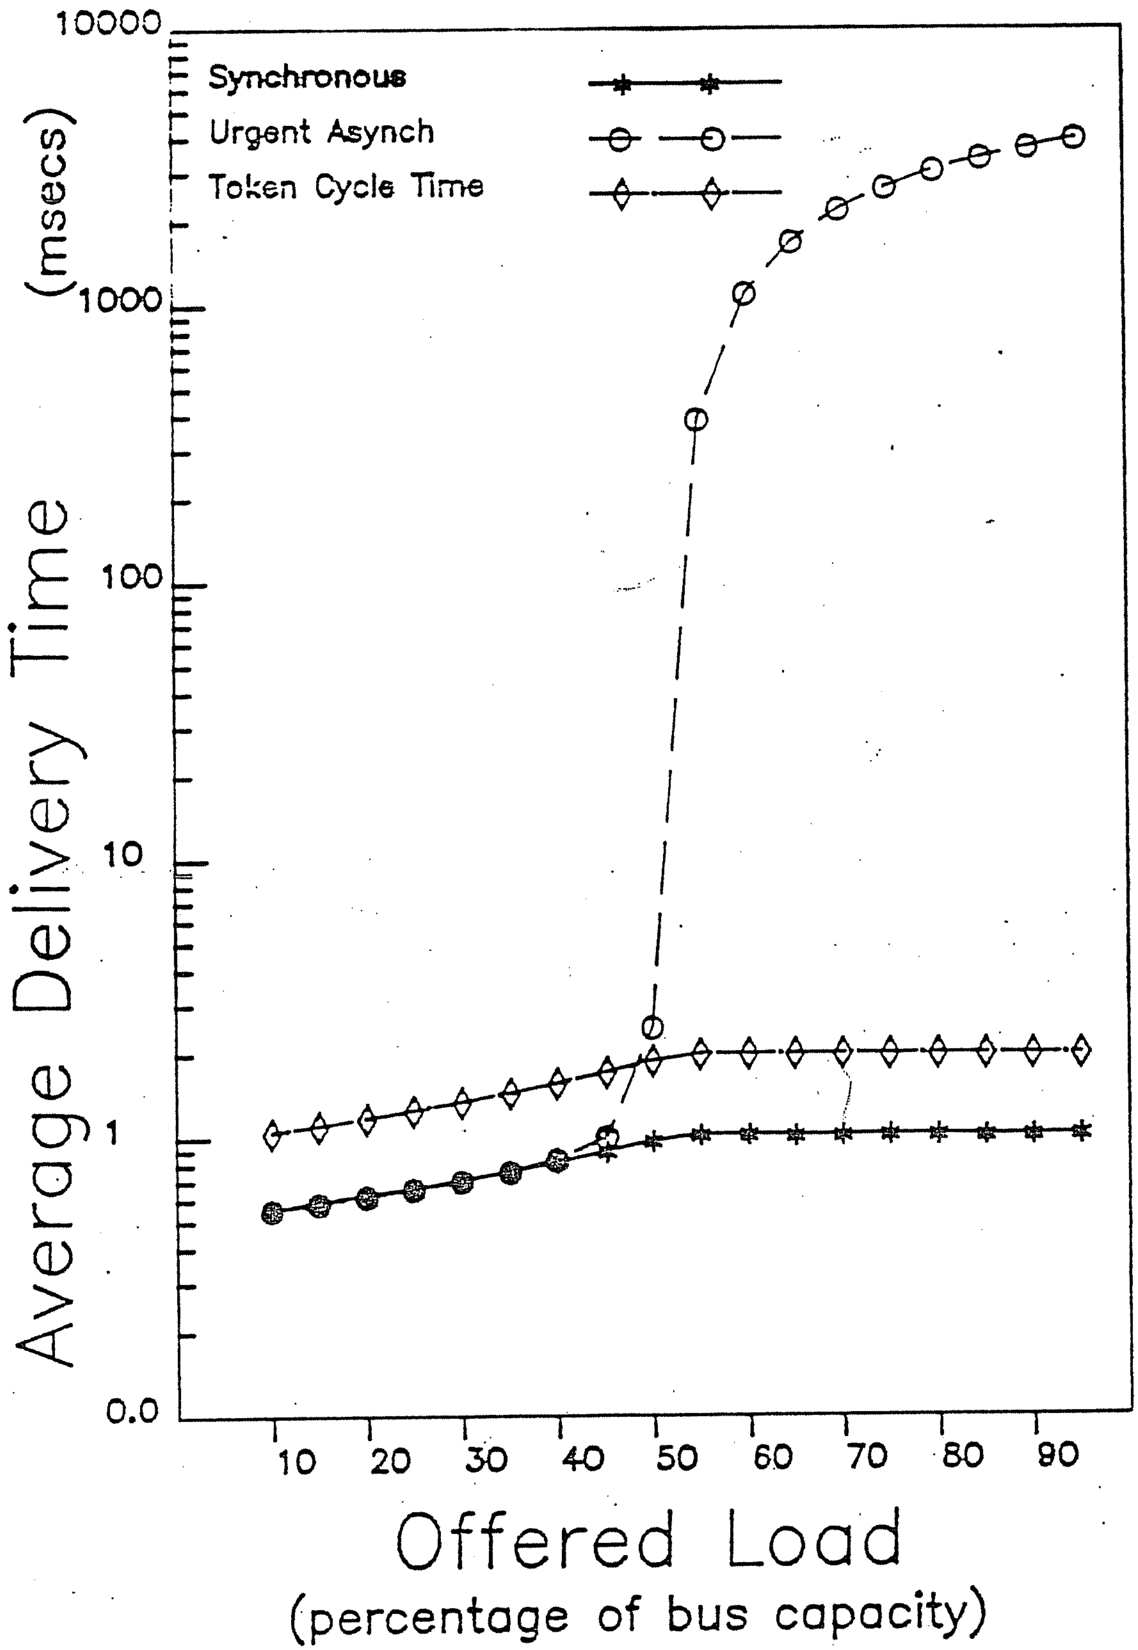

(slide graph6)

OBSERVATIONS

When token cycle time does not approach $N \cdot HPTHT$ or any of the *TRT*'s, the priority system does not function; all messages are transmitted

Setting the *HPTHT* implies an upper bound on the fraction of bus capacity used to carry data

Setting the *TRT*'s effectively implements the priority operation

Setting the *TRT*'s intelligently is fairly difficult and requires both knowledge and insight

NON-HOMOGENEOUS LOADS

The offered load is divided:

- one station generates one-half the total offered load
- 63 stations collectively generate the other half

Graph shows the contributions of queueing delay and network access delay to total message delivery time for each class of station

