ANSI/IEEE Std 802.4-1990
(Revision of ANSI/IEEE 802.4-1985)

(Adopted by ISO/IEC and redesignated as
ISO/IEC 8802-4: 1990)

Information processing systems—

Local area networks—

Part 4:
Token-passing bus access method and physical layer specifications

Adopted by the ISO/IEC and redesignated as ISO/IEC 8802-4: 1990

Sponsor

LAN/MAN Standards Committee
of the
IEEE Computer Society

Abstract: This Local Area Network (LAN) standard, ISO/IEC 8802-4 : 1990 (ANSI/IEEE Std 802.4-1990), deals with all elements of the token-passing bus access method and its associated physical signaling and media technologies. To facilitate interconnection of stations by way of a LAN using the token-passing bus access method, this standard specifies the characteristics of the transmission medium; the signaling method used; the frame formats transmitted; the actions of a station upon receipt of a frame; the services provided at the conceptual interface between the Medium Access Control (MAC) sublayer and the Logical Link Control (LLC) sublayer; and the actions, entities, and values used by management. There are four medium characteristics and signaling methods: 5 and 10 Mb/s phase-coherent FSK; 1, 5, and 10 Mb/s broadband; 10 and 20 Mb/s fiber optic; and 1 Mb/s phase-continuous FSK.

First Printing

ISBN 1-55937-027-0
Library of Congress Catalog Card Number 90-082603

The Institute of Electrical and Electronics Engineers, Inc.
345 East 47th Street, New York, NY 10017-2394, USA

August 17, 1990 *SH12948*

Foreword to IEEE Std 802.4-1990
(Revision of IEEE Std 802.4-1985)

(This Foreword is not a part of ISO/IEC 8802-4 : 1990 or of IEEE Std 802.4-1990.)

This standard is part of a family of standards for Local Area Networks (LANs). The relationship between this standard and other members of the family is shown below. (The numbers in the figure refer to IEEE Standard numbers.)

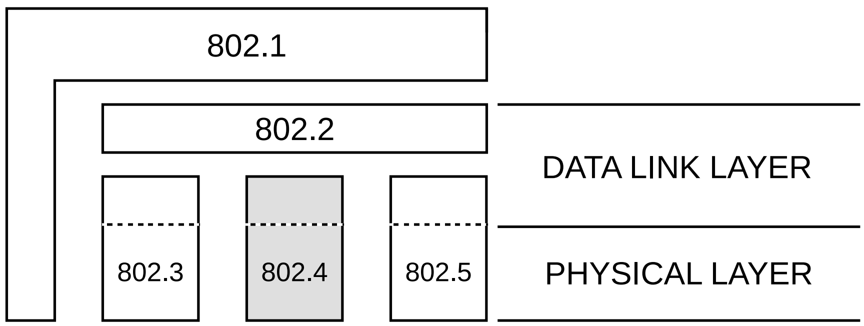

This family of standards deals with the physical and data link layers as defined by the ISO Open Systems Interconnection Basic Reference Model (ISO 7498 : 1984). The access standards define three types of medium access technologies and associated physical media, each appropriate for particular applications or system objectives. The standards defining these technologies are

(1) IEEE Std 802.3-1988 [ISO 8802-3], a bus utilizing CSMA/CD as the access method,
(2) IEEE Std 802.4-1990 [ISO/IEC 8802-4], a bus utilizing token passing as the access method,
(3) IEEE Std 802.5-1989, a ring utilizing token passing as the access method.

Other access methods (for example, metropolitan area networks and integrated voice-data networks) are under investigation.

IEEE Std 802.2-1989 [ISO 8802-2], the Logical Link Control standard, is used in conjunction with the medium access standards

IEEE 802.1[1] (a series of related standards) describes the relationship among the family of 802 standards and their relationship to the ISO Open Systems Interconnection Basic Reference Model in more detail. IEEE 802.1 will also contain networking management standards and information on internetworking.

The reader of this standard is urged to become familiar with the complete family of standards.

[1]IEEE Std 802.1A-1990, Overview and Architecture of Network Standards; IEEE Std 802.1D-1990, MAC (Media Access Control) Bridges; and IEEE Std 802.1E-1990, System Load Protocol have been approved as IEEE Standards, but are not yet published. Other projects in the 802.1 series are currently under development.

Summary of Changes

This standard is a major revision of IEEE Std 802.4-1985 (ISO/DIS 8802/4). This revision incorporates the results of over four years of work by the IEEE 802.4 Working Group and by other organizations.

This section of the Foreword summarizes the changes for the reader's convenience.

The management sections, Sections 3 and 9, have been completely rewritten. These revised sections are now in conformance with the work of the IEEE 802 Committee and in alignment, to the extent possible, with other groups working on management standards.

The access control machine (ACM), the heart of the Medium Access Control (MAC) protocol, has been revised. Errors and ambiguities that existed in the previous version were corrected, and enhancements were made to improve error recovery. The intention in revising the ACM was to retain complete interoperability with the ACM described by the previous standard. We believe that intention has been realized.

The name of the request_with_response option has been changed from "immediate response" option to avoid conflicts with other uses of the phrase immediate response (see 6.1.1 and 6.6.2).

Additional MAC Capabilities (6.7) has been added to give implementors suggestions for providing extra MAC features and to give guidance to conformance testers.

The specification of an interface within the Physical Layer for a separate modem has been included as Section 10.

Section 11 has been reserved to retain compatible section numbers and for future additions.

The phase-coherent FSK Physical Layer and medium, described in Sections 12 and 13, underwent revision. Based on implementation experience, the working group found it necessary to make substantive technical changes to the standard. The transmit and receive levels were changed, the preamble pattern and length were changed, and a receiver blanking specification was added. The result of these changes is that the implementations of this revised standard will not interoperate with implementations of the previous standard.

The broadband Physical Layer and medium, described in Sections 14 and 15, underwent minor revisions to clarify the specification and to provide guidance for conformance testing.

A fiber optic Physical Layer and medium specification has been added as Sections 16 and 17.

The 1 Mb/s phase-continuous FSK Physical Layer and medium specification was moved from Sections 10 and 11 to Sections 18 and 19, to avoid renumbering the more commonly referenced Physical Layer and media sections. The coding of the end delimiter coding pattern was changed to increase the reliability (Hamming distance) of the protocol.

This standard was submitted to ISO/IEC JTC1 for consideration as a revision and addendum to the previous edition of the token bus LAN standard, IEEE Std 802.4-1985 (ISO/DIS 8802/4). To facilitate processing of that document, portions that were not appropriate for an international standard were prefaced with a note enclosed in braces{…}.

These same portions contained within the current edition are not a part of the International Standard and are stated as such. These portions are peculiar to the IEEE version of this standard and consist of areas relating to

(1) References to national standards
(2) Recommended frequency allocations for North American CATV systems (see 14.8.4)
(3) Recommendations and guidelines related to safety concerns

This standard contains state-of-the-art material. The area covered by this standard is undergoing evolution. Revisions to this standard may occur either to clarify existing material, to correct possible errors, or to incorporate new, related material.

Readers wishing to know the state of revisions should contact the

Secretary, IEEE Standards Board
445 Hoes Lane, P.O. Box 1331
Piscataway, NJ 08855-1331, USA

The following were voting members of the IEEE 802.4 Token-Passing Bus Access Method Working Group at the time of approval of this standard:

Paul S. Eastman, *Chair* **Louis F. Wojnaroski**, *Vice Chair*
Rhonda Alexis-Dirvin, *Secretary*

Ernest E. Bergmann
Clyde A. Boenke
Shalom Bresticker
Charles Brill
Mike Bukowski
Mike Bush
Hsin-Hwai Chen
Jade Y. Chien
Robert S. Crowder
C. David Daly
Sanjay Dhawan
H. N. "Hank" Dorris
Robert H. Douglas
Daniel Dove
Richard B. Formeister
Anatoly V. Fridland
Darrell Furlong
T. R. (Tom) Gandhi
Maris Graube
David J. Greenstein
Gerhard Hammer
Richard F. Hunter
Dittmar Janetzky
Lawrence R. Johnson
Walter Kammerer
Bradford S. Kellar
Robert W. Kilgore
William H. Kind
Michael T. Klein
Simon Korotwitz
Leonid Koshevoy
Detlev Leisengang
Daniel N. LeBlanc
Bruce A. Loyer
Wayne A. Mack
Charles H. March
Gunther J. Martin
Michael A. Masleid
Stephen G. McChrystal
Dale D. Murry
David L. Nicholson
Toshio Ogawa
Ilan Pardo
Michael J. Perkins
Thomas L. Phinney
Fred Rhine
David W. Riley
Ken Ross
Chandos A. Rypinski
Toshio Saito
Walter Schreuer
Rich Seifert
Yehuda Shvager
Daniel P. Stokesberry
Orest Lev Storoshchuk
Catherine F. Summers
Chuck Thurwachter
Lourens Van Der Jagt
Earl J. Whitaker
Roger Wilmarth
Yong W. Yi

Special thanks to the following task group chairmen whose task groups contributed to this revision:

802.4A, Broadband	—	**Donald J. Loughry**
802.4B, Carrier Band	—	**Michael T. Klein**
802.4C, MAC Management	—	**Robert H. Douglas**
802.4D, Physical Management	—	**Michael J. Perkins**
802.4F, MAC Revision	—	**Louis F. Wojnaroski**
and past Chair	—	**Richard M. Collins**
802.4G, DTE-DCE Interface	—	**Clyde A. Boenke**
802.4H, Fiber Optic Media	—	**Robert S. Crowder**
802.4J, Conformance Testing	—	**Paul S. Eastman**
and past Chair	—	**Robert L. Husak** (deceased)

Special thanks to our past chairman, 1983–87, under whose guidance we developed the current draft and reached committee consensus:

Robert H. Douglas, *past Chair*

The final conditions for approval of this standard were met on February 27, 1990. This standard was conditionally approved by the IEEE Standards Board on February 15, 1990, with the following membership:

