ADVANCES IN TELECOMMUNICATION NETWORKS SERIES

LOCAL AREA & MULTIPLE ACCESS NETWORKS

Editor: RAYMOND L. PICKHOLTZ

COMPUTER SCIENCE PRESS

LOCAL AREA & MULTIPLE ACCESS NETWORKS

Editor: RAYMOND L. PICKHOLTZ
The George Washington University

COMPUTER SCIENCE PRESS

Printed in the United States of America

Computer Science Press, Inc.
1803 Research Boulevard
Rockville, Maryland 20850

1 2 3 4 5 6 91 90 89 88 87 86

Library of Congress Cataloging-in-Publication Data

Local area and multiple access networks.

(Advances in telecommunication networks)
Includes bibliographies and index.
1. Local area networks (Computer networks)
I. Pickholtz, Raymond L. II. Series.
TK5105.7.L58 1986 004.6'8 86-4218
ISBN 0-88175-143-X
ADVANCES IN TELECOMMUNICATION NETWORKS SERIES
ISSN 0888-2223

Introduction

This book is the outgrowth of an idea proposed by Dr. W. Chou to put together a volume of original papers, written by well-known experts, on the subject of Local Area Networks (LAN). While each paper was to be self-contained, the papers, which are now the chapters of this book, connect through the theme, which was on LANs. The subject of Local Area Networks is highly volatile and the technologies which have emerged will impact the entire telecommunications environment in a most significant way. Indeed, since the first conception of this book several years ago, many things have happened. First, a set of standards (generally designated as IEEE 802 but also partially adopted by other international standards organizations) has emerged. Second, there has been a significant interest in expanding the methods and techniques learned in LAN design to Metropolitan Area Networks and in the integration of voice and data and in interfacing (or replacing) LANs with electronic digital private branch exchanges (PBX). Finally, there has been a major increase in the sophistication of our fundamental and theoretical understanding of the issues of multiple access networks in general, of which LANs are a part. This volume now reflects fairly accurately the current state of affairs and addresses these new developments.

The first chapter by Stallings concerns the establishment of the industry standards of Local Area Networks. The IEEE 802 standard is concerned primarily with the physical and logical link layers of the open systems interconnection (OSI) model, actually has several aspects and implementations. Stallings' chapter gives a readable and authoritative overview of this set of standards which has become synonymous with LANs.

In Chapter 2, Professor Kleinrock writes about an issue is at the core of all multiple-access communications and that is channel efficiency. His paper provides both insight as well as vitality for LAN design.

M. Ferguson's Chapter 3 presents a thorough mathematical analysis of the mean waiting time of token ring systems which is the technique used in 802.5.

In Chapter 4, Chlamtac shows how available programmable LAN VLSI chips can be used to incorporate new user features which can be tailored to user defined applications to enhance the 802.3 standard.

Chou, et al. in Chapter 5 show how the LAN concept in the 802 standard can be extended to Metropolitan Area Networks and give a set of methods for performance evaluation.

In Chapter 6, Nishida, et al. introduce a novel scheme for congestion control in interconnected LANs. Such interconnected LANs may be viewed as a way of

Contents

Chapter 1

A Tutorial on the IEEE 802 Local Network Standard

WILLIAM STALLINGS*

1.1 INTRODUCTION

The key to the development of the LAN market is the availability of a low cost interface. The cost of connecting requirement to a LAN must be much less than the cost of the equipment alone. This requirement, plus the complexity of the LAN protocols, dictates a VLSI solution. However, chip manufacturers will be reluctant to commit the necessary resources unless there is a high-volume market. A LAN standard would assure the needed volume and in addition would enable equipment from a variety of manufacturers to intercommunicate. This was the rationale of the IEEE Project 802 [1], a committee established by the IEEE Computer Society in February of 1980 to prepare local area network standards. The work of the committee has now come to fruition. Several portions of the standard are now working their way through national and international standards organizations, and the remainder of the standard will follow soon. The purpose of this paper is to provide a tutorial on the technical content of the standard.