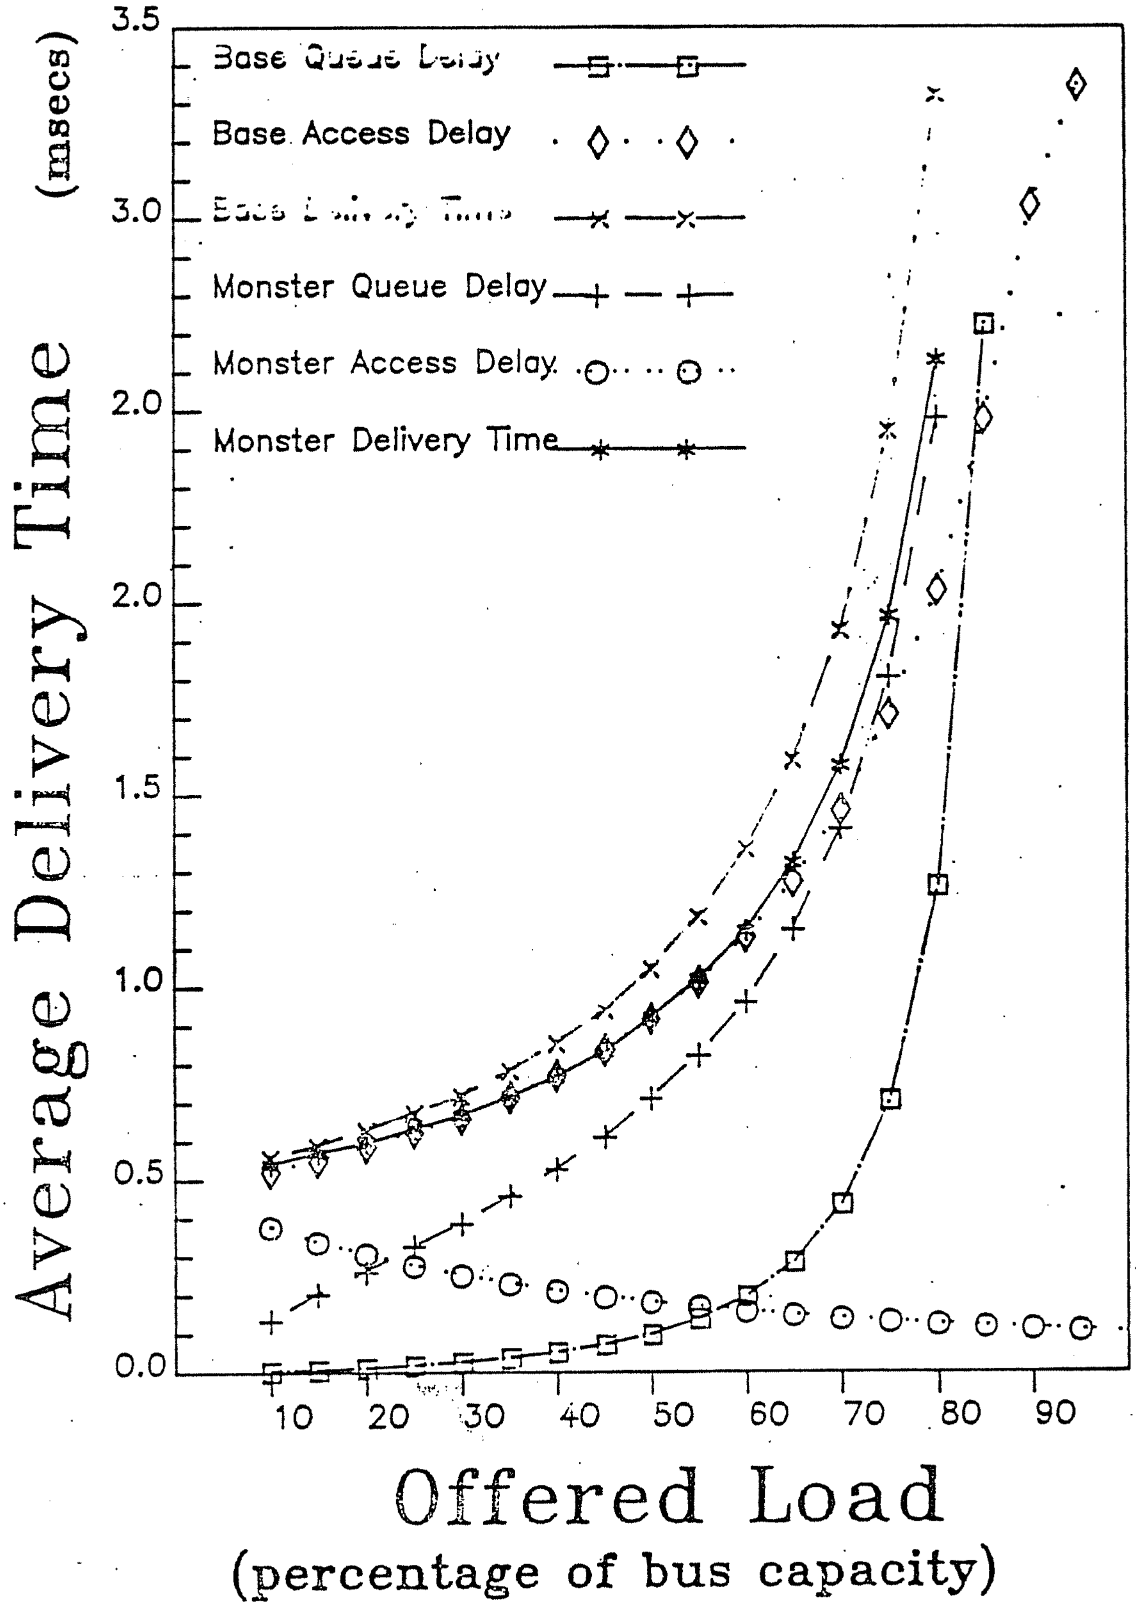
Average Delivery Time (msecs)
3.5
3.0
2.0
2.0
1.5
1.0
0.5
0.0
Base Access Delay
Monster Queue Delay
Monster Access Delay
Monster Delivery Time
10
20
30
40
50
60
70
80
90
Offered Load
(percentage of bus capacity)

