

Chapter 4

(Week 7)

The Medium Access Control Sublayer

ANDREW S. TANENBAUM

COMPUTER NETWORKS

FOURTH EDITION

PP. 247-292

4.1. THE CHANNEL ALLOCATION PROBLEM

4.2. MULTIPLE ACCESS PROTOCOLS

4.3. ETHERNET

4.4. WIRELESS LANS

4.5. BROADBAND WIRELESS

4.6. BLUETOOTH

4.7. DATA LINK LAYER SWITCHING

4.8. SUMMARY

Networks can be divided into two categories:

- 1) Those using **point-to-point** connections and
- 2) Those using **broadcast** channels.

This chapter deals with **broadcast** networks and their protocols.

In any **broadcast network**, the key issue is how to determine who gets to use the channel when there is competition for it.

To make this point clearer, consider a conference call in which six people, on six different telephones, are all connected together so that each one can hear and talk to all the others.

It is very likely that when one of them stops speaking, two or more will start talking at once, leading to chaos.

In a face-to-face meeting, chaos is avoided by external means, for example, at a meeting, people raise their hands to request permission to speak.

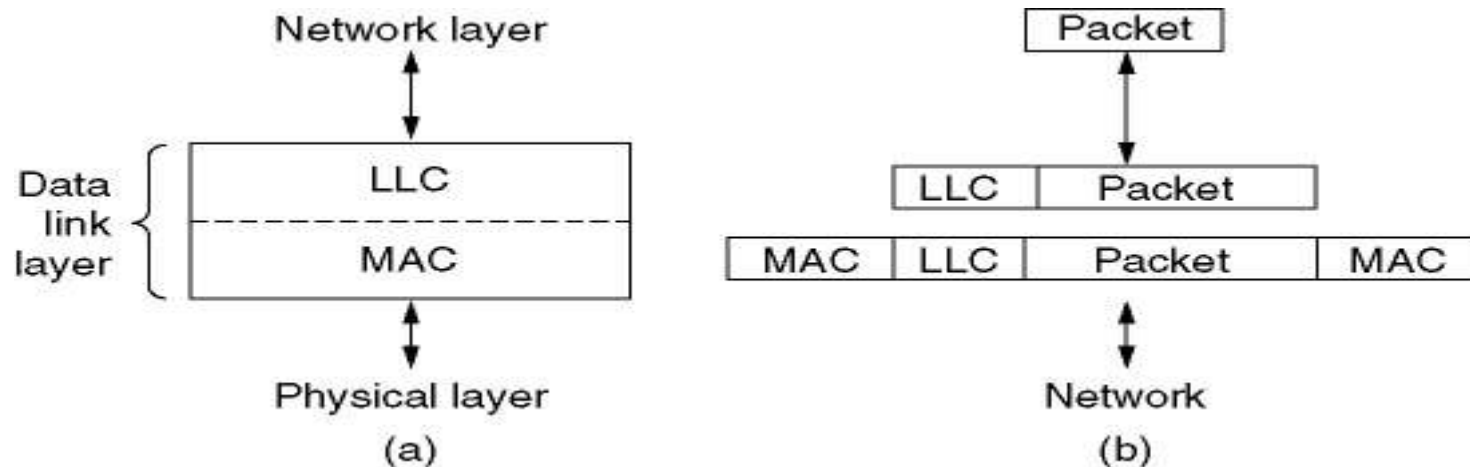
When only a single channel is available, determining **who should go next** is much harder.

Many protocols for solving the problem are known and form the contents of this chapter.

In the literature, broadcast channels are sometimes referred to as **multi-access channels** or **random access channels**

The protocols used to determine who goes next on a multi-access channel belong to a sublayer of the data link layer called the MAC (**Medium Access Control**) sublayer.

Technically, the MAC sublayer is the bottom part of **the data link layer**.



The MAC sublayer is especially important in **LANs**, nearly all of which use a multi-access channel as the basis of their communication.

WANs, in contrast, use point-to-point links, except for satellite networks.

Because multi-access channels and LANs are so closely related, in this chapter we will discuss **LANs in general**, as well as **satellite and some other broadcast networks**

4.1. THE CHANNEL ALLOCATION PROBLEM

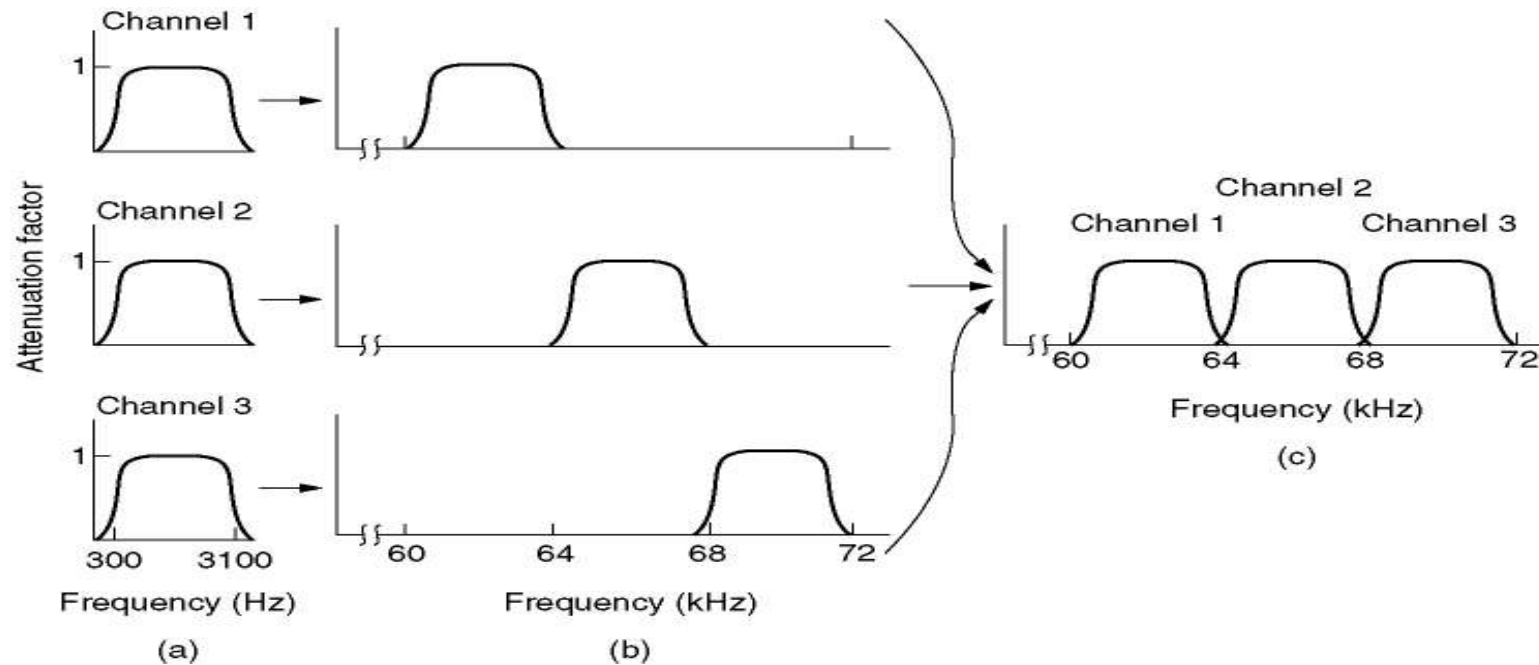
The central theme of this chapter is how to allocate a single broadcast channel among competing users.