Contents

SECTION PAGE

1. Introduction and Overview ... 17
   - 1.1 Scope ... 18
   - 1.2 Definitions ... 18
   - 1.3 References ... 18
   - 1.4 Conformance ... 20
   - 1.5 Overview of the Token Method ... 21
   - 1.6 MAC Sublayer Internal Structure ... 22
   - 1.7 PLE and Medium ... 24
   - 1.8 Access Method Characteristics ... 28
   - 1.9 Standard Organization ... 29
2. LLC-MAC Interface Service Specification ... 31
   - 2.1 Overview of the LLC-MAC Service ... 31
   - 2.2 Detailed Interactions with the LLC Entity ... 32
3. MAC Sublayer Management ... 35
   - 3.1 Overview ... 36
   - 3.2 MAC Management Facilities ... 36
4. Frame Formats ... 43
   - 4.1 Frame Components ... 44
   - 4.2 Enumeration of Frame Types ... 50
   - 4.3 Appendix—Recommendation for a Hierarchical Structure for Locally Administered Addresses ... 52
5. Elements of MAC Sublayer Operation ... 55
   - 5.1 Basic Operation ... 56
   - 5.2 Access Control Machine (ACM) States ... 65
   - 5.3 Interface Machine (IFM) Description ... 70
   - 5.4 Receive Machine (RxM) Description ... 71
   - 5.5 Transmit Machine (TxM) Description ... 76
   - 5.6 Regenerative Repeater Machine (RRM) Description ... 76
6. MAC Sublayer Definitions and Requirements ... 79
   - 6.1 MAC Definitions ... 79
   - 6.2 Transmission Order ... 80
   - 6.3 Delay Labeling ... 82
   - 6.4 Miscellaneous Requirements ... 82
   - 6.5 Use of Address Bits in Contention Algorithms ... 84
   - 6.6 Options within MAC Sublayer ... 85
   - 6.7 Additional MAC Capabilities ... 86
   - 6.8 Delegation of Right to Transmit ... 88
7. Access Control Machine (ACM) Description ... 89
   - 7.1 Variables and Functions ... 89
   - 7.2 Access Control Machine (ACM) Formal Description ... 99

8. MAC Sublayer-Physical Layer Interface Service Specification 135
8.1 Overview of the LAN Physical Layer Service 135
8.2 Detailed Specifications 136

9. Physical Layer Entity (PLE) Management 141
9.1 Overview 142
9.2 Physical Management Facilities 142
9.3 Additional Management 146

10. Exposed DTE-DCE Interface 147
10.1 Overview of DTE-DCE Interface 148
10.2 PHY-UNITDATA Request and Indication 148
10.3 DCE Management 151
10.4 Electrical Characteristics 158
10.5 Mechanical Characteristics 162
10.6 Appendix—Clarification of Management Mode 166

11. RESERVED 169

12. Single-Channel Phase-Coherent-FSK Bus Physical Layer Entity (PLE) 171
12.1 Nomenclature 172
12.2 Object 173
12.3 Compatibility Considerations 173
12.4 Medium Overview 173
12.5 PLE Overview 173
12.6 Application of Management 175
12.7 Functional, Electrical, and Mechanical Specifications 175
12.8 Environmental Specifications 183
12.9 Labeling 184

13. Single-Channel Phase-Coherent-FSK Bus Medium 185
13.1 Nomenclature 186
13.2 Object 187
13.3 Compatibility Considerations 187
13.4 Overview 188
13.5 Functional, Electrical, and Mechanical Specifications 188
13.6 Environmental Specifications 191
13.7 Transmission_path_delay Considerations 192
13.8 Documentation 192
13.9 Network Sizing 192
13.10 Appendix—Guidelines for Configuring the Medium 193

14. Broadband Bus Physical Layer Entity (PLE) 197
14.1 Nomenclature 198
14.2 Object 199
14.3 Compatibility Considerations 199
14.4 Operational Overview of Single-Cable Medium 200
14.5 Operational Overview of a Dual-Cable Medium 200

14.6 General Overview 201
14.7 Application of Management 203
14.8 Functional, Electrical, and Mechanical Specifications 203
14.9 Environmental Specifications 218
14.10 Labeling 218
14.11 Appendix—Provisions for Two MAC-Symbol/PHY-Symbol Signaling 219
14.12 Appendix—Detailed Scrambling and Descrambling Process 223

15. Broadband Bus Medium 225
15.1 Nomenclature 226
15.2 Object 228
15.3 Compability Considerations 228
15.4 Overview 229
15.5 Functional, Electrical, and Mechanical Specifications 230
15.6 Environmental Specifications 233
15.7 Transmission_path_delay Considerations 233
15.8 Documentation 234
15.9 Network Sizing 234

16. Fiber Optic Physical Layer Entity (PLE) 235
16.1 Nomenclature 236
16.2 Object 238
16.3 Compatibility Considerations 238
16.4 Operational Overview 239
16.5 Physical Layer Overview 239
16.6 Application of Management 240
16.7 Functional, Optical, Electrical, and Mechanical Specifications 240
16.8 Environmental Specifications 246
16.9 Labeling 246
16.10 Appendix—Alternative Fiber Optic Medium 247
16.11 Appendix—Comparison of Parameter Values of Standard and Alternative Fiber Optic Medium 248

17. Fiber Optic Medium 249
17.1 Nomenclature 250
17.2 Object 252
17.3 Compatibility Considerations 252
17.4 Overview 253
17.5 Functional, Optical, and Mechanical Specifications 253
17.6 Safety Requirements 255
17.7 Transmission_path_delay Considerations 255
17.8 Documentation 256
17.9 Appendix—Alternative Fiber Optic Medium 256
17.10 Appendix—Comparison of Parameter Values of Standard and Alternative Fiber Optic Medium 257
17.11 Appendix—Network Size and Configuration 257

SECTION PAGE

18. Single-Channel Phase-Continuous-FSK Bus Physical Layer Entity (PLE) 259
18.1 Nomenclature 260
18.2 Object 261
18.3 Compatibility Considerations 261
18.4 Medium Overview 262
18.5 Physical Layer Overview 262
18.6 Application of Management 263
18.7 Functional, Electrical, and Mechanical Specifications 263
18.8 Environmental Specifications 267
18.9 Labeling 268

19. Single-Channel Phase-Continuous-FSK Bus Medium 269
19.1 Nomenclature 270
19.2 Object 271
19.3 Compatibility Considerations 271
19.4 General Overview 272
19.5 Functional, Electrical, and Mechanical Specifications 272
19.6 Environmental Specifications 274
19.7 Transmission_path_delay Considerations 274
19.8 Documentation 275
19.9 Network Sizing 275
19.10 Appendix—Guidelines for Configuring the Medium 275

FIGURES

Fig 1-1 Relationship of Adjacent Protocol Layers 17
Fig 1-2 Logical Ring on Physical Bus 21
Fig 1-3 MAC Sublayer Functional Partitioning 22
Fig 2-1 Relation to LAN Model 31
Fig 3-1 Relation to LAN Model 35
Fig 5-1 Relation to LAN Model 55
Fig 5-2 Token Rotation Time "Priority" Example 63
Fig 5-3 MAC Finite State Machine Diagram 65
Fig 5-4 Receive Machine 72
Fig 5-5 sil/act Detector Finite State Machine 73
Fig 5-6 SD Detector Finite State Machine 73
Fig 5-7 ED Detector Finite State Machine 74
Fig 5-8 NB/BQ Finite State Machine 75
Fig 5-9 Bus Repeater 77
Fig 6-1 MAC Protocol_data_unit Transmission Order 81
Fig 6-2 Logical Token-Passing Ring 83
Fig 8-1 Relation to LAN Model 135
Fig 9-1 Relation to LAN Model 141
Fig 10-1 Relation to LAN Model 147
Fig 10-2 Single-ended Interface Connection 159
Fig 10-3 Indication Channel Timing 160

SECTION PAGE

Fig 10-4 Request Channel Timing Using TXCLK 160
Fig 10-5 Typical Driver—Differential Interface 161
Fig 10-6 Typical Receivers—Differential Interface 161
Fig 10-7 Request Channel Timing Using RETURNCLK 162
Fig 10-8 Typical Derivation of RETURNCLK 162
Fig 10-9 Cable Detection Circuits 163
Fig 10-10 Single-ended Interface Connector Receptacle 163
Fig 10-11 Typical Station Management 166
Fig 10-12 Writing to a Register 167
Fig 10-13 Typical Modem Failure Sequence 168
Fig 12-1 Relation to LAN Model 171
Fig 12-2 Physical Hardware Partitioning 172
Fig 12-3 MAC-Symbol Encodings 176
Fig 12-4 Eye Patterns 178
Fig 12-5 Receiver Blanking 181
Fig 13-1 Relation to LAN Model 185
Fig 13-2 Physical Hardware Partitioning 186
Fig 13-3 Limits for Percent PDD 190
Fig 14-1 Relation to LAN Model 197
Fig 14-2 Physical Hardware Partitioning 198
Fig 14-3 Limits for Amplitude Distortion 211
Fig 14-4 Transmit Spectrum Mask 212
Fig 14-5 $1 + X^{-6} + X^{-7}$ Scrambler 223
Fig 14-6 Exemplary Scrambler/Descrambler Circuitry 223
Fig 15-1 Relation to LAN Model 225
Fig 15-2 Physical Hardware Partitioning 226
Fig 16-1 Relation to LAN Model 235
Fig 16-2 Physical Hardware Partitioning 236
Fig 16-3 Optical Wave Shape Template 242
Fig 17-1 Relation to LAN Model 249
Fig 17-2 Physical Hardware Partitioning 250
Fig 17-3 *n* by *m* Splitter/Combiner Primitive 253
Fig 17-4 Star Topology 258
Fig 18-1 Relation to LAN Model 259
Fig 18-2 Physical Hardware Partitioning 260
Fig 19-1 Relation to LAN Model 269
Fig 19-2 Physical Hardware Partitioning 270

TABLES

Table 10-1 Request Channel Encoding—MAC Mode 150
Table 10-2 Indication Channel Encoding—MAC Mode 151
Table 10-3 Request Channel Encoding—Management Mode 154
Table 10-4 Indication Channel Encoding—Management Mode 154
Table 10-5 Single-ended Connector Pinout 164
Table 10-6 Differential Connector Pinout 165
Table 12-1 Data Rate Versus Signaling Frequencies 177

FIGURE PAGE

Table 14-1 Relative Power Level Reporting 203
Table 14-2 Baseband Pulse Coding Rules 206
Table 14-3 Usual North American Mid-split Channels — Nomenclature and Pairing 208
Table 14-4 Required Transmit Level 209
Table 14-5 Breakpoints of Amplitude Distortion Limits 210
Table 14-6 Required Noise Floor and In-band Signal Power 214
Table 14-7 Minimum Length of *Pad_idle* Preamble 214
Table 16-1 Data Rate Versus Media Transmission Rate 241
Table 16-2 Alternate Test Fiber Parameter Values 247
Table 16-3 Parameter Values for Standard and Alternate Test Fiber 248
Table 17-1 Alternate Test Fiber Parameter Values 256
Table 17-2 Parameter Values for Standard and Alternate Test Fiber 257
Table 17-3 Passive Star Power Loss Budget 258

APPENDIX

A A Model Used for the Service Specification 279
A1. Service Hierarchy 279
A2. N-Layer Interface 279
A3. Specification of Services 280
A4. Classification of N-Layer Service Primitives 280
A5. Interaction Behavior 280

APPENDIX FIGURES

Fig A1 Service Hierarchy Relationships 279
Fig A2 Service Primitive Interactions 280

Information processing systems—
Local area networks—

Part 4:
Token-passing bus access method and physical layer specifications

1. Introduction and Overview

This portion of the Local Area Network (LAN) standards deals with all elements of the token-passing bus access method and its associated physical signaling and media technologies. The access function coordinates the use of the shared medium among the attached stations and has the relationship to the other protocol functions given by Fig 1-1.