OBSERVATIONS

Highly loaded stations receive better service than other stations

The quality of service a station receives depends upon the load generated by the remainder of the network

SECTION 6

IEEE 802.5 TOKEN RING

- Ring architecture
- Similarities and differences from the token bus
- Token passing
- Frame formats
- Priorities
- Acknowledgements

TOKEN RING

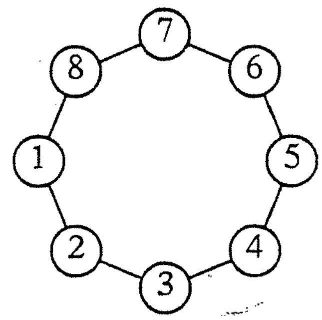

Nodes are physically sequential

Data paths are all point-to-point

Message is purged by its transmitter

Acknowledgement bit at tail of message

Supports eight priorities

TOKEN

START	TOKEN	END

START and END are delimiters

Free token is "0111 1111"

Busy token is "0111 1110"

Free token converted to busy by flipping last bit

Token length 24 bits

TOKEN

Token is a unique data item (one per ring)

Data frames may now exceed the ring circumference

Inefficient at low loads, but does not matter

Efficient at high loads where it does matter

Inherently fair

Adopted by Prime, Apollo, and IBM

FRAME FORMAT

PRE	SD	FC	DA	SA	DATA	FCS	ED

PREAMBLE - pattern to set receiver's clock

SD - start delimiter

FC - frame control

DA - destination address

SA - source address

DATA - information to be transmitted

FCS - frame check sequence, 32 bits

ED - end delimiter

ADDRESSES

I/G	ADDRESS
1	15

I/G	L/G	ADDRESS
1	1	46

I/G	G	MANUFACTURER CODE	SERIAL NUMBER
1	1	n	46-n

ACCESS CONTROL

PPP	T	M	RRR

PPP -- priority of current message

T -- token bit

M -- monitor bit

RRR -- reservation priority

PRIORITY

Supports 8 priority levels

Priority operation:

- station has message of priority P_q
- await next access control field
- if P_q > RRR then
  - — stack current RRR
  - — set new RRR equal to P_q
- await token
- if $P_q \geqslant$ PPP then
  - — transmit message
  - — pop old priority (saved RRR)
  - — set current RRR equal to old priority

FRAME STATUS

A	C	RR	A	C	RR
1	1	2	1	1	2

"A" and "C" are transmitted as "0"

If address is recognized, "A" is set

If frame is copied, "C" is set

Frame returns to transmitter

Transmitter now recognizes:

"A"	"C"	
0	0	station non-existent or not active
1	0	station busy or frame in error
1	1	acknowledgement

SECTION 7

GENERAL MOTORS MAP PROTOCOL

- Summary of usage
- Dependence on 802.4 token bus

GENERAL MOTORS MAP

Purpose:

- Define a MAP message standard which supports application-to-application communication
- Identify application functions to be supported by the message standard
- Recommend protocol(s) that meet GM's functional requirements