We will first look at static and dynamic schemes in general. Then we will examine a number of specific algorithms.

4.1.1. Static Channel Allocation in LANs and MANs

The traditional way of allocation a single channel, such as a telephone trunk, among multiple competing users is **Frequency Division Multiplexing** (FDM). If there are **N** users, the bandwidth is divided into **N** equal sized portions, each user being assigned one portion.

Frequency Division Multiplexing



- (a) The original bandwidths.
- (b) The bandwidths raised in frequency.
- (b) The multiplexed channel.

Since each user has a private frequency band, there is no interference between users

When there is only a small and fixed number of users, each of which has a heavy (buffered) load of traffic (e.g., carriers' switching offices), FDM is a simple and efficient allocation mechanism.

However, when the number of senders is large and continuously varying, or the traffic is bursty, FDM presents some problems.

If the spectrum is cut up into N regions, and fewer than N users are currently interested in communicating, a large piece of valuable spectrum will be wasted.

If more than N users want to communicate, some of them will be denied permission, for lack of bandwidth, even if some of the users who have been assigned a frequency band hardly ever transmit or receive anything.

However, even assuming that the number of users could somehow be held constant at N , dividing the single available channel into static subchannels is inherently inefficient.

The basic problem is that when some users are quiescent, their bandwidth is simply lost. They are not using it, and no one else is allowed to use it either.

Furthermore, in most computer systems, data traffic is extremely bursty (peak traffic to mean traffic ratios of 1000:1 are common).


Consequently, most of the channels will be idle most of the time.

Precisely the same arguments that apply to FDM also apply to time division multiplexing (TDM).

Each user is statically allocated every **N**th time slot.

If a user does not use the allocated slot, it just lies fallow.

Time Division Multiplexing



The T1 carrier (1.544 Mbps).

Since none of the traditional static channel allocation methods work well with bursty traffic, we will now explore dynamic methods.

4.1.2. Dynamic Channel Allocation in LANs and MANs

Before we get into the first of the many channel allocation methods to be discussed in this chapter, it is worthwhile carefully formulating the allocation problem.

Underlying all the work done in this area are five key assumptions, described below.

1. Station Model.

The model consists of **N** independent **stations** (e.g., computers, telephones, or personnel communicators), each with a program or user that generates frames for transmissions. Stations are sometimes called **terminals**.

Once a frame has been generated, the station is blocked and does nothing until the frame has been successfully transmitted.

2. Single Channel Assumption.

A single channel is available for all communication. All stations can transmit on it and all can receive from it.

3. Collision Assumption.

If two frames are transmitted simultaneously, they overlap in time and the resulting signal is garbled. This event is called a **collision**. All stations can detect collisions.

4.(a) Continuous Time.

Frame transmission can begin at any instant. There is no master clock dividing time into discrete intervals.

(b) Slotted Time.

Time is divided into discrete intervals (slots). Frame transmissions always begin at the start of a slot. A slot may contain 0, 1, or more frames, corresponding to an idle slot, a successful transmission, or a collision, respectively.

5.(a) Carrier Sense.

Stations can tell if the channel is in use before trying to use it. If the channel is sensed as busy, no station will attempt to use it until it goes idle.

(b) No Carrier Sense.

Stations cannot sense the channel before trying to use it. They go ahead and transmit. Only later can determine whether the transmission was successful.

4.2. Multiple Access Protocols

Many Algorithms for Allocating a Multiple Access Channel are Known.

Method	Description
FDM	Dedicate a frequency band to each station
WDM	A dynamic FDM scheme for fiber
TDM	Dedicate a time slot to each station
Pure ALOHA	Unsynchronized transmission at any instant
Slotted ALOHA	Random transmission in well-defined time slots
1-persistent CSMA	Standard carrier sense multiple access
Nonpersistent CSMA	Random delay when channel is sensed busy
P-persistent CSMA	CSMA, but with a probability of p of persisting
CSMA/CD	CSMA, but abort on detecting a collision
Bit map	Round robin scheduling using a bit map
Binary countdown	Highest numbered ready station goes next
Tree walk	Reduced contention by selective enabling
MACA, MACAW	Wireless LAN protocols
Ethernet	CSMA/CD with binary exponential backoff
FHSS	Frequency hopping spread spectrum
DSSS	Direct sequence spread spectrum
CSMA/CA	Carrier sense multiple access with collision avoidance

Channel allocation methods and systems for a common channel.

In the following sections we will study a small sample of the more interesting ones and give some examples of their use.

- ALOHA
- Carrier Sense Multiple Access Protocols
- Collision-Free Protocols
- Limited-Contention Protocols
- Wavelength Division Multiple Access Protocols
- Wireless LAN Protocols

ALOHA

Although Abramson's work, called ALOHA system, used ground-based radio broadcasting, the basic idea is applicable to any system in which uncoordinated users are competing for the use of a single shared channel.

We will discuss two versions of ALOHA here: **pure** (1970) and **slotted** (1972). They differ with respect to whether time is divided into discrete slots into which all frames must fit. Pure ALOHA does not require global time synchronization; slotted ALOHA does.

Pure ALOHA (1)

The basic idea of an ALOHA system is simple: let users transmit whenever they have data to be sent.

There will be collisions, of course, and the colliding frames will be damaged.

However, due to the feedback property of broadcasting, a sender can always find out whether its frame was destroyed by listening to the channel, the same way other users do.


Pure ALOHA (2)

If listening while transmitting is not possible for some reason, acknowledgements are needed.

If the frame was destroyed, the sender just waits a random amount of time and sends it again.