1.1.1 An Architecture for Local Networks

The task of the IEEE 802 was to specify the means by which devices could communicate over a local area network. The committee characterized its work in this way [2]:

> A Local Network is a data communications system which allows a number of independent devices to communicate with each other. This Standard defines a set of interfaces and protocols for the Local Network.

*President, Comp/Comm Consulting, Great Falls, VA.

> A Local Network is distinguished from other types of data networks in that the communication is usually confined to a moderate size geographic area such as a single office building, a warehouse, or a campus. The network can generally depend on a communications channel of moderate to high data rate which has a consistently low error rate. The network is generally owned and used by a single organization. This is in contrast to long distance networks which interconnect facilities in different parts of the country or are used as a public utility. The Local Network is also different from networks which interconnect devices on a desktop or components within a single piece of equipment.
>
> The objective of the Local Network Standard is to ensure compatibility between equipment made by different manufacturers such that data communications can take place between the devices with a minimum effort on the part of the equipment users or the builders of a system containing the equipment. To accomplish this, the Standard will provide specifications which establish common interfaces and protocols for local area data communications networks.

Two conclusions were quickly reached. First, the task of communication across the local network is complex, and therefore it needs to be broken up into more manageable subtasks. Second, no single technical approach will satisfy all requirements.

The first conclusion, reflected in a "local network reference model," is compared in Figure 1.1 to the better-known open systems interconnection (OSI) model. The model has three layers:

1. *Physical.* This layer is concerned with the nature of the transmission medium, and with the details of device attachment and electrical signaling.
2. *Medium access control.* A local network is characterized by a collection of devices all needing to share a single transmission medium. A means to control access is needed so that only one device attempts to transmit at a time.
3. *Logical link control.* This layer is concerned with establishing, maintaining, and terminating a logical link between devices.

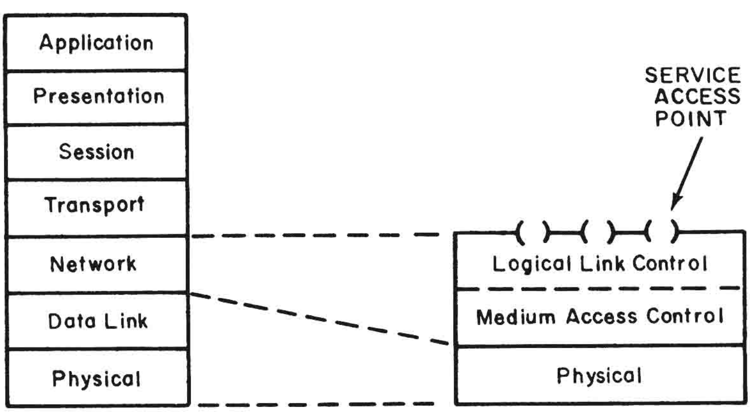

Figure 1.1 IEEE 802 Reference model.

The second conclusion was reluctantly reached when it became apparent that no single standard would satisfy all committee participants. There was support for both ring and bus topologies. Within the bus topology, there was support for two access methods (CSMA/CD and token bus) and two media (baseband and broadband). The response of the committee was to standardize all serious proposals rather than settling on just one. The result is depicted in Figure 1.2.

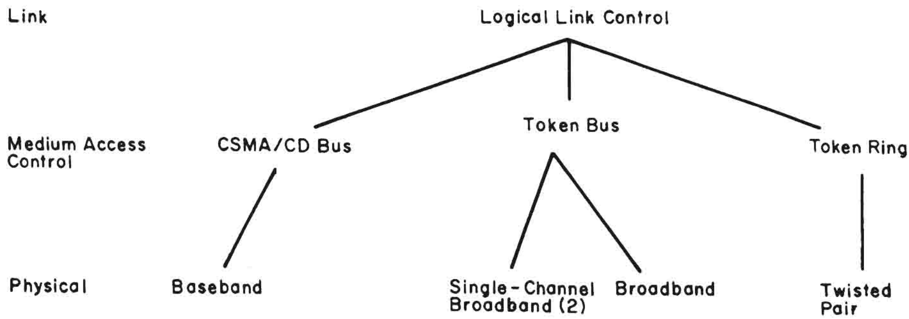

Figure 1.2 Structure of IEEE 802 local network standard.