Fig 1-1
Relationship of Adjacent Protocol Layers

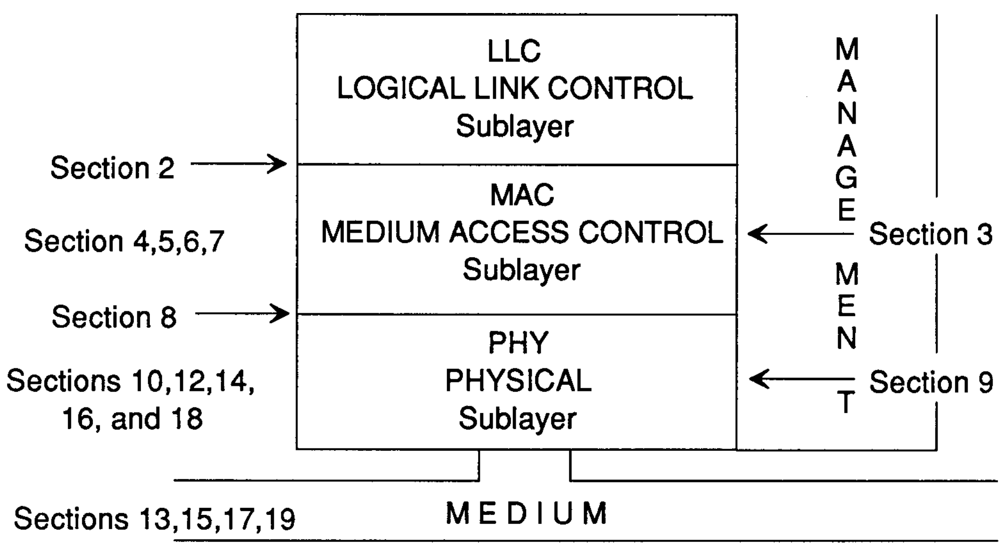

1.1 Scope. For the purpose of compatible interconnection of stations by way of a LAN using the token-passing bus access method this standard

(1) Specifies the electrical and/or optical and physical characteristics of the transmission medium,
(2) Specifies the electrical or optical signaling method used,
(3) Specifies the frame formats transmitted,
(4) Specifies the actions of a station upon receipt of a frame,
(5) Specifies the services provided at the conceptual interface between the Medium Access Control (MAC) sublayer and the Logical Link Control (LLC) sublayer above it,
(6) Specifies the actions, entities, and values used to manage the MAC sublayer and Physical Layer Entity (PLE).

Within this standard the operation of a station is specified in terms of the layered model shown in Fig 1-1 and described in ISO 7498 : 1984 [13].[1] Additional information on the 802 family of LAN standards can be found in IEEE Std 802-1990 [9].[2]

Figure 1-1 also shows which sections of the standard specify interfaces between layers and which sections specify the operation of the layers themselves.

Recommendations for frequency allocations for community antenna television (CATV) systems are subjects for national standardization and are outside the scope of this International Standard.

1.2 Definitions. The definitions used in this standard are consistent with ISO 2382-9 : 1984 [12]. A more specific information processing vocabulary pertaining to LAN systems is in development as WP 2382/25.

1.3 References.[3] When the following standards referred to in this standard are superseded by an approved revision, the latest revision shall apply. References to IEEE and ISA standards are not a part of this International Standard.

[1] ANSI/ISA-S72.01-1986, PROWAY-LAN Industrial Data Highway.[4]

[2] CCITT Recommendation X.150, Principles of Maintenance Testing for Public Data Networks Using Data Terminal Equipment (DTE) Data Circuit-Terminating (DCE) Test Loops. In Vol. VIII.3 of the *CCITT Blue Books*—Data communication network transmission, signaling and switching, network aspects,

[1]The numbers in brackets correspond to those of the references listed in 1.3.

[2]IEEE Std 802-1990, Overview and Architecture of Network Standards; IEEE Std 802.1D-1990, MAC (Media Access Control) Bridges; and IEEE Std 802.1E-1990, System Load Protocol have been approved as IEEE Standards, but are not yet published. Other projects in the 802.1 series are currently under development.

[3]See Index to locate citations in the text of the following publications.

[4]ISA documents are available from the Instrument Society of America, Standards Dept., 67 Alexandria Drive, P.O. Box 12277, Research Triangle Park, NC 27709.

maintenance and administrative arrangements. Geneva: International Telecommunications Union, 1989.[5]

[3] CISPR Publication 22 (1985), Limits and Methods of Measurement of Radio Interference Characteristics of Information Technology Equipment.[6]

[4] IEC Publication 169-8 (1978), Radio Frequency Connectors, Part 8: RF Coaxial Connectors with Inner Diameter of Outer Conductor 6.5 mm (0.256 in) with Bayonet Lock and Characteristic Impedance of 50 Ω (type BNC).

[5] IEC Publication 807-2 (1985), Detailed Specification for a Range of Connectors and Round Contacts (Fixed Solder Contact Types).

[6] IEC Publication 825 (1984), Radiation Safety of Laser Products, Equipment Classification, Requirements and User's Guide.

[7] IEC Publication 950 (1986), Safety of Information Technology Equipment, Including Electrical Business Equipment (Amendment 1, 1988).

[8] IEEE C37.90.1-1989, IEEE Standard Surge Withstand Capability (SWC) Tests for Protective Relays and Relay Systems (ANSI).[7]

[9] IEEE Std 802-1990, IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture.

[10] Reserved for future use.[8]

[11] IEEE Std 802.7-1989, IEEE Recommended Practices for Broadband Local Area Networks (ANSI).

[12] ISO 2382-9 : 1984, Data processing—Vocabulary—Part 09: Data communication.[9]

[13] ISO 7498 : 1984, Information processing systems—Open systems interconnection—Basic reference model.

[14] ISO/IEC 7498-4 : 1989, Information processing systems—Open systems interconnection—Basic reference model—Part 4: Management framework.

[15] ISO 8802-2 : 1989 (IEEE Std 802.2-1989), Information processing systems—Local area networks—Part 2: Logical link control.

[5]CCITT documents are available in the US from the US Dept. of Commerce, National Technical Information Service, 5285 Port Royal Rd., Springfield, VA 22161. CCITT documents are also available from CCITT General Secretariat, International Telecommunications Union (ITU), Sales Section, Place des Nations, CH-1211, Genève 20, Switzerland/Suisse.

[6]CISPR and IEC documents are available in the US from the Sales Department, American National Standards Institute, 1430 Broadway, New York, NY 10018. These documents are also available from International Electrotechnical Commission, 3 rue de Varembé, Case postale 131, CH1211, Genève 20 Switzerland/Suisse.

[7]IEEE documents are available from the Service Center, Institute of Electrical and Electronics Engineers, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331.

[8]When the following IEEE project is approved and published, it will become a part of this reference section: P802.1B, Local Area Networks and Metropolitan Networks Management.

[9]ISO documents are available in the US from the Sales Department, American National Standards Institute, 1430 Broadway, New York, NY 10018. ISO documents are also available from the ISO Office,1 rue de Varembé, Case postale 56, CH-1211, Genève 20, Switzerland/Suisse.

[16] ISO 8824 : 1987, Information processing systems—Open systems interconnection—Specification of abstract syntax notation one (ASN.1).

[17] Reserved for future use.[10]

[18] Reserved for future use.[11]

[19] Reserved for future use.[12]

[20] Reserved for future use.[13]

1.4 Conformance. Implementations that claim conformance with this standard shall

(1) Offer the mandatory LLC-MAC interface services specified in Section 2,

(2) Support the mandatory layer management actions, entities, and values specified in Sections 3 and 9,

(3) Generate, transmit, receive, and recognize the frames and sequences specified in Section 4,

(4) Support the mandatory capabilities, constraints, and behavior of the medium access protocol specified in Sections 6 and 7,

(5) Offer the mandatory MAC sublayer-Physical Layer interface services specified in Section 8,

(6) Support the mandatory capabilities, features, and constraints of at least one of the specified PLEs and medium specifications defined in Sections 12 and 13, 14 and 15, 16 and 17, and 18 and 19,

(7) Support at least one of the data rates specified for the chosen PLE, and

(8) Support the capabilities, behavior, or values defined for all options claimed to be supported by the implementation.

Unless otherwise stated herein, all specifications shall apply over the entire operating environmental range for which the manufacturer claims conformance.

A number of options are specified in this standard. An implementation shall indicate which, if any, of these options are supported.

[10] When the following ISO/IEC projects is completed, approved, and published, it will become a part of this reference section: JTCI DIS 9314-3, Information processing systems—Fibre distributed data interface (FDDI)—Part 3: Physical layer medium dependent (PMD) requirements.

[11] When the following ISO/IEC project is completed, approved, and published, it will become a part of this reference section: JTC1 DIS 9595, Information processing systems—Open systems interconnection—Management information services definition.

[12] When the following ISO/IEC project is completed, approved, and published it will become a part of this reference section: JTC1 DIS 9596, Information processing systems—Open Systems Interconnection—Management information protocol specification.

[13] When the following ISO/IEC project is completed, approved, and published, it will become a part of this reference section: JTC1 DIS 10039, Information technology—Telecommunications and information exchange between systems—Medium access control service definition.

1.5 Overview of the Token Method

1.5.1 The Essence of the Token Access Method

(1) A *token* controls the right of access to the physical medium; the station which holds (possesses) the token has momentary control over the medium.