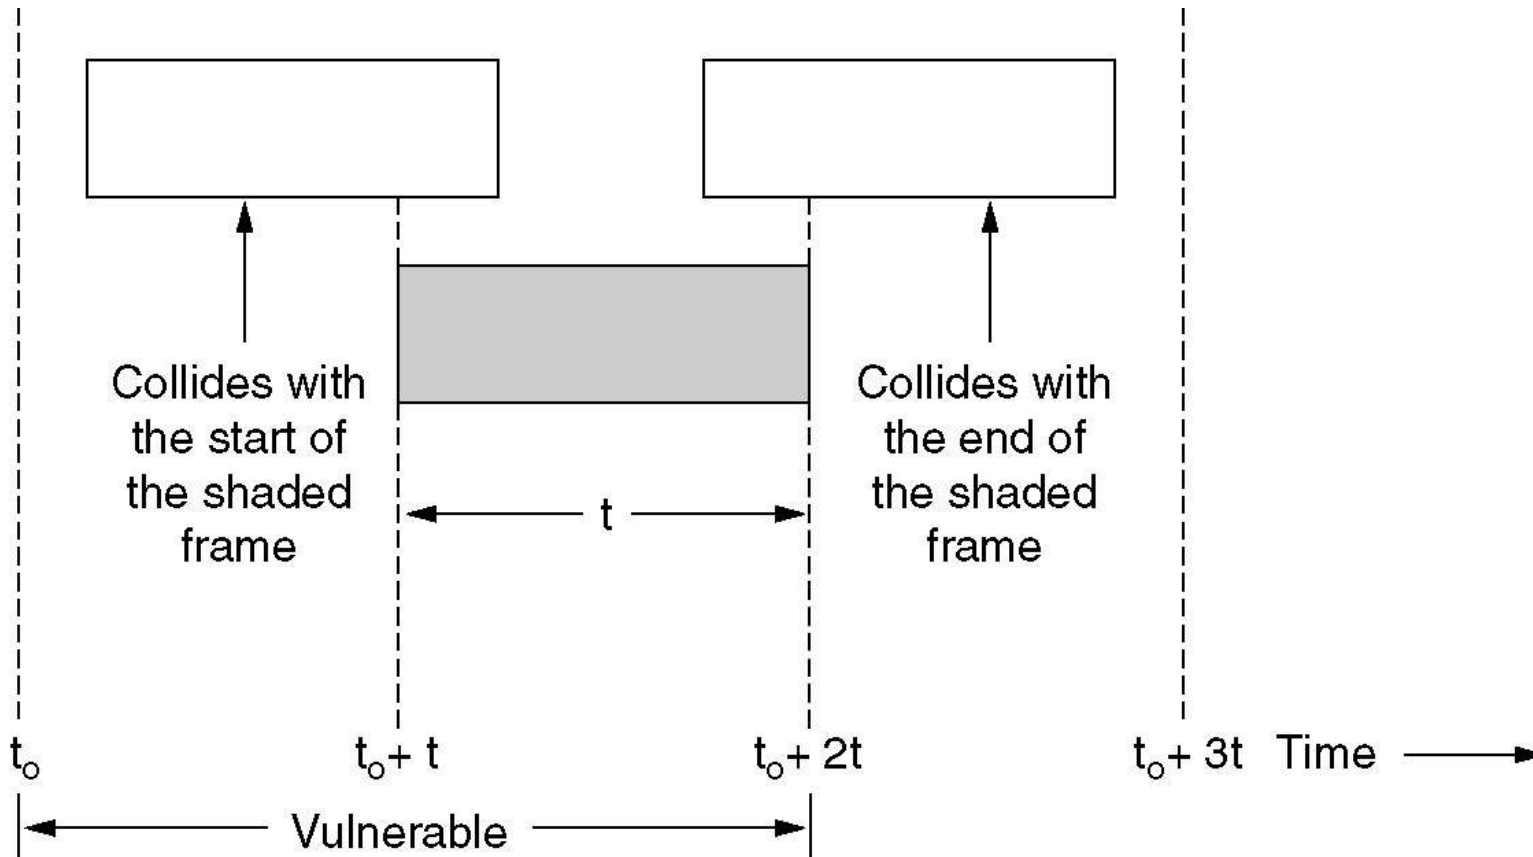
Systems in which multiple users share a common channel in a way that can lead to conflicts are widely known as **contention systems**.

Pure ALOHA (3)



In pure ALOHA, frames are transmitted at completely arbitrary times.

Pure ALOHA (4)



Vulnerable period for the shaded frame.

Pure ALOHA (5)

- With pure ALOHA the best channel utilization that can be achieved is

$$S = Ge^{-2G}$$

for pure ALOHA

$$G = 0.5$$

- and

$$S = 1/2e \text{ or } 18 \text{ percent}$$

Slotted ALOHA (1)

In 1972, Roberts published a method for doubling the capacity of an ALOHA system. His proposal was to divide time into discrete intervals, each interval corresponding to one frame.

This approach requires the users to agree on slot boundaries. One way to achieve synchronization would be to have one special station emit a pip at the start of each interval, like a clock.

Slotted ALOHA (2)

In Roberts' method, which has come to be known as **slotted ALOHA**, in contrast to Abramson's **pure ALOHA**, a computer is not permitted to send whenever a carriage return is typed. Instead, it is required to wait for the beginning of the next slot

Thus, the continuous pure ALOHA is turned into a discrete one.

Slotted ALOHA (3)

- With slotted ALOHA the best channel utilization that can be achieved is

$$S = Ge^{-G}$$


for slotted ALOHA

$$G = 1$$

- and

$$S = 1/e \text{ or } 36 \text{ percent}$$

Pure and Slotted ALOHA



Throughput versus offered traffic for ALOHA systems.

4.2.2. Carrier Sense Multiple Access Protocols (1)

With slotted ALOHA the best channel utilization that can be achieved is $1/e$. This is hardly surprising, since with stations transmitting at will, without paying attention to what the other station are doing, there are bound to be many collisions.

Carrier Sense Multiple Access Protocols (2)

In local area networks, however, it is possible for station to detect what other station are doing, and adapt their behavior accordingly. This networks can achieve a much better utilization than $1/e$.

In this section we will discuss some protocols for improving performance

Carrier Sense Multiple Access Protocols (3)

Protocols in which stations listen for a carrier (i.e., a transmission) and act accordingly are called **carrier sense protocols**.

A number of them have been proposed.

- 1) 1-persistent CSMA (Carrier Sense Multiple Access)
- 2) nonpersistent CSMA
- 3) p-persistent CSMA
- 4) CSMA with Collision Detection (CSMA/CD)

1-persistent CSMA (Carrier Sense Multiple Access)

When a station has data to send, it first listens to the channel to see if anyone else is transmitting at that moment. If the channel is busy, the station waits until it becomes idle. When the station detects an idle channel, it transmits a frame. If a collision occurs, the station waits a random amount of time and starts all over again. The protocol is called 1-persistent because the station transmits with a probability of 1 when it finds the channel idle.