1.1.2 Status

The work of the IEEE 802 committee is currently organized into the following subcommittees:

1. IEEE 802.1 Higher Layer Interface Standard (HILI).
2. IEEE 802.2 Logical Link Control Standard (LLC).
3. IEEE 802.3 CSMA/CD.
4. IEEE 802.4 Token Bus.
5. IEEE 802.5 Token Ring.
6. IEEE 802.6 Metropolitan Area Network (MAN).

The HILI subcommittee is working on a variety of related issues such as higher layer interfaces, internetworking, addressing, and network management.

Work has been completed on LLC, CSMA/DC, token bus, and token ring for an initial standard, and approved IEEE standards have been adopted for each [3–6]. Work on new options and features continues in each subcommittee.

The work on metropolitan area networks is just beginning to make progress. The subcommittee is attempting to develop a small number of reasonable alternatives for further study.

The acceptance of the IEEE 802 standards has been remarkably widespread. The National Bureau of Standards, which issues Federal Information Processing Standards (FIPS) for U.S. government procurements, has issued a FIPS for CSMA/CD and LLC [7]. The others will probably follow. The International Organization for Standardization (ISO) has decided to adopt the IEEE 802 documents in toto as a draft proposal. This is the first step in the development of international standards. The influential European Computer Manufacturers Association (ECMA), which had been actively drafting its own LAN standards, has now officially deferred to IEEE 802.

1.2 LOGICAL LINK CONTROL (LLC)

1.2.1 Overview

The link layer for LANs bears some resemblance to the more common link layers extant. Like all link layers, the LAN link layer is concerned with the transmission of a frame of data between two stations, with no intermediate switching nodes.

It differs from traditional link layers in three ways:

1. It must support the multiaccess nature of the link (this differs from multidrop in that there is no primary node).
2. It is relieved of some details of link access by the medium access control (MAC) layer.
3. It must provide some network layer (layer 3) functions.

Figure 1.3 will help clarify the requirements for the link layer. We consider two stations or systems that communicate via a LAN (bus or ring). Higher layers (the equivalent of transport and above) provide end-to-end services between the stations. Below the link layer, a medium access control (MAC) layer provides the necessary logic for gaining access to the network for frame transmission and reception.

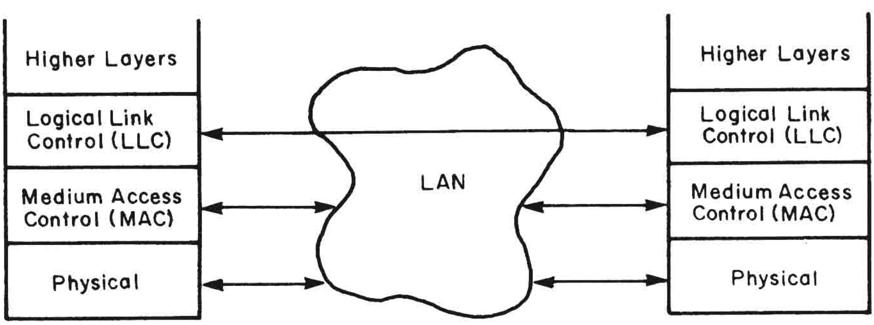

Figure 1.3 LAN communication architecture.

At a minimum, LLC should perform those functions normally associated with the link layer:

1. *Error control.* End-to-end error control and acknowledgement. The link layer should guarantee error-free transmission across the LAN.
2. *Flow control.* End-to-end flow control.
3. *Sequencing.* Frames are delivered in the order in which they are sent.