(2) The token is passed by stations residing on the medium. As the token is passed from station to station a logical ring is formed.

(3) Steady state operation consists of a data transfer phase and a token transfer phase.

(4) Ring maintenance functions within the stations provide for ring initialization, lost token recovery, new station addition to the logical ring, and general housekeeping of the logical ring. The ring maintenance functions are replicated among all the token-using stations on the network.

Shared media generally can be categorized into two major types. These types are *broadcast* and *sequential*. This standard deals exclusively with the broadcast type. On a broadcast medium, every station may receive all signals transmitted. Media of the broadcast type are usually configured as a physical bus.

In Fig 1-2, note that the token medium access method is always sequential in a logical sense. That is, during normal, steady state operation, the right to access the medium passes from station to station. Furthermore, note that the physical connectivity has little impact on the order of the logical ring and that stations can respond to a query from the token holder even without being part of the logical ring. (For example, stations H and F can receive frames and could respond but cannot initiate a transmission since they will never be sent the token.)

The MAC sublayer provides sequential access to the shared bus medium by passing control of the medium from station to station in a logically circular fashion. The MAC sublayer determines when the station has the right to access the shared medium by recognizing and accepting the token from the predecessor station, and it determines when the token shall be passed to the successor station.

Fig 1-2
Logical Ring on Physical Bus

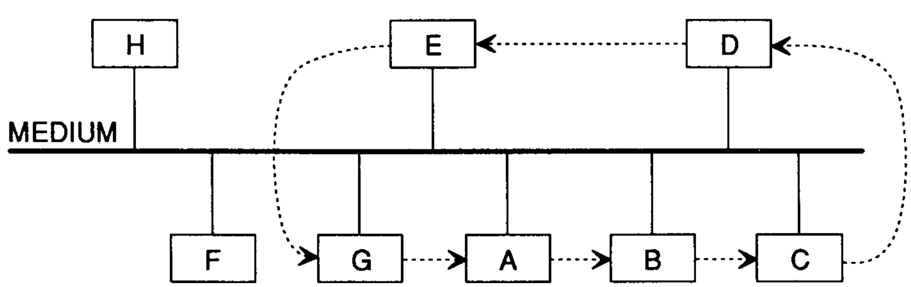

1.5.2 General Sublayer Functions

(1) Lost token timer
(2) Distributed initialization
(3) Token holding timer
(4) Limited data buffering
(5) Node address recognition
(6) Frame encapsulation (including token preparation)
(7) Frame check sequence (FCS) generation and checking
(8) Valid token recognition
(9) Ring member addition/deletion
(10) Node failure error recovery

1.6 MAC Sublayer Internal Structure. The MAC sublayer performs several functions that are loosely coupled. The descriptions and specifications of the MAC sublayer in this standard are organized in terms of one of several possible partitionings of these functions. The partitioning used here is illustrated in Fig 1-3, which shows five asynchronous logical "machines," each of which handles some of the MAC functions, as discussed in 1.6.1 through 1.6.5.

Of these five machines, the access control machine (ACM) is the most critical and the most complex; it is the key control mechanism for the token-passing bus access method, and it cooperates closely with ACMs in other stations given only limited information regarding the state of the network. Because of its impor-

Fig 1-3
MAC Sublayer Functional Partitioning

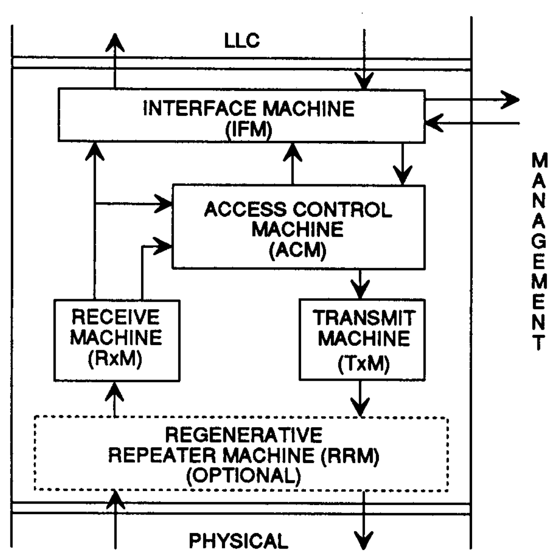

tance, and because much of its operation is not easily inferred from its functional requirements, the explanation and specification of the ACM are the major concerns of Sections 5, 6, and 7.

The interface machine (IFM) and receive machine (RxM) participate heavily in the operation of the MAC sublayer protocol. They are, however, discussed only in sufficient detail so that the reader will understand their role in the MAC sublayer in support of the ACM. The level of discussion was chosen to avoid any ambiguities which might compromise the coexistence of conformant stations on a single bus.

1.6.1 Interface Machine (IFM). This machine acts as an interface and buffer between the LLC and MAC sublayers and between network management and the MAC sublayer. It interprets all incoming MA-UNITDATA and other service primitives and generates appropriate outgoing service primitives. This machine handles the mapping of "quality of service" parameters from the LLC view to the MAC view, where this is necessary. It handles queuing of service requests, for example, requests to send an LLC protocol data unit (PDU). Finally, it performs the "address recognition" function on received data frames, accepting only those addressed to this station.

1.6.2 Access Control Machine (ACM). This machine cooperates with the ACMs of all other stations on the bus in handling the token to control transmission access to the shared bus. The ACM may (optionally) manage multiple MAC access classes so as to provide different levels of "quality of service" to the LLC sublayer. The ACM is also responsible for initialization and maintenance of the logical ring, including the admission of new stations. Finally, it has responsibility for the detection of and, where possible, recovery from faults and failures in the token-passing bus network.

1.6.3 Receive Machine (RxM). This machine accepts atomic symbol inputs from the Physical Layer and assembles them into frames which it validates and passes to the IFM and ACM. The RxM accomplishes this by recognizing the frame start and the frame end delimiters (SD and ED), checking the FCS and validating the frame's structure. The RxM also identifies and indicates the reception of *noise_bursts* and the *bus_quiet* condition.

1.6.4 Transmit Machine (TxM). This machine generally accepts a frame from the ACM and transmits it, as a sequence of atomic_symbols in the proper format, to the PLE. The TxM builds a MAC protocol-data-unit by prefacing each frame with the required preamble and SD, and appending the FCS and ED. When operating with a Regenerative Repeater Machine (RRM), the TxM's operation may be some what different.

1.6.5 Regenerative Repeater Machine (RRM). This machine is an optional MAC component that is present only in special "repeater" stations, for example, a broadband remodulator. In such repeater stations, the RRM, when appropriate, repeats the incoming atomic_symbol stream, from the PLE, back to the PLE for retransmission; in such cases the PLE is understood to be connected to at least two different segments of a single bus. The RxM and TxM may cooperate with the RRM in such repeating operations.

1.7 PLE and Medium. The broad outlines for the PLE sections of this standard and also an introduction to the various physical signaling techniques and media defined for LANs using the token-passing MAC protocol are given here.

Four bus media and corresponding PLEs are specified for use with the token-passing MAC protocol.

Each PLE and corresponding bus medium is described in a pair of consecutive sections (that is, Sections 12 and 13, Sections 14 and 15, Sections 16 and 17, and Sections 18 and 19), consisting of the following:

(1) A section specifying a specific PLE, including how the generic management objects (Section 9) in that PLE are particularized, and

(2) A section specifying the medium appropriate to that PLE.

Four different PLEs with corresponding media suitable for use with the token-passing bus MAC protocol are specified in this standard. They are distinguished primarily by the different forms of media and signaling specified for each type of PLE. The remainder of this subsection outlines the salient points of each type of PLE and corresponding medium.

1.7.1 Summary of Phase-Coherent Frequency Shift Keying (FSK)

Topology. Omnidirectional bus

Cable. 75 Ω coaxial cable, such as RG-11 type or semirigid

Connector at Station. 75 Ω female F-series, to be specified by IEC (see NOTES in 13.5.1)

Recommended Cable Configuration. CATV trunk of either RG-11 type or semi rigid and flexible drop cables of up to 50 m in length

Trunk Connection Unit. 75 Ω nondirectional 20 dB passive impedance-matching tap

Repeaters. Active regenerative repeaters used for high-fanout branching and extension of the system beyond the limits of a single segment of the cable

Transmit Level. + 63 to + 66 dB (1 mV, 75 Ω) [dBmV]

Receiver Sensitivity. + 10 dBmV to + 66 dBmV

Data Rates. 5 Mb/s and 10 Mb/s

Signaling. Direct encoding of *data* and *non_data* symbols, each as an integral number of cycles of constant frequency, with frequency changes only at the zero crossings of the waveforms. Two frequencies are used:

(1) The lower is 1 Hz/(b/s) (that is, 5 MHz at 5 Mb/s, 10 MHz at 10 Mb/s)

(2) The higher is 2 Hz/(b/s) (that is, 10 MHz at 5 Mb/s, 20 MHz at 10 Mb/s)

The symbol representations are

zero—two full cycles of the higher frequency

one—one full cycle of the lower frequency

pairs of *non_data*—one full cycle of the higher frequency, one full cycle of the lower frequency, and another one full cycle of the higher frequency

Pad_idle. Alternating *one* and *zero* symbols, starting with *one*

Clock Recovery. From zero crossings in the received signal

Transmit Data Timing. Phase-locked to transmit frequencies. All ports of a regenerative repeater use the same transmit data timing and, hence, identical transmit frequencies.

1.7.2 Summary of Multilevel Duobinary Amplitude Modulation/Phase Shift Keying (AM/PSK)

Topology. Directional bus with active repeater

Cable. 75 Ω coaxial cable, such as RG-6 type or semirigid

Connector at Station. 75 Ω female F-series, to be specified by IEC (see NOTE in 15.5.1)

Recommended Cable Configuration. CATV-like semirigid trunk and flexible drop cable

Trunk Connection Unit. 75 Ω directional passive impedance-matching tap

Repeaters. Regenerative repeater used as system data-rate clock source, as central assessor of contention and noise, and to retransmit all signaling received on the directional medium

Amplifiers. Standard CATV bidirectional (or unidirectional for dual-cable configuration) broadband amplifiers used for extension of the system beyond the basic signal loss budget

Channel Bandwidths. 1.5 MHz, 6 MHz, and 12 MHz

Transmit Level

In 1.5 MHz bandwidth, + 41 dB (1 mV, 75 Ω) [dBmV]
In 6 MHz bandwidth, + 47 dB (1 mV, 75 Ω) [dBmV]
In 12 MHz bandwidth, + 50 dB (1 mV, 75 Ω) [dBmV]

Receiver Sensitivity

In 1.5 MHz, –13 to + 4 dB (1 mV, 75 Ω) [dBmV]
In 6 MHz, –7 to +10 dB (1 mV, 75 Ω) [dBmV]
In 12 MHz, –4 to +13 dB (1 mV, 75 Ω) [dBmV]

Data Rates. 1 Mb/s in 1.5 MHz, 5 Mb/s in 6 MHz, 10 Mb/s in 12 MHz

Channel Frequency Assignments. Channel frequency assignments are a subject for national standardization. The following North American recommended channel assignments are included for information only. They are not a part of this International Standard.