Nonpersistent CSMA (Carrier Sense Multiple Access)

Before sending, a station senses the channel. If no one else is sending, the station begins doing so itself. However, if the channel is already in use, the station does not continually sense it for the purpose of seizing it immediately upon detecting the end of the previous transmission. Instead, it waits a random period of time and then repeats the algorithm. Consequently, this algorithm leads to better channel utilization but longer delays than 1-persistent CSMA.

p-persistent CSMA (Carrier Sense Multiple Access)


It applies to slotted channel and works as follows. When a station becomes ready to send, it sense the channel. If it is idle, it transmits with a probability p . With a probability $q=1-p$, it defers until the next slot. If that slot is also idle, it either transmits or defers again, with probabilities p and q . This process is repeated until either the frame has been transmitted or another station has begun transmitting.

CSMA with Collision Detection(1)

Another improvement over ALOHA is for station to abort their transmissions as soon as they detect a collision. In other words, if two stations sense the channel to be idle and begin transmitting simultaneously, they will both detect the collision almost immediately. Rather than finish transmitting their frames, which are irretrievably garbled anyway, they should abruptly stop transmitting as soon as the collision is detected. Quickly terminating damaged frames saves time and bandwidth.


This protocol is widely used on LANs in the MAC sublayer.

CSMA with Collision Detection (2)



CSMA/CD can be in one of three states: contention, transmission, or idle.

CSMA PROTOCOLS' PERFORMANCE



Comparison of the channel utilization versus load for various random access protocols.

Collision-Free Protocols (1)

Although collisions do not occur with CSMA/CD once a station has unambiguously captured the channel, they can still occur during the contention period. These collisions adversely affect the system performance, especially when the cable is long and the frames are short. And CSMA/CD is not universally applicable.

Collision-Free Protocols (2)

In this section, we will examine two protocols that resolve the contention for channel without any collisions at all, not even during the contention period.

- 1) a bit-map protocol
- 2) binary countdown

Collision-Free Protocols (3)

Assumptions for these two protocols:

- 1) there are exactly n stations, each with a unique address from 0 to $n-1$ “wired” into it.
- 2) propagation delay is negligible

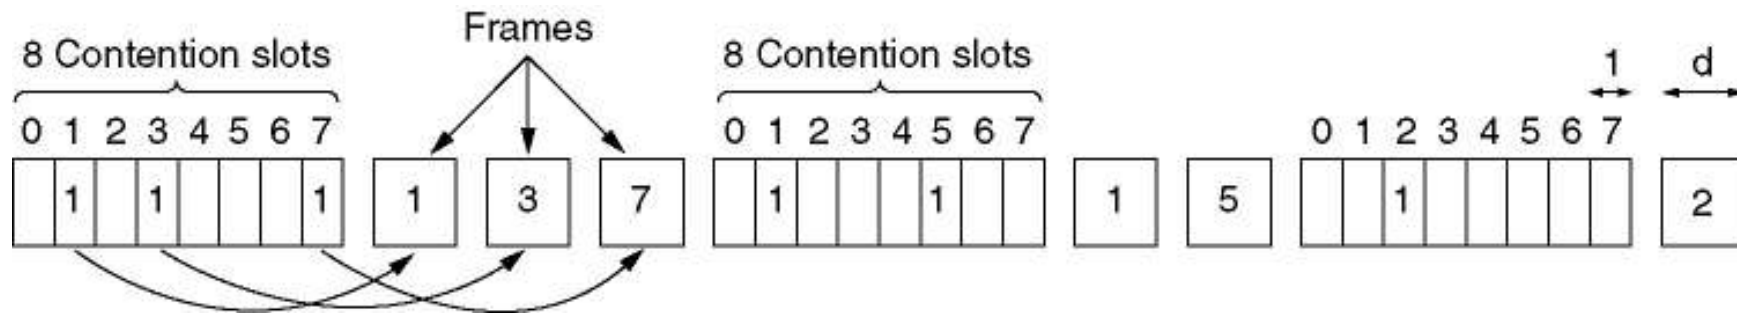
Collision-Free Protocols (4)

BASIC BIT-MAP PROTOCOL(1)

- In this protocol, each contention period consists of exactly n slots.
- If station 0 has a frame to send, it transmits a 1 bit during the zeroth slot.
- No other station is allowed to transmit during this slot.
- In general, station j may announce that it has a frame to send by inserting a 1 bit into slot j .

Collision-Free Protocols (5)

BASIC BIT-MAP PROTOCOL(2)



The basic bit-map protocol.

Collision-Free Protocols (6)

BINARY COUNTDOWN (1)

A problem with the basic bit-map protocol is that the overhead is 1 bit per station, so it does not scale well to networks with thousands of stations.

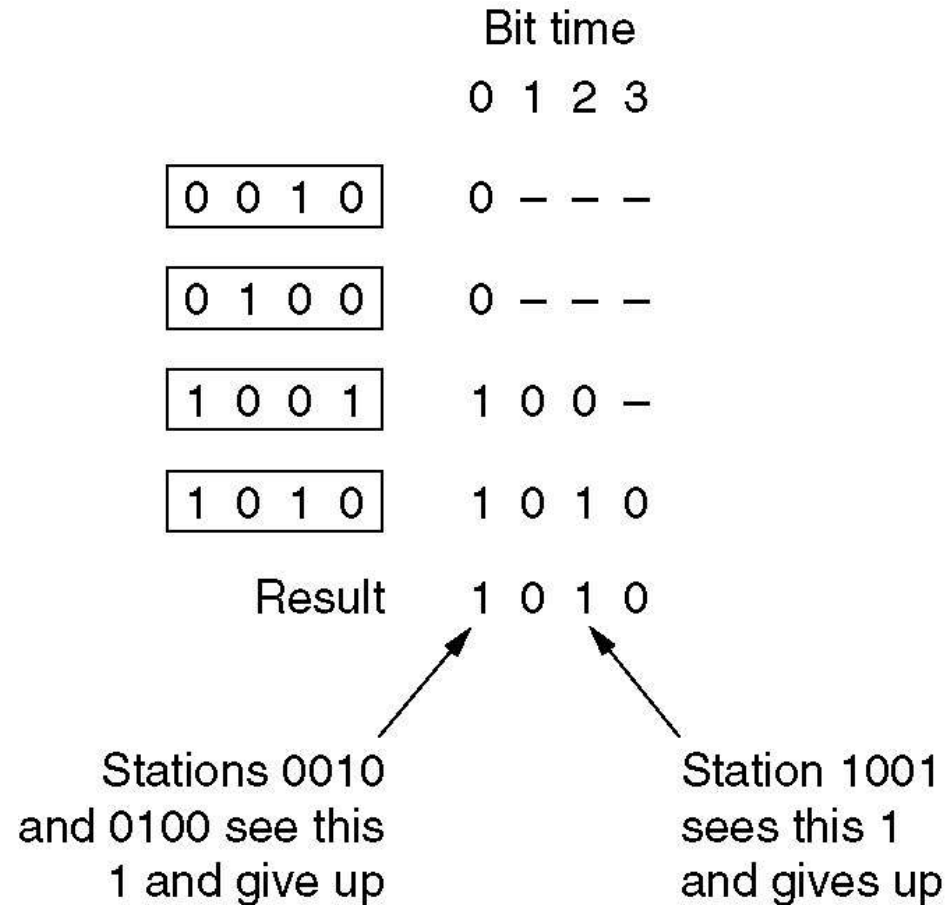
Collision-Free Protocols (7)