These functions can be provided in much the same way as for HDLC and other point-to-point link protocols—by the use of sequence numbers.

Because of the lack of intermediate switching nodes, a LAN does not require a separate network layer; rather, the essential layer-3 functions can be incorporated into layer 2:

1. *Datagram.* Some form of connectionless service is needed for efficient support of highly interactive traffic.
2. *Virtual circuit.* A connection-oriented service is also usually needed.
3. *Multiplexing.* Generally, a single physical link attaches a station to a LAN; it should be possible to provide data transfer with multiple end points over that link.

Because there is no need for routing, the above functions are easily provided. The datagram service simply requires the use of source and destination address fields. The station sending the datagram must designate the destination address, so that the frame is delivered properly. The source address must also be indicated so that the recipient knows where the frame came from.

Both the virtual circuit and multiplexing capabilities can be supported with the concept of the service access point (SAP). Figure 1.4 shows three stations attached to a LAN [8]. Each station has an address. Further, the link layer supports multiple SAPs, each with its own address. The link layer provides communication between SAPs. Assume that a process or application X in station A wishes to send a message to a process in station C. X may be a report generator program in minicomputer A. C may be a printer and a simple printer driver. X attaches itself to SAP 1 and requests a connection to station C, SAP 1 (station C may have only one SAP if it is a single printer). Station A's link layer then sends to the LAN a "connection-request" frame which includes the source address (A,1), the destination address (C,1), and some control bits indicating that this is a connection request. The LAN delivers this frame to C, which, if it is free, returns a "connection-accepted" frame. Henceforth, all data from X will be assembled into a frame by A's LLC, which includes source (A,1) and destination (C,1) addresses. Incoming frames addressed to (A,1) will be rejected unless they are from (C,1); these might be acknowledgment frames, for example. Similarly, station C's printer is declared busy and C will only accept frames from (A,1).

Thus, a connection-oriented service is provided. At the same time, process Y could attach to (A,2) and exchange data with (B,1). This is an example of multiplexing. In addition, various other processes in A could use (A,3) to send datagrams to various destinations.

One final function of the link layer should be included to take advantage of the multiple access nature of the LAN:

4. *Multicast, broadcast.* The link layer should provide a service of sending a message to multiple stations or all stations.

With these requirements in mind, we turn to the 802 specification.