For 10 Mb/s, channels 3' and 4' (59.75–71.75 MHz) and P and Q (252–264 MHz)

For 5 Mb/s, channels 3' (59.75–65.75 MHz) and P (252–258 MHz), or 4' (65.75–71.75 MHz) and Q (258–264 MHz)

For 1 Mb/s, any of the eight equally spaced 1.5 MHz subchannels of reverse channels 3' and 4', paired with the corresponding forward channels P and Q

Scrambler. A self-sync scrambler with a generator polynomial of $1 + X^{-6} + X^{-7}$ is applied to all framed data prior to encoding for transmission, both to increase the average number of transitions within a transmission and to randomize spectral components of the transmitted modulation.

Signaling. *Data* and *non_data* symbols are encoded so as to specify the amplitude of the eventual modulation. A single form is defined in this standard: One MAC-symbol per PHY-symbol, together with compatibility considerations for an additional form described in 14.11.

In both forms of signaling, the middle level is used only to signal the *non_data* symbols found in frame delimiters and reported *silence*, and to break up long sequences of a single other signaling level (see 14.8.2.1 (4) (c)). Such long sequences are unlikely since a scrambler is applied to framed data before it is encoded for transmission.

For one MAC-symbol per PHY-symbol signaling, the symbol representations at the receiver are

{*0*} = *zero*—zero amplitude
{*4*} = *one*—"maximum" amplitude
{*2*} = *non_data*—"middle" amplitude of one half "maximum"

Modulation. Multilevel duobinary AM/PSK

***Pad_idle*.** Alternating {*4*}} and {*0*} symbols, starting with {*4*}

Reported *Silence*. The repeated sequence of symbols transmitted by the remodulator to report that it is not receiving any signaling. The sequence has a period of four symbols and the repetition may break off after any symbol of the sequence. Listening modems can set their automatic gain control (AGC), and can determine the signaling mode of the system by the sequence used.

For one MAC-symbol per PHY-symbol signaling, the sequence is:

{*2*} {*2*} {*0*} {*4*}

For a possible future two MAC-symbols per PHY-symbol signaling mode, the following sequence is reserved:

{*2*} {*2*} {*4*} {*0*}

Implementations that detect this reserved repeating sequence should either switch to the enhanced mode of operation or inhibit transmission.

Clock Recovery. From level transitions within the received signaling

Transmit Data Timing. Sourced by remodulator; frequency-locked to received data timing at all other stations

1.7.3 Summary of Fiber Optic Media

Topology. Directional bus using active or passive stars

Cable. The standard assumes a silica fiber optic waveguide with the following nominal characteristics: a core diameter of 62.5 μm, outside diameter of 125 μm, and an effective numerical aperture of 0.275.

NOTE: Operation using a 50 μm alternate test fiber is described in Appendixes 16.10 and 17.9.

Connector at Station. The cable plant interface connector (CPIC) is a duplex connector as defined in ISO/IEC 9314-3 [16].

Repeaters. Active regenerative repeaters used for high-fanout star topologies

Transmit Characteristics. –7 to –11 dBm effective optical launch power with center wavelength between 800 and 910 nm

Receiver Sensitivity

Moderate sensitivity: –11 to –31 dBm effective power with silence level of –40 dBm

High sensitivity: –21 to –41 dBm effective power with silence level of –50 dBm

Data Rates. 5 Mb/s, 10 Mb/s, and 20 Mb/s

Signaling. Manchester encoding of *data* and *non_data* symbols

The symbol representations are

{*H L*} = *zero*

{*L H*} = *one*

{*L L H H*} and {*H H L L*} = pairs of *non_data* symbols

***Pad_idle*.** Alternating one and zero symbols, starting with *one*

Clock Recovery. From transitions generated by the Manchester encoding

Transmit Data Timing. Phase-locked to transmit frequencies. All ports of a regenerative repeater use the same transmit data timing and, hence, identical transmit frequencies.

1.7.4 Summary of Phase-Continuous FSK

Topology. Omnidirectional bus

Trunk Cable. 75 Ω coaxial cable, such as types RG-6, RG-11, and semirigid

Drop Cable. 35–50 Ω coaxial cable stub, less than 350 mm long

Connector at Station. 50 Ω male BNC-series, see IEC 169-8 (1978) [4]

Recommended Cable Configuration. Long unbranched trunk cable with very short “stub” drop cables

Trunk Connection Unit. 75 Ω tee connector

Repeaters. Active regenerative repeaters used for branching and extension of the system beyond the basic signal loss budget

Transmit Level. + 54 to + 60 dB (1 mV, 37.5 Ω)

Receiver Sensitivity. + 24 dB (1 mV, 37.5 Ω)

Data Rate. 1 Mb/s

Signaling. Manchester encoding of *data* and *non_data* symbols

The symbol representations are

{*H L*} = *zero*—initial high, final low level

{*L H*} = *one*—initial low, final high level

{*L L H H*} and {*H H L L*} = pairs of *non_data* symbols, initial pair of low level followed by a final pair of high level or vice versa

Modulation. Phase-continuous FSK (a form of frequency modulation), with Manchester representation:

(1) Frequency of high level = 6.25±0.08 MHz
(2) Frequency of low level = 3.75±0.08 MHz

***Pad_idle*.** Alternating *one* and *zero* symbols, starting with *one*

Clock Recovery. From transitions generated by the Manchester encoding

1.7.5 Alternate Physical Layer and Medium for Industrial Control. Alternate Physical Layers and media for industrial control are specified in ANSI/ ISA-S72.01-1986 [1]. These alternate Physical Layers and media are compatible at the MAC sublayer-Physical Layer interface.

1.8 Access Method Characteristics. An understanding of the basic characteristics of the token-passing access method is useful, so as to better understand where and when a token-passing bus is an appropriate LAN technology.

Some of the important features of this medium access method are as follows:

(1) The method is efficient in the sense that, under high offered load, the coordination of the stations requires only a small percentage of the medium's capacity.
(2) The method is fair in the sense that it offers each station an equal share of the medium's capacity. It does not, however, require any station to use its full share.
(3) The method permits multiple classes of service.
(4) The method coordinates the stations' transmissions so that they minimize and control their interference with each other.
(5) The method imposes no additional requirements on the medium or the PLE's capabilities over those necessary for transmission and reception of multibit, multiframe sequences at the specified bit error rate.
(6) In the absence of system noise, the method provides computable, deterministic, worst-case bounds on access delay for the highest priority class of service for any given network and loading configuration.
(7) Periods of controlled interference are distinguishable; system noise measurements are possible during the remaining periods.
(8) The method places minimal constraints on how a station that momentarily controls the medium may use its share of the medium's capacity.

(9) The method efficiently supports the proposed LLC Type 3 Service by allowing the token holding station to wait for the receiving station to respond to the token holder's transmission.

(10) Though not specified by this standard, the method permits the presence of large numbers of low-cost, reduced-function stations in the network, together with one or more full-function stations. (At least one full-function station is needed to make the system operational.) An example of a reduced function station is one that does not contain access control logic.

1.9 Standard Organization. This standard is organized in 19 sections, which are summarized below.

Section 1 (this section) begins with a general discussion of the token-passing bus access method. The MAC sublayer functional partitioning used in subsequent discussions is introduced here. The PLE and media options are also surveyed. Finally, features of the token-passing bus access method are reviewed.

Section 2 details the logical interfaces between the LLC sublayer and the MAC sublayer, and the services and command interfaces (such as sending a frame) provided to the LLC sublayer.

Section 3 details the management parameters, actions, and events within the MAC sublayer.

Section 4 details the general MAC frame structure, including delimiters, addressing, and the FCS. All of the frame formats that the MAC handles, including MAC control frames, are enumerated.

Section 5 discusses the basic concepts of the access protocol and provides an informal description of the actions in each state of the ACM. The other state machines of the MAC sublayer are also described in Section 5.

Section 6 contains definitions of essential MAC terms and components.

Section 7 specifies the MAC ACM by means of a state machine model. This is the definitive specification of the token-passing bus MAC operation. Section 7 also describes the MAC sublayer variables, functions, and procedures used in the state machine.

Section 8 details the logical interface between the MAC sublayer and the Physical Layer. Included here are descriptions of the interface symbols, requests, and responses.

Section 9 details the management parameters, actions, and events within the PLE.

Section 10 defines the logical, electrical, and mechanical interface within the PLE between the station and a separated modem.

Section 11 is reserved.

Sections 12 and 13 detail the PLE and medium, respectively, for a single-channel (that is, omnidirectional) phase-coherent-FSK coaxial cable bus at 5 Mb/s or 10 Mb/s.

Sections 14 and 15 detail the PLE and medium, respectively, for a dual-channel (that is, head-ended) broadband duobinary AM/PSK coaxial cable bus at 1 Mb/s, 5 Mb/s, or 10 Mb/s.

Sections 16 and 17 detail the PLE and medium, respectively, for a fiber optic bus at 5 Mb/s, 10Mb/s, or 20 Mb/s.

Sections 18 and 19 detail the PLE and medium, respectively, for a single-channel (that is, omnidirectional) phase-continuous-FSK coaxial cable bus at 1 Mb/s.

2. LLC-MAC Interface Service Specification

This section specifies the services provided to the Logical Link Control (LLC) sublayer, and to the alternative sublayer specified in ANSI/ISA-S72.01-1986 [1], at the boundary between the LLC functions and the Medium Access Control (MAC) sublayer of the Data Link Layer of the reference model. This standard specifies these services in an abstract way. It does not specify or constrain the implementation entities and interfaces within a computer system. The relationship of this section to other sections of this standard and to LAN specifications is illustrated in Fig 2-1.