Goals:

- The driving force behind the MAP effort is the need for compatibility of communications to integrate the many factory floor devices
- It is the intention of MAP to promote a multi-vendor environment

MAP ARCHITECTURE

Layer 1 -- IEEE 802.4 Broadband

Layer 2 -- IEEE 802.2 Class 1

Layer 3 -- Null

Layer 4 -- ISO Transport Class 4

Layer 5 -- ISO Session Kernel

Layer 6 -- Null

Layer 7 -- ISO CASE Kernel

NETWORK MANAGEMENT

MAP network management will support:

- Monitoring
- Control
- Configuration
- Problem determination
- Recovery

PROGRAMMABLE DEVICES

MAP must support programmable devices through all 7 layers

- programmable controllers
- robots
- CNC machines
- weld controllers

The minimum set of network functions include

- program upload and download
- storage and retrieval of data
- status reporting
- remote diagnostics

INTERCONNECTION

Interconnection of devices is accomplished by

- via direct connection to the broadband LAN if possible
- via bridges to connect segments of a single LAN
- via gateways to dissimilar LANs or long-haul networks
- via routers to interconnect several networks at a common point

Result is CATANET

CATANET

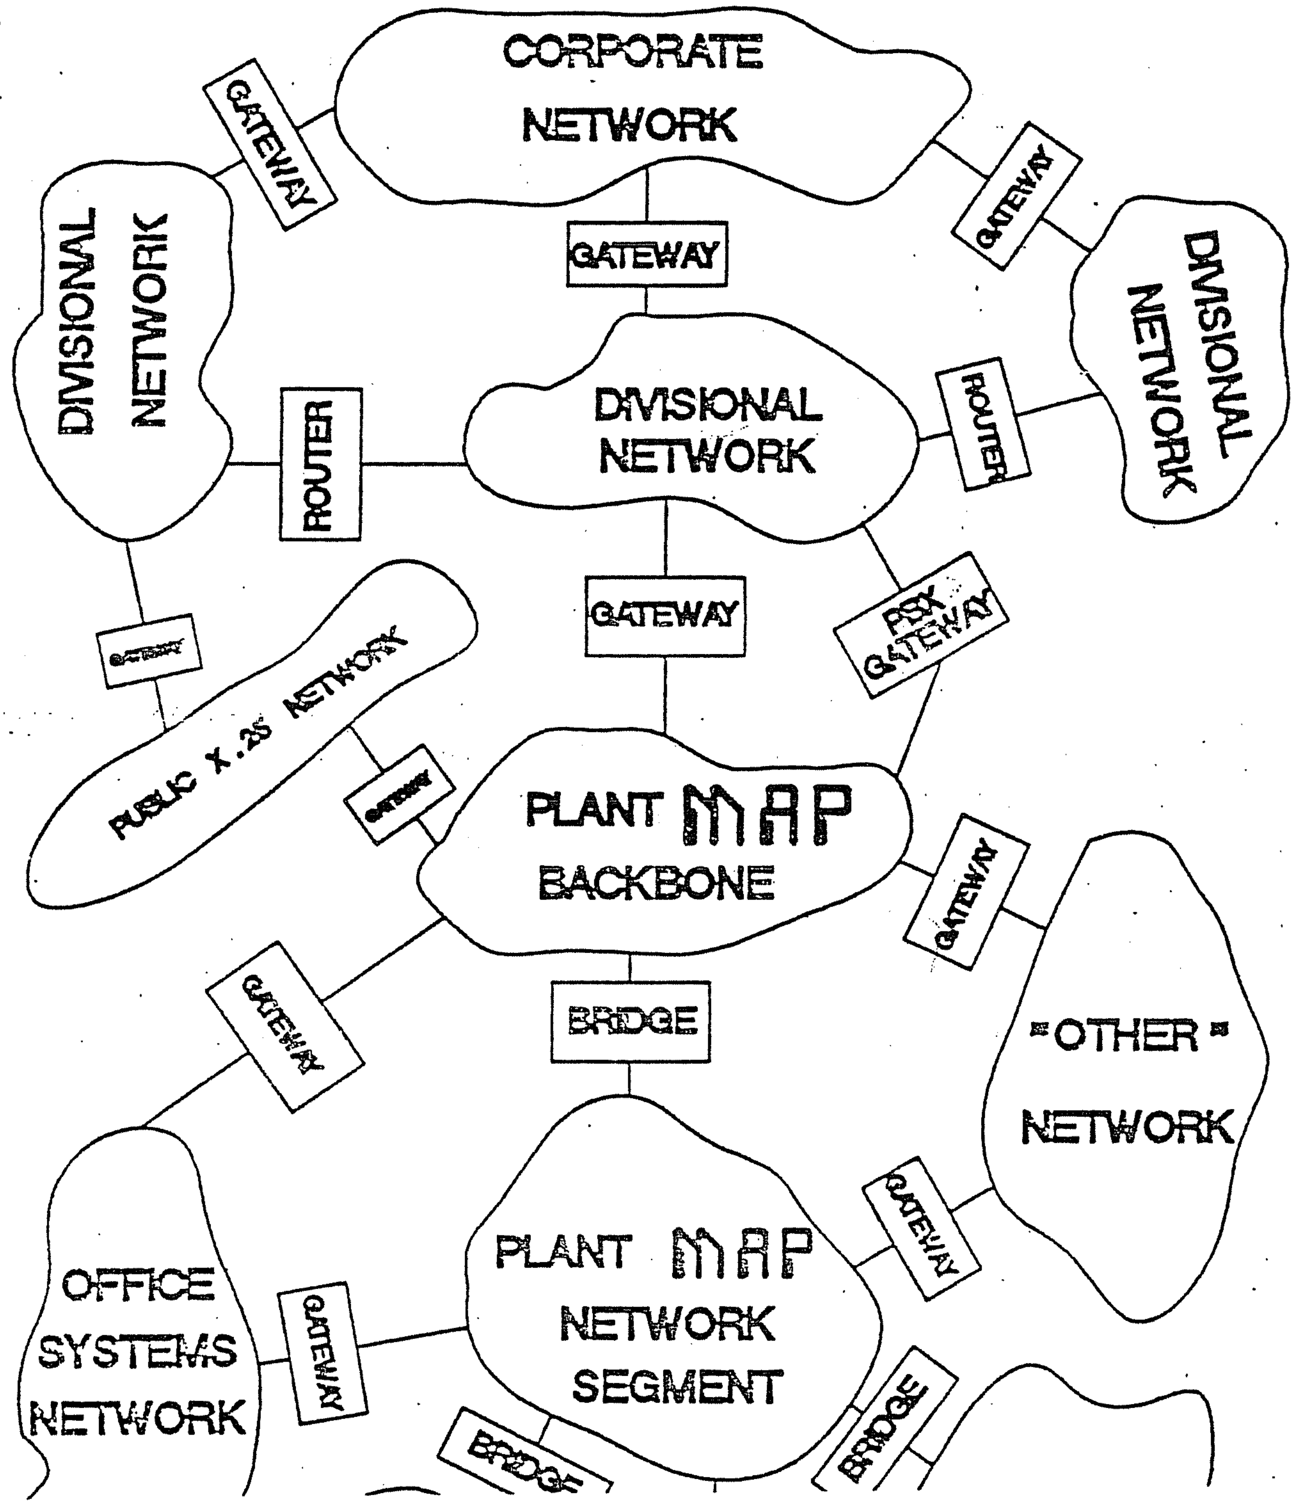

(slide catanet)

IMPLEMENTATION PLAN

Multi-vendor connections via a centralized computer node

Multi-vendor connections via a distributed LAN

Gateways to selected programmable controllers

Gateways to wide area networks

Reduce ISO layers 1,2,3,4 to hardware

Add ISO layers 5,6

In 1988, achieve *plug compatibility* by a majority of suppliers

SECTION 8

SUMMARY