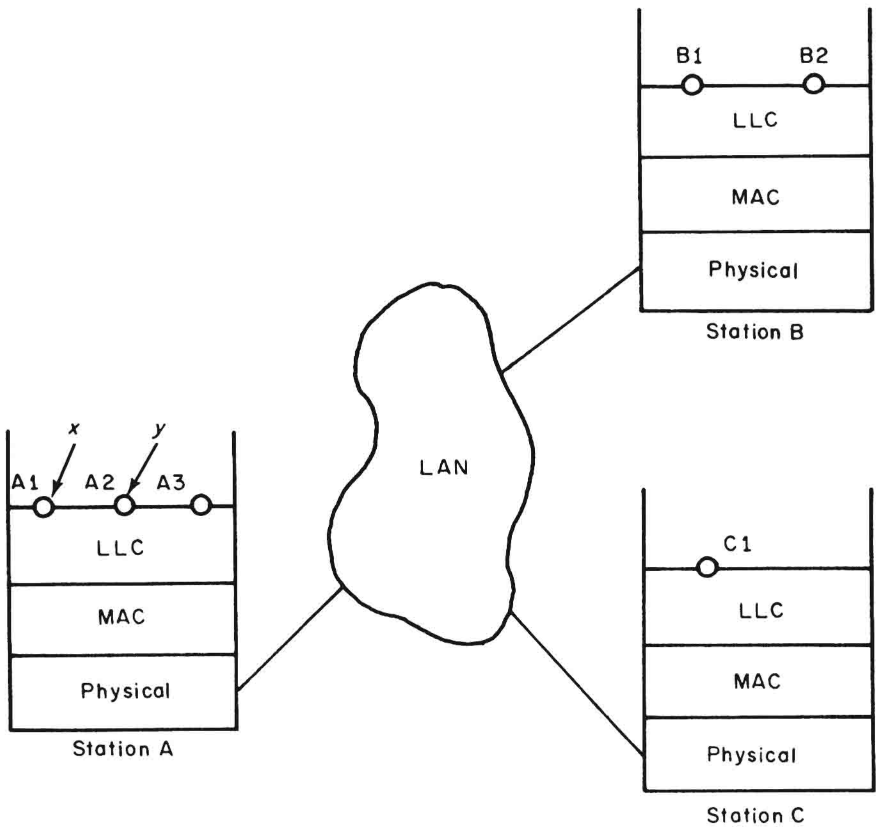

Figure 1.4 LAN link control scenario.

1.2.2 IEEE 802 LLC Specification

The IEEE 802 Logical Link Control (LLC) is a good example of a LAN link control layer. It is well thought out, and offers a variety of services. This section summarizes the features of LLC.

The LLC can be specific in three parts:

1. The interface with the station, specifying the services that LLC (and hence the LAN) provides to the network subscriber.
2. The LLC protocol, specifying the LLC functions.
3. The interface with MAC, specifying the services that LLC requires to perform its function.

A variety of functions were mentioned in the previous section. Not all of these functions are needed in all environments. Accordingly, the 802.2 standard defines two types of data link control operation. The first is a connectionless operation which provides minimum service with minimum protocol complexity.

This type is useful and efficient when higher layers (e.g., network, transport) provide error control, flow control, and sequencing functions. It is also useful when the guaranteed delivery of data is not required. The second type is a connection-oriented operation which provides the functions referred to above using a protocol similar to HDLC. These two types of operations are reflected in the specifications of both the LLC services and the LLC protocol.

LLC Services

The LLC standard specifies two alternative forms of service to higher-layer entitles: unacknowledged connectionless service and connection-oriented service. These services are defined by specifying the service primitives and parameters exchanged between an LLC entity and its users. These are listed in Table 1.1.

The Unacknowledged Connectionless Service is a datagram-style of service that simply allows for sending and receiving LLC frames, with no form of acknowledgment to assure delivery. It supports point-to-point, multipoint, and broadcast addressing, and multiplexing.

The Unacknowledged Connectionless Service provides for only two primitives across the interface between the next highest layer and LLC (not counting management service primitives). L-DATA.request is used to pass a frame to LLC for transmission. L-DATA.indication is used to pass a frame up from LLC upon reception. The local-address and remote-address parameters specify the local and remote LLC users, respectively. These parameters are logically equivalent to the

Table 1.1
LLC Primitives

Unacknowledged connectionless service
L-DATA.request
L-DATA.indication
Connection-oriented service
L-DATA-CONNECT.request
L-DATA-CONNECT.indication
L-DATA-CONNECT.confirm
L-CONNECT.request
L-CONNECT.indication
L-CONNECT.confirm
L-DISCONNECT.request
L-DISCONNECT.indication
L-DISCONNECT.confirm
L-RESET.request
L-RESET.indication
L-RESET.confirm
L-CONNECTION-FLOWCONTROL.request
L-CONNECTION-FLOWCONTROL.indication

LLC service access point (SAP) and the MAC address. This point is elaborated below. The 1-sdu parameter specifies the LLC service data unit; this is the block of data exchanged between LLC and its user. The service-class parameter specifies the desired priority. This parameter is passed through the LLC entity to the MAC entity, which has the responsibility of implementing a priority mechanism. Token bus and token ring are capable of this, but the 802.3 CSMA/CD specification is not.

The Connection-Oriented Service provides a virtual-circuit-style connection between service access points. It provides a means by which a user can request or be notified of the establishment or termination of a data-link connection. It also provides flow control, sequencing, and error recovery. It supports point-to-point addressing and multiplexing.