NOTES: (1) The exact relationship of the layers described in this standard to the layers defined by the Open Systems Interconnection (OSI) Basic Reference Model (ISO 7498 : 1984 [13]) is for future study.

(2) Work is in process to produce a single service specification that is common to all the MAC sublayers; see ISO/IEC 10039 [20]. When that standard is completed, approved, and published, Section 2 will be deleted.

2.1 Overview of the LLC-MAC Service

2.1.1 General Description of Services Provided. This section informally describes the services provided to the LLC sublayer by the token-passing MAC sublayer, both of the Data Link Layer. These services provide only connectionless

Fig 2-1
Relation to LAN Model

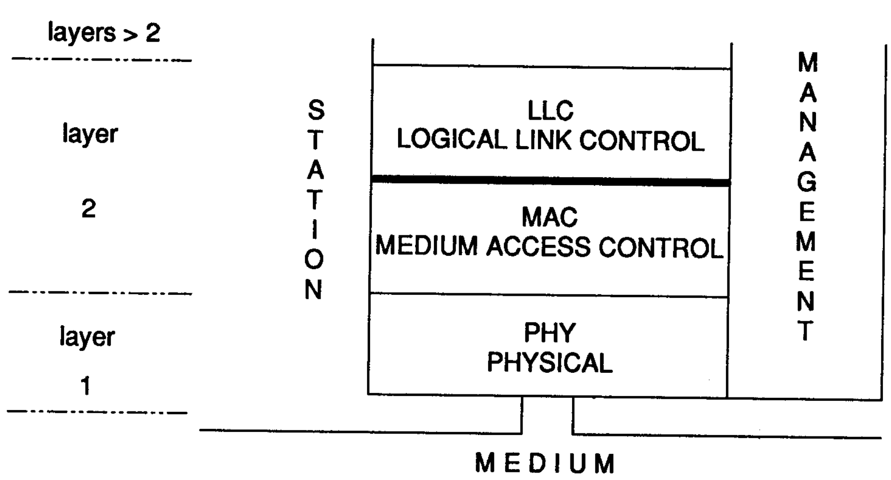

data transfer services between peer LLC entities. They provide the means by which LLC entities can exchange MAC service data units (M_SDU) without the establishment of an underlying point-to-point connection. The data transfer can be point-to-point or multipoint, unacknowledged or acknowledged.

2.1.2 Model Used for the Service Specification. The model and descriptive method are detailed in Appendix A.

2.1.3 Overview of Interactions. The primitives associated with this connection less data transfer service are

(1) MA-UNITDATA request
(2) MA-UNITDATA indication
(3) MA-UNITDATA confirm

An MA-UNITDATA request primitive is passed to the MAC sublayer to request that an M_SDU be sent. (All M_SDUs are sent using connectionless procedures.) An MA-UNITDATA indication primitive is passed from the MAC sublayer to indicate the arrival of an M_SDU. An MA-UNITDATA confirm primitive is passed from the MAC sublayer to indicate the status of the previous associated MA-UNITDATA request primitive.

2.1.4 Basic Services and Options. All services are mandatory and are required in all implementations.

2.2 Detailed Interactions with the LLC Entity. This subsection describes in detail the primitives and parameters associated with the connectionless data transfer services provided to LLC by MAC. Note that the parameters are specified in an abstract sense. The parameters specify the information that shall be available to the receiving entity. A specific implementation is not constrained in the method of making this information available. For example, the M_SDU parameter associated with some of the data transfer service primitives may be provided by actually passing the MAC service data unit, by passing a descriptor, or by other means. The values of some selection parameters may be implied by an implementation.

2.2.1 MA-UNITDATA request

2.2.1.1 Function. This primitive is the service request primitive for the connectionless data transfer service.

2.2.1.2 Semantics. The primitive shall provide parameters as follows:

```
MA-UNITDATA request    (
                       destination address,
                       source address,
                       M_SDU,
                       desired_quality
                       )
```

The destination_address parameter specifies either an individual or a group MAC-entity address. The source_address parameter specifies the originating MAC-entity address, typically the local station. The M_SDU parameter specifies the MAC service data unit to be transmitted by the MAC sublayer entity for the requesting LLC sublayer entity.

The desired_quality parameter specifies the desired quality of service. The semantics of this parameter include both MAC-level priority, with a range of 0 (lowest) to 7 (highest) (see 6.6.1.2), and MAC-level delivery confirmation service, with values of request_with_no_response, request_with_response, and response.

2.2.1.3 When Generated. This primitive is passed from the LLC sublayer entity to the MAC sublayer entity to request that the MAC sublayer entity compose and transmit the specified frame at the desired quality of service on the LAN.

2.2.1.4 Effect on Receipt. Receipt of this primitive causes the MAC sublayer entity to attempt to compose and transmit the specified frame.

2.2.1.5 Additional Comments. A value of request_with_response for the delivery-confirmation component of the quality parameter indicates that the next MA-UNITDATA indication should itself have a quality parameter specifying response, in which case that next MA-UNITDATA indication shall be associated with this MA-UNITDATA request.

A value of response for the delivery-confirmation component of the quality parameter indicates that the immediately prior MA-UNITDATA indication shall itself have had a quality parameter specifying request_with_response.

A group destination_address should not be used when specifying request_with_response.

2.2.2 MA-UNITDATA indication

2.2.2.1 Function. This primitive is the service indication primitive for the connectionless data transfer service.

2.2.2.2 Semantics. The primitive shall provide parameters as follows:

```
MA-UNITDATA indication    (
                           destination_address,
                           source_address,
                           M_SDU,
                           quality
                           )
```

The destination_address and source_address parameters specify the DA and SA fields of a frame (see Section 4) as received by the local MAC entity, and thus the MAC entities purportedly involved in the communication. The M_SDU parameter specifies the MAC service data unit as received by the local MAC sublayer entity.

The quality parameter specifies the delivered quality of service. The semantics of this parameter include both MAC-level priority, with a range of 0 (lowest) to 7 (highest) (see 6.6.1.2), and MAC-level delivery-confirmation service, with values of request_with_no_response, request_with_response, and response.

2.2.2.3 When Generated. This primitive is passed from the MAC sublayer entity to the LLC sublayer entity to indicate the arrival of a data frame from the Physical Layer entity (PLE). Such frames are reported only when they are free of detected errors and their (individual or group) destination address designates the local MAC entity.

2.2.2.4 Effect on Receipt. The effect of receipt of this primitive by the LLC entity is specified in ISO 8802-2 : 1989 (IEEE Std 802.2-1989) [15].

2.2.2.5 Additional Comments. In the absence of undetected errors, the contents of the M_SDU parameter are logically complete and unchanged relative to the M_SDU parameter in the associated MA-UNITDATA request at the sending station.

NOTE: This is a guarantee of transparency.

A value of request_with_response for the delivery-confirmation component of the quality parameter indicates that the receiving LLC sublayer entity should immediately respond with an MA-UNITDATA request which itself has a quality parameter specifying response.

A value of response for the delivery-confirmation component of the quality parameter indicates that this MA-UNITDATA indication may be associated with a prior MA-UNITDATA request which itself had a quality parameter specifying request_with_response, and which was issued by the same LLC sublayer entity.

2.2.3 MA-UNITDATA confirm

2.2.3.1 Function. This primitive provides the LLC sublayer with status information for the previous associated MA-UNITDATA request primitive.

2.2.3.2 Semantics. The primitive shall provide parameters as follows:

```
MA-UNITDATA confirm    (
                       destination_address,
                       source_address,
                       status,
                       provided_quality
                       )
```

The destination_address and source_address parameters specify the destination address and source address fields of the associated MA-UNITDATA request primitive. The status parameter indicates the status of the service provided for a previous associated MA-UNITDATA request primitive.

The provided_quality parameter specifies the quality of service actually provided for the previous request. The semantics of this parameter include both MAC-level priority, with a range of 0 (lowest) to 7 (highest) (see 6.6.1.2), and MAC-level delivery confirmation service.

2.2.3.3 When Generated. This primitive is passed from the MAC sublayer entity to the LLC sublayer entity to indicate the status of the service provided for the previous associated LLC data transfer request.

A success indication is passed when the request has been completed successfully, either because no response was requested or because the requested response was received. A failure indication is passed upon local failure or when a requested response is not received, even after the specified number of retries. When the quality parameter of that request specifies request_with_response, a failure indication is also passed if the allowed number of retries has occurred with no response.

2.2.3.4 Effect on Receipt. The effect of receipt of this primitive by the LLC sublayer entity is specified in ISO 8802-2 : 1989 (IEEE Std 802.2-1989) [15].

2.2.3.5 Additional Comments. It is assumed that sufficient information is available to the LLC sublayer to associate the status with the appropriate request.

3. MAC Sublayer Management

This section specifies the abstract entities (parameters, events, and actions) used in managing the Medium Access Control (MAC) sublayer. The relationship of this section to other sections of this standard and to Local Area Network (LAN) specifications is illustrated in Fig 3-1.

There are two ways to manage a MAC sublayer: locally or remotely.

Local management is required to initially configure a MAC and to operate it when remote management is not available. Since the operation of local MAC management, other than the need to manage the entities described in this section, does not affect station interoperability, the specification of local management is outside the scope of this standard.

This standard uses the concept of remote management as defined in the Open Systems Interconnection (OSI) Management Framework described in ISO/IEC 7498-4 : 1989 [14]. Remote management uses a communications protocol to manipulate and observe the managed entities. This section is designed to be used with the companion remote management standard described by IEEE P802.1B [9]. Other protocols, such as described by ISO/IEC 9595 [18] and ISO/IEC 9596 [19], may also be used.