BINARY COUNTDOWN (2)

- We can do better than that by using binary station addresses.
- A station wanting to use the channel now broadcasts its address as a binary bit string, starting with the high-order bit.
- All addresses are assumed to be the same length.
- The bits in each address position from different stations are Boolean ORed together.

Collision-Free Protocols (8)

BINARY COUNTDOWN (3)



The binary countdown protocol. A dash indicates silence.

Limited-Contention Protocols (1)

We have now considered two basic strategies for channel acquisition in a cable network:

- 1) contention, as in CSMA
- 2) collision-free methods

Limited-Contention Protocols (2)

Each strategy can be rated as to how well it does with respect to the two important performance measures:

- 1) delay at low load
- 2) channel efficiency at high load.

Limited-Contention Protocols (3)

Under conditions of light load, contention (i.e., pure or slotted ALOHA) is preferable due to its low delay.

As the load increases, contention become increasingly less attractive, because the overhead associated with channel arbitration becomes greater.

Limited-Contention Protocols (4)

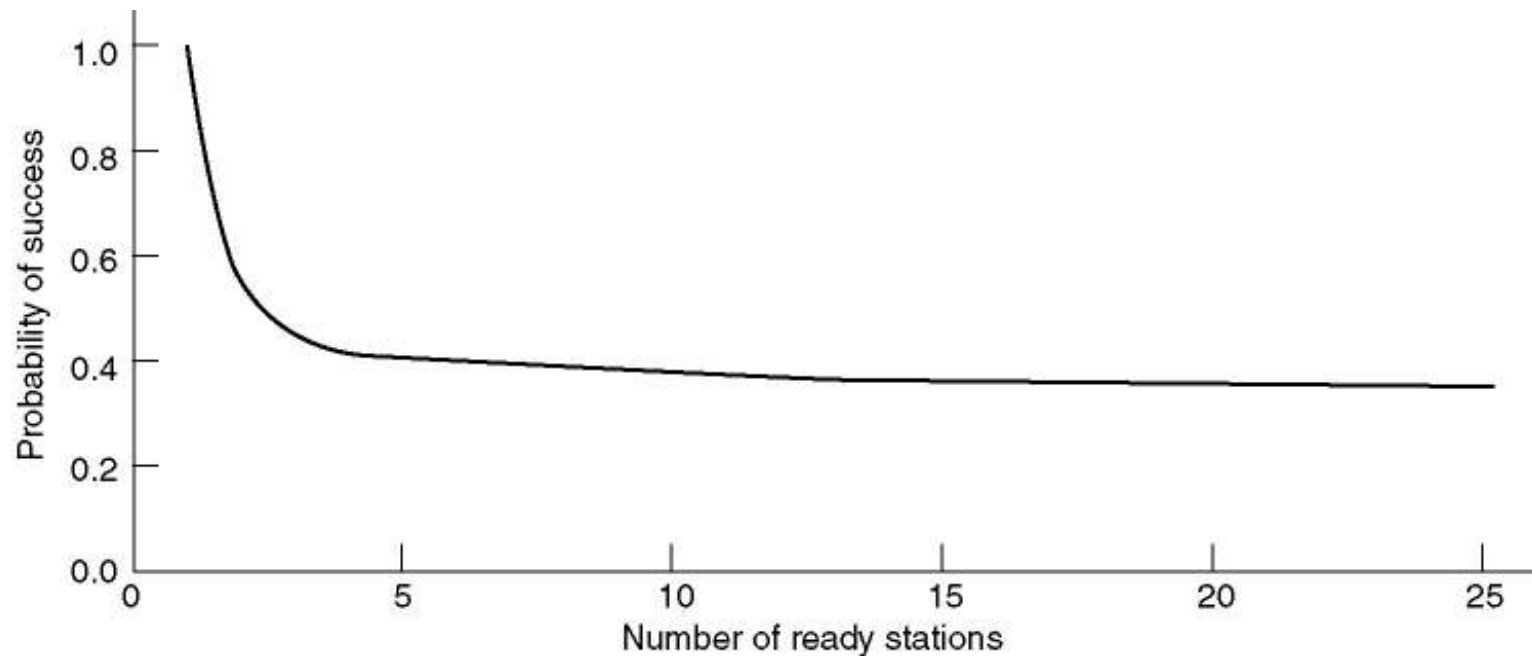
Just the reverse is true for the collision-free protocols.

At low load, they have high delay, but as the load increase, the channel efficiency improves rather than gets worse as it does for contention protocols.

Limited-Contention Protocols (5)

Obviously, it would be nice if we could combine the best properties of the contention and collision-free protocols, arriving at a new protocol that used contention at low load to provide low delay, but used a collision-free technique at high load to provide good channel efficiency.

Limited-Contention Protocols (6)



Acquisition probability for a symmetric contention channel.

Limited-Contention Protocols (7)

From previous figure, it is fairly obvious that the probability of some station acquiring the channel can be increased only by decreasing the amount of competition.

Limited-contention protocols do precisely that.

Limited-Contention Protocols (8)

They first divide the stations into groups.

Only the members of group 0 are permitted to complete for slot 0.

If one of them succeeds, it acquires the channel and transmits its frame.

If the slot lies fallow or if there is a collision, the members of group 1 contend for slot 1, etc

The Adaptive Tree Walk Protocol (1)

As more and more stations are assigned to the same slot in limited - contention protocols, the probability of a collision grows, but the length of the bit-map scan needed to give everyone a chance shrinks.

The limiting case is a single group containing all stations (i.e., slotted ALOHA)

The Adaptive Tree Walk Protocol (2)

What we need is a way to assign stations to slots dynamically, with many stations per slot when the load is low and few (or even just one) station per slot when the load is high.