The Connection-Oriented Service includes L-DATA-CONNECT.request and L-DATA-CONNECT.indication, with meanings analogous to those above, plus L-DATA-CONNECT.confirm, which conveys the result (acknowledged, failure) of the previous associated L-DATA-CONNECT.request via the status parameter. If the transfer was successful, this primitive indicates that the remote LLC entity received the 1-sdu and positively acknowledged it via the LLC protocol.

The next six primitives listed in Table 1.1 are used to establish and subsequently tear down a logical connection between SAPs. The status parameter indicates whether or not the request was successful and, if not, the reason for failure. The reason parameter specifies the reason for an LLC-requested disconnection. The L-RESET primitives are used to reset a logical connection to an initial state. Sequence numbers are reset, and the connection is reinitialized.

The two flow control primitives regulate the flow of data across the SAP between LLC and the LLC user. The flow can be controlled in either direction. This is a local flow control mechanism which specifies the amount of data that may be passed across the SAP.

LLC Protocol

The LLC protocol is modeled after the HDLC balanced mode, and has similar formats and functions. These are summarized briefly in this section. The reader should be able to see how this protocol supports the LLC services defined above.

The LLC frame consists of four fields, as shown in Figure 1.5a. Unlike most link control formats, LLC requires both source and destination addresses to identify the two peer communicating entities. The source and destination are uniquely identified by a (node, SAP) pair. However, the node addresses are also used by MAC and are included in the outer MAC frame.

The format of the control field (Figure 1.5b) is identical to that of HDLC and the functioning is the same, with three exceptions:

1. LLC only makes use of the asynchronous balanced mode of operation, and does not employ HDLC's normal response mode or asynchronous response mode. This mode is used to support connection-oriented service. The set

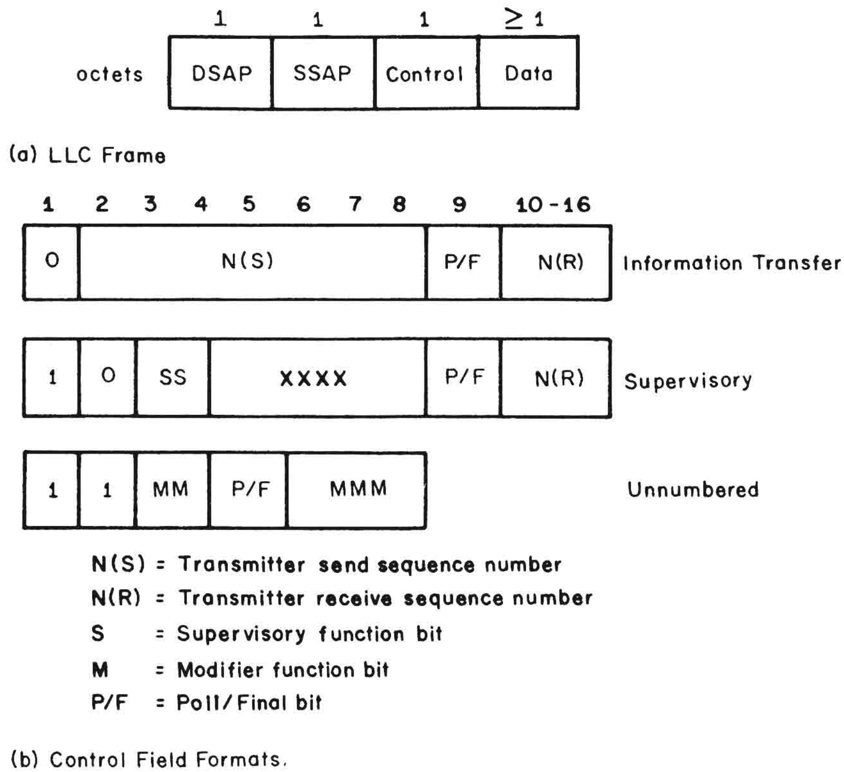

Figure 1.5 Logical link control format.

asynchronous balanced mode (SABM) command is used to establish a connection, and disconnect (DISC) is used to terminate the connection.