Fig 3-1
Relation to LAN Model

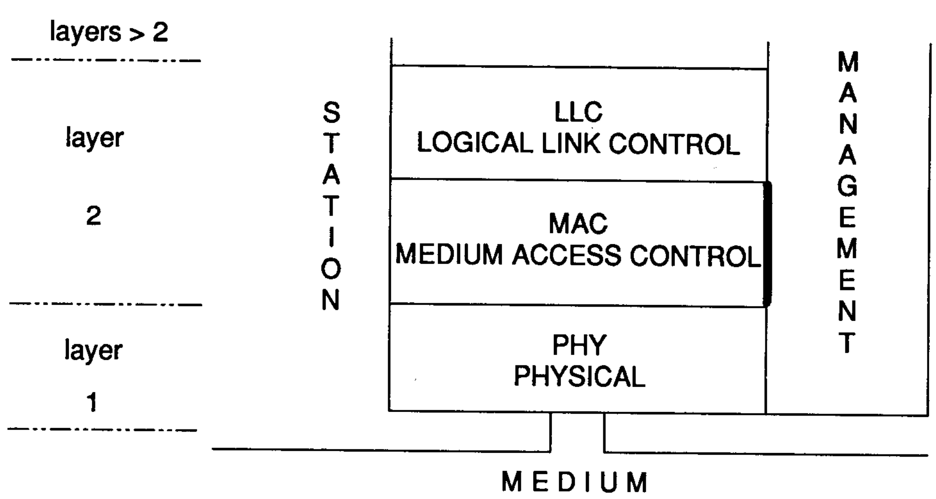

To implement a remote management protocol, the identification and value representation of the managed entities must be defined. Such a definition is given by a layer management interface (LMI) protocol.

NOTE: The definition of a MAC sublayer LIM protocol is under study. An addendum to this standard is planned that will contain this protocol.

Within this document the general term "management" is used to refer to either local or remote management, unless otherwise indicated.

3.1 Overview. This section describes the entities communicable via the LMI, shown as a thick line in Fig 3-1. The MAC sublayer entities described in this section consist of

(1) Parameters within the MAC sublayer written and read by the management entity,
(2) Actions initiated by the management entity that cause changes within the MAC sublayer,
(3) Events within the MAC sublayer entity that are passed to the management entity.

The notation used to specify the entities defined in this section is the Abstract Transfer Syntax and Notation (ASN.1), defined in ISO 8824 : 1987 [16]. This syntax differs from that used in Sections 4 through 7 to describe the token bus MAC. The difference is in the spelling of the entity names. The ASN.1 notation does not allow the underscore character, "_", and uses a combination of uppercase and lowercase characters to improve readability.

For example, this section defines a parameter *inRing*. Section 7 defines a parameter *in_ring*. While the syntax of the names is different, the parameters are equivalent.

3.2 MAC Management Facilities. This subsection contains a description of use of the management parameters, actions, and events within the token-passing bus MAC sublayer.

3.2.1 Organization. There are three types of management entities:

(1) Parameters
(2) Actions
(3) Events

The MAC sublayer parameters accessible by the management entity are described first. These parameters are organized into groups. The groups not only collect the variables, they also define collections of parameters that can be accessed in a common manner.

NOTE: In implementing control of access to parameters, the intent of this standard is to implement the granularity of access control only to the group level. Thus, all parameters within a group share a common access level.

Any or all of the parameters in a group may be passed in a single remote management request. The semantics and grouping of the parameters so communicated are specified in the following subsections.

3.2.2 ResourceTypeID. This subsection discusses the parameters of the MAC sublayer LMI that are common to any layer. These parameters create a "name plate" that identifies the basic attributes of the MAC sublayer.

The ResourceTypeID is common to all MAC sublayers. Implementation of this group shall be required. The ResourceTypeID group is defined in IEEE P802.1B [9].

3.2.2.1 resourceType. This parameter identifies the resource as a MAC sublayer.

3.2.2.2 standardRevision. This parameter identifies the revision of this standard implemented by the MAC sublayer.

3.2.2.3 ImeOptions. This parameter identifies the options supported by the implementation. The options defined by this standard are priority and request_with_response.

3.2.3 MAC Characteristics Group. The MAC characteristics group defines attributes of the MAC sublayer that are specific to a token-passing bus MAC sublayer. The implementation of the parameters in the MAC characteristics group shall be optional except for maxSDUSize, which is required. All parameters shall be read-only.

3.2.3.1 maxSDUSize. The maximum size service data unit that the sublayer implementation is capable of accepting.

The value of maxSDUSize is between 516 (to allow 512 octet frames) and 8191 octets.

3.2.3.2 transmitTimeDelay. The estimated maximum delay incurred by the MAC in sending a data frame, including the time in the MAC transmit queue. The value of this parameter is for use by higher layers in computing a time-out delay to use in detecting lost frames. The type is an array of time values indexed by MAC service class (see 6.6.1.2).

The units of this parameter and methods of computing the values are for future study.

3.2.4 MAC Management Group. The MAC management parameters describe the current state of the MAC sublayer. The implementation of access by the management entity to these parameters shall be optional. Access by the management entity to these parameters shall be read-only to prevent unpredictable changes to MAC operation. In order to change the value of one of these parameters, the management entity must change the corresponding value in the initialization group (see 3.2.7) and then perform the reset action to force the station to reinitialize.

3.2.4.1 state. MAC sublayer state. The current state of the MAC entity. The MAC state is changed by management actions (see 3.2.11) and events within the MAC access control machine (ACM) (see Section 7).

The range of values of the state parameter is offline, tokenBus, tokenBusBridge, and tokenBusRepeater.

NOTE: The operation of a station in the token-passing bus bridge mode is for future study. Bridge operation is envisioned to include acceptance of frames addressed to stations on networks other than the one to which the station is connected.

3.2.4.2 testStatus. The status of a previously initiated test. Negative, implementation-dependent values shall be treated as "failure" by entities that lack further information on the implementation-dependent meaning.

See 7.1.1 for a detailed description of the following parameters of the MAC management group.

3.2.4.3 ts. This station's address.

3.2.4.4 slotTime. The network slot_time.

3.2.4.5 minPostSilencePreambleLength. The minimum number of octets of preamble to prefix to a frame after silence. The value is obtained from the PLE and used by the transmit machine (TxM).

3.2.5 Station Performance Group. These parameters affect the operation of the station and the station's performance in the network. Management may read and write the station performance parameters in order to alter network operation without having to re-initialize the station. The implementation of access by the management entity to these parameters shall be optional.

See 7.1.1 for a detailed description of the parameters of the station performance group.

3.2.5.1 maxInterSolicitCount

3.2.5.2 maxRetryLimit

3.2.5.3 hiPriTokenHoldTime

3.2.5.4 targetRotationTime

3.2.5.5 ringMaintenanceTimerInitialValue

3.2.5.6 inRingDesired

3.2.6 Station Monitoring Group. These parameters reflect the current status of the network and the station's status with respect to the network.

The implementation of access by the management entity to these parameters shall be optional. The station monitoring parameters shall be read-only with the exception of maxTokenRotationTime, which shall be read-write.

See 7.1.6 for a description of the parameters of the station monitoring group not discussed in the following paragraphs.

3.2.6.1 ns. The address of the next station in the logical ring.

3.2.6.2 ps. The address of the previous station in the logical ring.

3.2.6.3 nsKnown

3.2.6.4 inRing

3.2.6.5 soleActiveStation

3.2.6.6 lastTokenRotationTime. The observed token rotation time for the last token rotation, timed from token arrival to token arrival. Units are octet times.

A value of zero indicates that the value is not currently available; for example, possibly when the station is not a member of the logical ring. A negative value indicates a range overflow.

3.2.6.7 maxTokenRotationTime. The maximum token rotation time. The value is created by comparing the old value of maxTokenRotationTime with the value of lastTokenRotationTime whenever the token is received. If the old maxTokenRotationTime value is less then the lastTokenRotationTime, the value of maxTokenRotationTime is set equal to the value of lastTokenRotationTime. Units are octet times.

The value of this parameter may be set to zero by the management entity to begin a time interval over which the maximum token rotation time is to be measured.

A value of zero indicates that the value is not currently available. A negative value indicates a range overflow.

The creation of other statistics from the maximum and instantaneous token rotation time variables is for future study.

3.2.6.8 approximateRingSize. The approximate number of stations in the logical ring.

A possible method of obtaining the value of approximateRingSize is for the station to count the number of valid token frames that it hears between sending and receiving the token. Such a simplistic approach is prone to error when tokens are repeated or heard in error. Implementors can improve the approximation with more sophisticated techniques if desired.

A value of zero means the value is currently unavailable, for example, possibly when the station is not a member of the logical ring. A negative value indicates a range overflow.

3.2.7 Initialization Group. The initialization parameters are copied to the MAC sublayer operational parameters when the station is initialized. The initialization parameters shall be required in implementations where the corresponding operational parameters exist. The initialization parameters shall be readable and writeable by the management entity.

See 7.1.7 for a detailed description of the parameters of the initialization group.

3.2.7.1 initTS

3.2.7.2 initSlotTime

3.2.7.3 initMaxInterSolicitCount

3.2.7.4 initMaxRetryLimit

3.2.7.5 initHiPriTokenHoldTime

3.2.7.6 initTargetRotationTime

3.2.7.7 initRingMaintenanceTimerInitialValue

3.2.7.8 initInRingDesired

3.2.8 Interface Machine (IFM) Counters Group. These parameters monitor the operation of the IFM. In counting frame length, octets in a frame are counted between the start delimiter and the end delimiter, exclusively. Service class to access class correspondence is specified in 6.6.1.2.

The implementation of the IFM counters and access by the management entity shall be optional. Access by the management entity to these parameters shall be read-only.

3.2.8.1 messagesSent. The number of data frames sent at the corresponding access class.

3.2.8.2 octetsSent. The number of octets in data frames sent at the corresponding access class, including control fields and excluding delimiters.

3.2.8.3 messagesReceived. The number of valid data frames received and accepted at the corresponding access class.

3.2.8.4 octetsReceived. The number of octets of valid data frames received and accepted at the corresponding access class, including control fields and excluding delimiters.

NOTES: (1) This standard does not include a count of frames discarded at the MAC-LLC interface due to lack of buffer space. The presumption is that the LLC sublayer is responsible for buffer management and discarding frames for lack of space.

(2) It is possible to have underruns while transmitting frames and/or overruns while receiving frames. Implementations may have counters that record the occurrence of these events. It is recommended that these counters be included in private management parameters of the IFM counters group.

3.2.9 Access Control Machine (ACM) Counters Group. These parameters monitor the operation of the ACM. The implementation of the ACM counters and access by the management entity shall be optional.

Implementation of a threshold facility associated with each counter shall be treated as a separate option. Thus, one of three conditions exists: No implementation, 32-bit counters without threshold, or 32-bit counters with threshold. If implemented, access by the management entity to these parameters shall be read only, except for the portions associated with the threshold mechanism. (The threshold facility can be used to generate an event to the management entity when the value of a counter exceeds a specified value.)

NOTE: The condition described in 6.7.4 can cause the token_pass_failure, who_follows_query, and solicit_any counters to increment during correct operation of the network.