The Adaptive Tree Walk Protocol (3)

It is convenient to think of the stations as the leaves of a binary tree.

In the first contention slot following a successful frame transmission, slot 0, all stations are permitted to try to acquire the channel.

The Adaptive Tree Walk Protocol (4)

If one of them does so, fine.


If there is a collision, then during slot 1 only those stations falling under node 2 in the tree may compete.

If one of them acquires the channel, the slot following the frame is reserved for those stations under node 3

The Adaptive Tree Walk Protocol (5)

If, on the other hand, two or more stations under node 2 want to transmit, there will be a collision during slot 1, in which case it is node 4's turn during slot 2.

Adaptive Tree Walk Protocol (5)

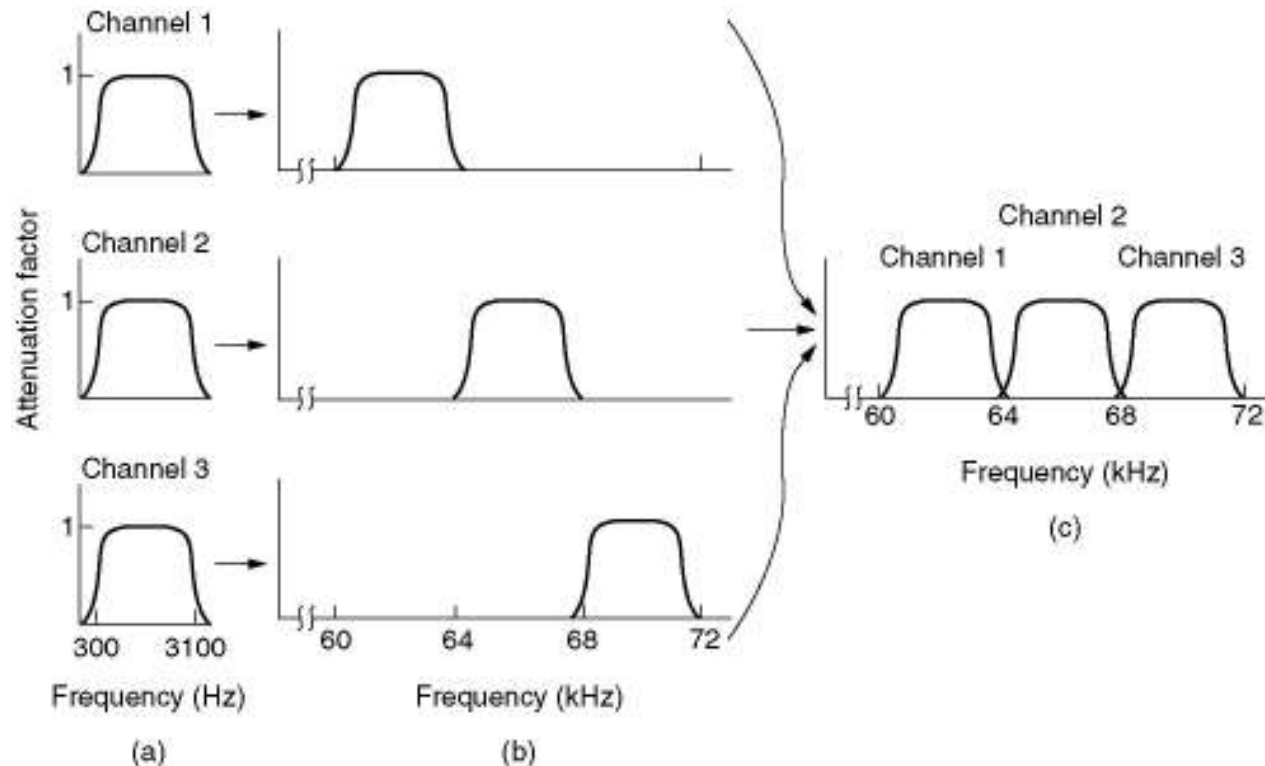


The tree for eight stations.

Wavelength Division Multiple Access Protocols(1)

To allow multiple transmission at the same time, the spectrum is divided into channels (wavelength bands).

Frequency Division Multiplexing



(a) The original bandwidths.

(b) The bandwidths raised in frequency.

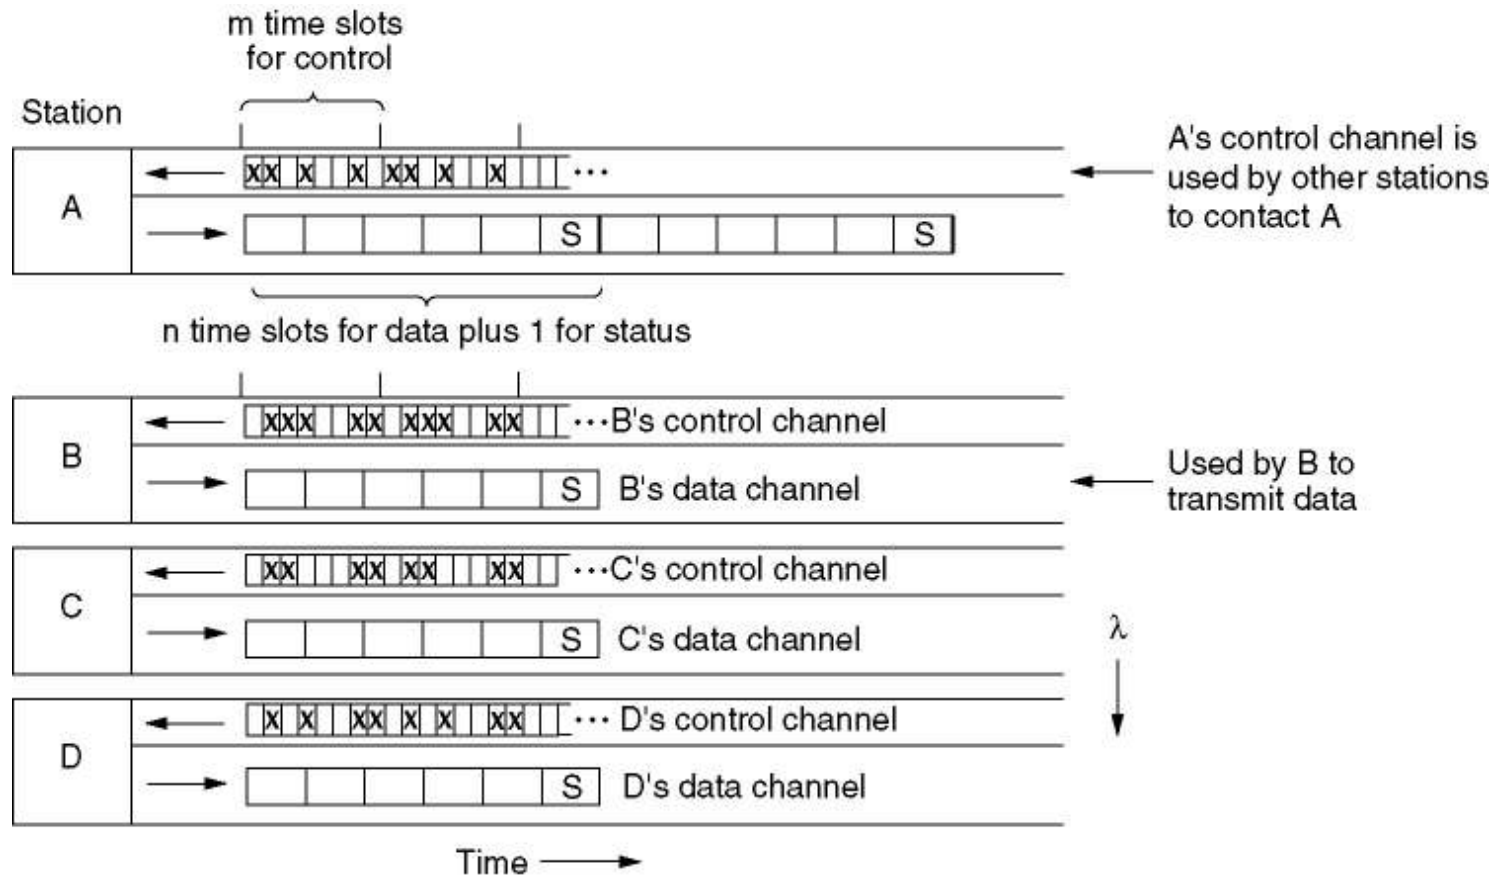
(c) The multiplexed channel.