2. LLC supports a connectionless (datagram) service by using the unnumbered information (UI) frame.
3. LLC permits multiplexing by the use of SAPs.

A brief summary follows.

As with HDLC, three frame formats are defined for LLC: information transfer, supervisory, and unnumbered. Their use depends on the type of operation employed. The types are Type 1 (connectionless) and Type 2 (connection-oriented).

With Type 1 Operation, protocol data units (PDUs) are exchanged between LLC entities without the need to establish a data link connection. There is no acknowledgment, flow control, or error control. This type of operation supports the Unacknowledged Connectionless Service.

Three unnumbered frame formats are used. The UI (unnumbered information) frame is used to send a connectionless data frame, containing data from an LLC user. The XID (exchange identification) frame is used to convey station class

(Class I—only connectionless; Class II—connection-oriented and connectionless; plus window size). The TEST (test) frame is used to request a TEST frame in response, to test the LLC-to-LLC path.

With Type 2 Operation, a data link connection is established between two LLC entities prior to data exchange. This type of operation supports Connection-Oriented Service and uses all three frame formats.

The information transfer frames are used to send data (as opposed to control information). N(S) and N(R) are frame sequence numbers that support error control and flow control. A station sending a sequence of frames will number them, modulo 128, and place the number in N(S). N(R) is a piggybacked acknowledgment. It enables the sending station to indicate which number frame it expects to receive next. These numbers support flow control since, after sending 127 frames without an acknowledgment, a station can send no more. The numbers support error control, as explained below.

The supervisory frame is used for acknowledgment and flow control. The 2-bit SS field is used to indicate one of three commands: Receive Ready (RR), Receive Not Ready (RNR), and Reject (REJ). RR is used to acknowledge the last frame received by indicating in N(R) the next frame expected. This frame is used when there is no reverse traffic to carry a piggybacked acknowledgment. RNR acknowledges a frame, as with RR, but also asks the transmitting station to suspend transmission. When the receiving station is again ready, it sends a RR frame. REJ is used to indicate that the frame with number N(R) is rejected and that it and any subsequently transmitted frames must be sent again.

Unnumbered frames are used for control purposes in Type 2 Operation. The bit pattern defines one of the following commands:

1. SABME (set asynchronous balanced mode extended): Used by an LLC entity to request logical connection with another LLC entity.
2. DISC (disconnect): Used to terminate a logical connection; the sending station is announcing that it is suspending operations. Any outstanding unacknowledged I PDUs remain unacknowledged.

The foregoing frames are all commands, initiated by a station at will. The following frames are responses:

1. UA (unnumbered acknowledgment): Used to acknowledge SABME and DISC commands.
2. DM (disconnect mode): Used to respond to a frame in order to indicate that the station's LLC is logically disconnected.
3. FRMR (frame reject). Used to indicate that an improper frame has arrived—one that somehow violates the protocol.

The P/F bit is used to indicate that a response is requested to a command frame.

LLC-MAC Interface

The IEEE 802 LLC is intended to operate with any of the three MAC protocols (CSMA/CD, token bus, token ring). A single logical interface to any of the MAC layers is defined. The 802 standard does not define an explicit interface, but provides a "model." The basic primitives are:

1. MA-DATA.request. To request transfer of an LLC frame from local LLC to destination LLC. This includes information transfer, supervisory, and unnumbered frames.
2. MA-DATA.confirm. Response from local MAC layer to LLC's MA-DATA.request. It indicates the success or failure of the request, but has only *local* significance (i.e., it is not an end-to-end acknowledgment).
3. MA-DATA.indicate. To transfer incoming LLC frame from local MAC to local LLC.