Wavelength Division Multiple Access Protocols(4)

In this protocol, **WDMA**, each station is assigned two channels.

A narrow channel is provided as a control channel to signal the station, and a wide channel is provided so the station can output data frames.

Wavelength Division Multiple Access Protocols



Wavelength division multiple access.

Wavelength Division Multiple Access Protocols(6)

Each channel is divided into groups of time slots.

On both channels, the sequence of slots repeats endlessly, with slot 0 being marked in a special way so latecomers can detect it.

All channels are synchronized by a single global clock.

Wireless LAN Protocols

- To achieve true mobility, notebook computers need to use radio (or infrared) signals for communication.
- A system of notebook computers that communicate by radio can be regarded as a wireless LAN.
- These LANs have somewhat different properties than conventional LANs and require special MAC sublayer protocols.


Wireless LAN Protocols



A wireless LAN. (a) A transmitting. (b) B transmitting.

- A naive approach to using a wireless LAN might be to try CSMA: just listen for other transmissions and only transmit if no one else is doing so.
- The trouble is interference at the receiver, not at the sender.

Wireless LAN Protocols (2)



The MACA protocol. (a) A sending an RTS to B.

(b) B responding with a CTS to A.

MACA – Multiple Access with Collision Avoidance

MACAW – MACA for Wireless

RTS – Request To Send

CTS – Clear To Send

4.3 Ethernet

We have now finished our general discussion of channel allocation protocols, so it is time to see how these principles apply to real systems, in particular, LANs.

Both the **INTERNET** and **ATM** were designed for **Wide Area Networking**.

However, many companies universities, and other organizations have large numbers of computers that must be connected.

This need gave rise to the **Local Area Network**.

In this section we will say a little bit about the most popular LAN, **ETHERNET**

ETHERNET'S TOPICS

- Ethernet Cabling
- Manchester Encoding
- The Ethernet MAC Sublayer Protocol
- The Binary Exponential Backoff Algorithm
- Ethernet Performance
- Switched Ethernet
- Fast Ethernet
- Gigabit Ethernet
- IEEE 802.2: Logical Link Control
- Retrospective on Ethernet

First Ethernet architecture:


Transmission medium: thick coaxial cable (the ether) up to 2.5 km long (with repeaters every 500 meters)

Up to 256 machines could be attached to the system via transceivers screwed onto the cable

The system ran at 2.94 mbps

A cable with multiple machines attached to it in parallel is called **a multidrop cable**

Ethernet



a) Architecture of the original Ethernet.

The IEEE has standardized a number of local area networks and metropolitan area networks under the name of IEEE 802

IEEE 802 Standards

Number	Topic
802.1	Overview and architecture of LANs
802.2 ↓	Logical link control
802.3 *	Ethernet
802.4 ↓	Token bus (was briefly used in manufacturing plants)
802.5	Token ring (IBM's entry into the LAN world)
802.6 ↓	Dual queue dual bus (early metropolitan area network)
802.7 ↓	Technical advisory group on broadband technologies
802.8 †	Technical advisory group on fiber optic technologies
802.9 ↓	Isochronous LANs (for real-time applications)
802.10 ↓	Virtual LANs and security
802.11 *	Wireless LANs
802.12 ↓	Demand priority (Hewlett-Packard's AnyLAN)
802.13	Unlucky number. Nobody wanted it
802.14 ↓	Cable modems (defunct: an industry consortium got there first)
802.15 *	Personal area networks (Bluetooth)
802.16 *	Broadband wireless
802.17	Resilient packet ring


The 802 working groups. The important ones are marked with *. The ones marked with ↓ are hibernating. The one marked with † gave up.

Ethernet Cabling

Name	Cable	Max. seg.	Nodes/seg.	Advantages
10Base5	Thick coax	500 m	100	Original cable; now obsolete
10Base2	Thin coax	185 m	30	No hub needed
10Base-T	Twisted pair	100 m	1024	Cheapest system
10Base-F	Fiber optics	2000 m	1024	Best between buildings

The most common kinds of Ethernet cabling.


Ethernet Cabling (2)



Three kinds of Ethernet cabling.


(a) 10Base5, (b) 10Base2, (c) 10Base-T.