1.3 CSMA/CD

1.3.1 Overview

The simplest form of medium access control adopted for IEEE 802 is carrier sense multiple access with collision detection (CSMA/CD). The original baseband version of this technique was developed and patented by Xerox [9] as part of its Ethernet local network [10,11]. The original broadband version was developed and patented by MITRE [12] as part of its MITREnet local network [13].

We begin by looking at a similar but simpler protocol known as Carrier Sense Multiple Access (CSMA). With this scheme, a station wishing to transmit first listens to the medium to determine if another transmission is in progress. If the medium is idle, the station may transmit. It may happen that two or more stations attempt to transmit at about the same time. If this happens, there will be a collision; the data from both transmissions will be garbled and not be received successfully. To account for this, a station waits a reasonable amount of time after transmitting for an acknowledgment. If there is no acknowledgment, the station assumes that a collision has occurred and retransmits.

With CSMA, an algorithm is needed to specify what a station should do if the medium is found to be busy. Three approaches are depicted in Figure 1.6. One algorithm is *non-persistent* CSMA. A station wishing to transmit listens to the medium and obeys the following rules:

1. If the medium is idle, transmit.
2. If the medium is busy, wait an amount of time drawn from a probability distribution (the retransmission delay) and repeat step 1.

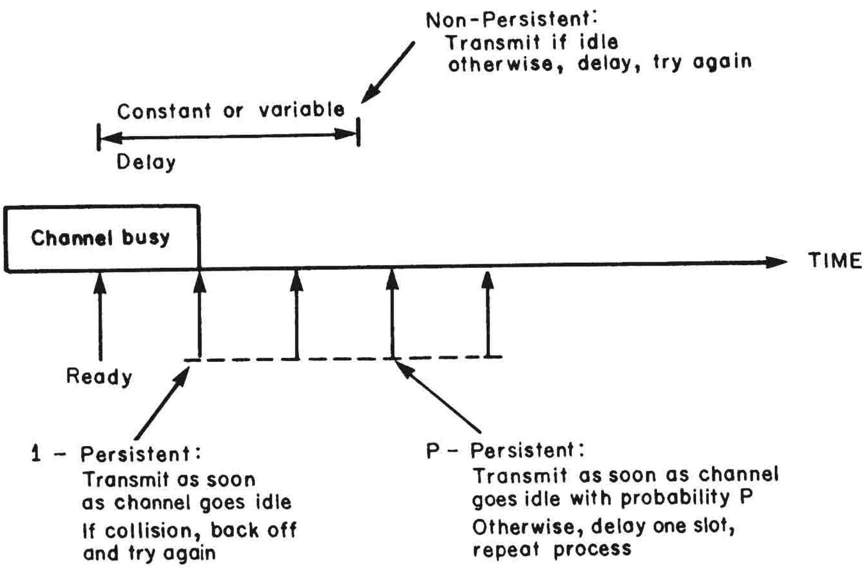

Figure 1.6 CSMA persistence and back-off.

The use of random retransmission times reduces the probability of collisions. The drawback is that even if several stations have a frame to send, there is likely to be some wasted idle time following a prior transmission.

To avoid channel idle time, the *1-persistent protocol* can be used. A station wishing to transmit listens to the medium and obeys the following rules:

1. If the medium is idle, transmit.
2. If the medium is busy, continue to listen until the channel is sensed idle, then transmit immediately.
3. If there is a collision (determined by a lack of acknowledgment), wait a random amount of time and repeat step 1.

Whereas non-persistent stations are deferential, 1-persistent stations are selfish. If two or more stations are waiting to transmit, a collision is guaranteed. Things only get sorted out after the collision.

A compromise that attempts to reduce collisions, like non-persistent, and to reduce idle time, like 1-persistent, is *p-persistent*. The rules are:

1. If the medium is idle, transmit with probability p, and delay one time unit with probability $(1-p)$. The time unit is typically equal to the maximum propagation delay.
2. If the medium is busy, continue to listen until the channel is idle and repeat step 1.
3. If transmission is delayed one time unit, repeat step 1.