Ethernet Cabling (3)



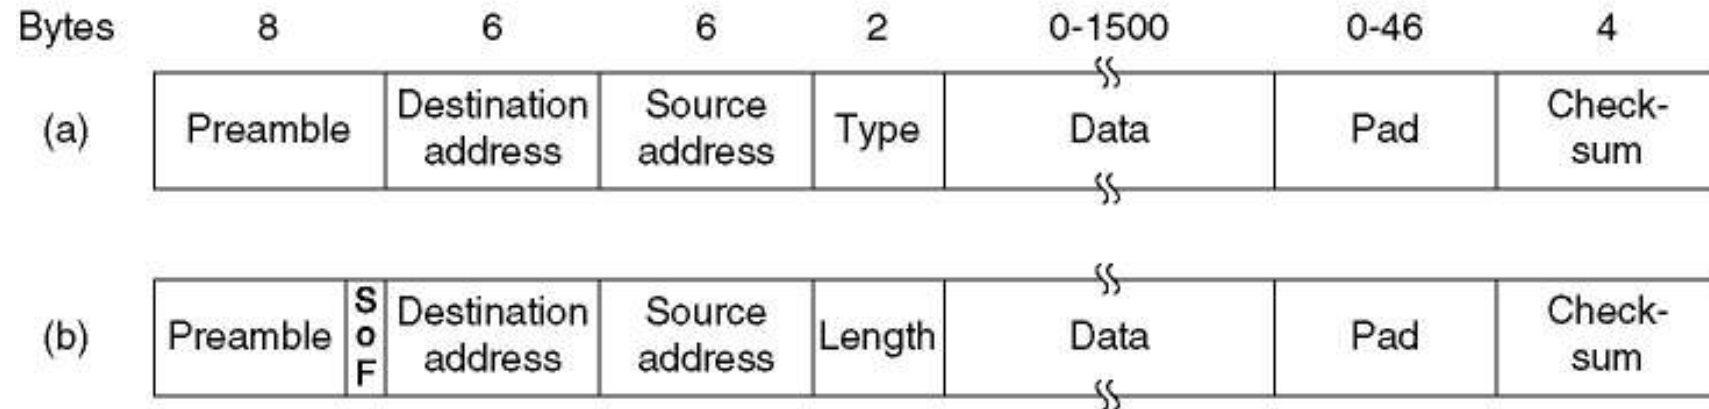
Cable topologies. (a) Linear, (b) Spine, (c) Tree, (d) Segmented.

Ethernet Cabling (4)




(a) Binary encoding, (b) Manchester encoding,
(c) Differential Manchester encoding.

Ethernet MAC Sublayer Protocol




Frame formats. (a) DIX Ethernet, (b) IEEE 802.3.

Ethernet MAC Sublayer Protocol (2)




Collision detection can take as long as 2τ .

Ethernet Performance



Efficiency of Ethernet at 10 Mbps with 512-bit slot times.

Switched Ethernet




A simple example of switched Ethernet.

Fast Ethernet

Name	Cable	Max. segment	Advantages
100Base-T4	Twisted pair	100 m	Uses category 3 UTP
100Base-TX	Twisted pair	100 m	Full duplex at 100 Mbps
100Base-FX	Fiber optics	2000 m	Full duplex at 100 Mbps; long runs

The original fast Ethernet cabling.

Gigabit Ethernet



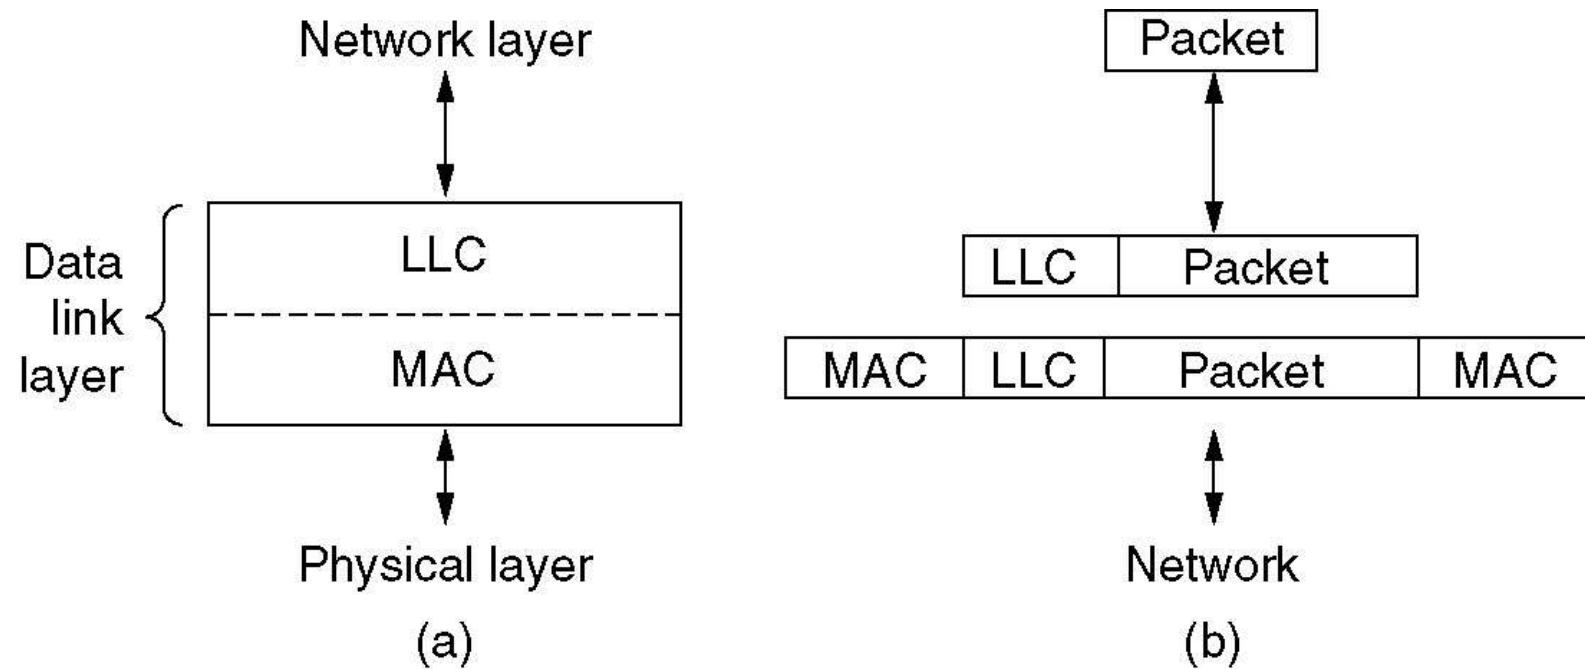
(a) A two-station Ethernet. (b) A multistation Ethernet.

Gigabit Ethernet (2)

Name	Cable	Max. segment	Advantages
1000Base-SX	Fiber optics	550 m	Multimode fiber (50, 62.5 microns)
1000Base-LX	Fiber optics	5000 m	Single (10 μ) or multimode (50, 62.5 μ)
1000Base-CX	2 Pairs of STP	25 m	Shielded twisted pair
1000Base-T	4 Pairs of UTP	100 m	Standard category 5 UTP

Gigabit Ethernet cabling.

IEEE 802.2: Logical Link Control



(a) Position of LLC. (b) Protocol formats.