3.2.9.1 whoFollowsQuery. The number of occurrences of the ACM who_follows_query state transition. (See 7.2.3.8.)

3.2.9.2 tokenPassFailed. The number of occurrences of the ACM token_pass_failed state transition when pass_state equals pass_token.

3.2.9.3 solicitAny. The number of occurrences of the ACM solicit_any state transition not due to the ACM do_solicit_any state transition.

3.2.9.4 noSuccessor. The number of occurrences of the ACM no_successor_8 state transition.

3.2.9.5 unexpectedFrame. The number of occurrences of the ACM unexpected_frame_6 and unexpected_frame_10 state transitions.

3.2.9.6 claimToken. The number of occurrences of the ACM no_token state transition.

3.2.10 Receiver Machine (RxM) Counters Group. These parameters monitor the operation of the RxM. The implementation of the RxM counters and access by the management entity shall be optional. Thus the implementation of the MAC variable noise_expected discussed in the next paragraph, shall also be optional. Access to the RxM counters shall be read-only, except for the portions associated with the optional threshold mechanism.

The MAC variable noise_expected is used only in the definition of the RxM counters and is not described elsewhere in this standard. The noise_expected variable attempts to track the state of the network so that errors reported by the modem can be classified.

Errors reported by the modem may actually be induced by “collisions” when multiple stations transmit during times when contention can occur. By tracking the state of the network, the ACM attempts to predict the times of possible contention. The variable noise_expected conveys this information.

To simplify the explanation, the ACM state tables in 7.2 do not include the variable noise_expected variable noise_expected. If required to implement the counters defined in this subsection, noise_expected shall be reset whenever the ACM

(1) Receives or sends a valid token frame
(2) Receives or sends a valid data frame

The noise_expected variable shall be set whenever the ACM

(a) Traverses the initialize arc
(b) Receives or sends a solicit_successor_1 frame
(c) Receives or sends a solicit_successor_2 frame
(d) Receives or sends a resolve_contention frame
(e) Receives or sends a who_follows frame
(f) Receives or sends a claim_token frame

Implementation of a threshold facility associated with each counter except validFrames shall be treated as a separate option. Thus, one of three conditions exists: No implementation, simple 32-bit counters, or thresholded 32-bit counters. If implemented, access by the management entity to these parameters shall be read only, except for the portions associated with the threshold mechanism.

3.2.10.1 validFrames. The number of valid frames received of all types, including MAC control frames. (See 4.1.)

3.2.10.2 modemErrors. The number of noise bursts that occur when noise_expected is reset. modemErrors is a count of noise bursts that are probably caused by errors in the PLE or as the result of noise introduced on the medium.

The following counters are defined in order to aid in diagnosing the cause of errors reported in modemErrors. The value of modemErrors is the sum of the values of the following counters:

(1) **fcsErrors.** The number of frames received with incorrect frame check sequence (FCS) and E-bit reset when noise_expected is reset. (This value represents a count of FCS errors not previously detected by a remodulator or other repeater.)
(2) **eBitErrors.** The number of frames received with the E-bit set and the FCS invalid when noise_expected is reset. (This value represents a count of FCS errors previously detected by a remodulator or other repeater.)
(3) **nonSilence.** The number of occurrences of non-silence followed by silence in which a start delimiter was not detected when noise_expected is reset. (This error is possibly caused by a noise “hit” on the start delimiter.)
(4) **frameFragments.** The number of occurrences of a start delimiter followed by a start delimiter, invalid symbol sequence, or silence without an intervening end delimiter. frameFragments is only incremented when noise_expected is reset. (This error is probably caused by premature termination of a frame by an abort sequence.)

The treatment of frames that are received with the correct FCS and the E-bit set is not specified by this standard. However, if such frames are discarded by the MAC as invalid, the count of such discarded frames shall be included in eBitErrors.

3.2.11 Action Definitions. Actions are requests by the management entity to perform some activity. The reset and initialize with a value of initializeTokenBus actions shall be required locally. All other actions shall be optional.

3.2.11.1 reset Action. This action causes the MAC sublayer to go to the OFFLINE state. The MAC sublayer shall remain in this state until a local initialize action occurs. A remote initialize action shall be ignored following a remote reset action.

The action may be used to disable the MAC sublayer if it is suspected of malfunctioning.

3.2.11.2 initialize Action. This action resets the MAC sublayer and then causes the station's ACM to traverse the initialize arc and enter the IDLE state in accordance with the action request parameter. A remote initialize action shall be ignored if the MAC sublayer is already in the OFFLINE state. The action request parameter may assume the values tokenBus, tokenBusBridge, or tokenBusRepeater.

Following an initialize action, the values of the counters specified in 3.2.8 through 3.2.10 are unspecified.

3.2.11.3 initiateTest Action. The initiateTest action commences an internal test of the MAC sublayer if the station is in the OFFLINE state. The result of the most recently initiated test is reported in the parameter testStatus. (See 3.2.4.2.)

The effect of the initiateTest action when the station is not in the OFFLINE state is for future study.

3.2.12 Event Definitions. Events are spontaneously reported by the sublayer to the management entity. Remote reporting of events shall be optional.

Local support of the thresholdReached event is required when thresholded counters are implemented. Support of the newSuccessor event shall be optional. Passing of the new successor's address shall be optional. Support of the noSuccessor event shall be optional. Local support of the duplicateAddressFault and faultyTransmitterFault events shall be required.

NOTE: In implementing routing of events, the intent of this standard is to provide a unique routing path for each event.

3.2.12.1 thresholdReached. A thresholded counter has passed the threshold value. The event value shall be the counter's parameter identifier.

3.2.12.2 newSuccessor. The station's successor has changed. The event value shall be the new successor address.

3.2.12.3 noSuccessor. The station now considers itself the sole member or not a member of the ring.

3.2.12.4 duplicateAddressFault. The station has detected another station with the same MAC address. The state transition causing this event is duplicate_address_1. (See 6.7.1 for possible additional detection methods.)

3.2.12.5 faultyTransmitterFault. The station has decided that it may have a faulty transmitter. The state transitions associated with such a decision are

(1) end_all_contention
(2) no_future

4. Frame Formats

This section defines the required Medium Access Control (MAC) frame formats. This includes all allowed frame formats and the arrangement of all frame subfields. The term frame as used here refers to the protocol data units (PDUs) exchanged by MAC sublayer entities. The MAC service data units received from the Logical Link Control (LLC) sublayer are contained within some of these MAC frames. The frame components and formats used by medium access control are also described in this section.

The MAC frames and abort sequences are described in the following subsections. First, the components of the frames are discussed, followed by the definition of the valid frame formats. All frames sent or received by the MAC sublayer shall conform to the following general format:

PREAMBLE	SD	FC	DA	SA	DATA_UNIT ...	FCS	ED

```
where
  PREAMBLE = pattern sent to set receiver's modem clock and level
             (1 or more octets)
        SD = start delimiter (1 octet)
        FC = frame control (1 octet)
        DA = destination address (2 or 6 octets)
        SA = source address (2 or 6 octets)
 DATA_UNIT = information (0 or more octets)
       FCS = frame check sequence (4 octets)
        ED = end delimiter (1 octet)
```

The number of octets between SD and ED, exclusive, of a data frame (see 4.1.3.2) shall be 8191 octets or fewer.

The abort sequence shall conform to the following format:

SD	ED

```
where
  SD = start delimiter (1 octet)
  ED = end delimiter (1 octet)
```

Within this section the following acronyms are used for the addresses of the station under discussion, its successor and its predecessor in the logical ring:

TS—this station's address
NS—next station's address
PS —previous station's address

4.1 Frame Components. This subsection describes the frame components that are shown in the previous illustrations in greater detail.

A valid frame consists of preamble, a start delimiter, a frame control octet, destination and source addresses, possibly a data_unit, a valid FCS, and an end delimiter, and does not meet any of the criteria of an invalid frame as described in 4.2.3. This definition excludes abort sequences from the set of valid frames.

4.1.1 Preamble. The preamble pattern precedes every transmitted frame. Preamble is sent by the MAC sublayer entity as an appropriate number of *pad_idle* symbols. Preamble may be decoded by the receiver as arbitrary data symbols that occur outside frame delimiters. Preamble is primarily used by the receiving modem to acquire signal level and phase lock by using a known pattern. The preamble pattern is chosen for each modulation scheme and data rate for this purpose. The parameter min_post_silence_preamble_length specifies the minimum amount of preamble on the first frame transmitted after a period of "transmitted" *silence*. See 7.1.1 for use of the parameter within the MAC sublayer. See 12.6, 14.7, 16.6, and 18.6 for the specification of the different physical layers.

A secondary purpose for the preamble is to guarantee a minimum ED to SD time period to allow stations to process the frame previously received. The minimum amount of preamble transmitted is a function of both the data rate and the modulation scheme. The standard requires that the duration of the preamble shall be at least 2 μs, regardless of data rate, and that an integral number of octets shall be sent. Thus, at a data rate of 1 Mb/s, one octet of preamble is required to meet the integral number of octet requirements, and at a data rate of 10 Mb/s, three octets are required to meet the minimum time requirement.

The maximum amount of preamble is constrained by the "jabber" control in the PLE. Additionally, for claim_token frames, all stations shall use the minimum number of preamble octets to ensure that all frames are of uniform specified length.

4.1.2 Start Delimiter (SD). The frame structure requires an SD, which begins the frame. The SD consists of signaling patterns that are always distinguishable from data.

The SD is coded as follows (see 12.7, 14.8, 14.11, 16.7, and 18.7 for representations of the symbol coding as present on the medium):

```
                          ┌──── first MAC-symbol transmitted
start delimiter (SD):    | N N 0 N N 0 0 0 |
                           1 2 3 4 5 6 7 8 ← bit positions
```

```
where
  N = non_data MAC-symbol
  0 = zero MAC-symbol
  1 = one MAC-symbol
```

4.1.3 Frame Control (FC) Field. The FC octet determines what class of frame is being sent from among the following general categories:

(1) MAC control
(2) LLC data

The FC format for each of these categories is illustrated by the following:

4.1.3.1 MAC Control Frame

```
      ┌──────── first MAC-symbol transmitted
      │
 ┌──────┬─────────────┐
 │ 0 0  │ C C C C C C │
 └──────┴─────────────┘
   1 2    3 4 5 6 7 8   ← bit positions
```

```
where
  CCCCCC = type of MAC_control frame as follows:

C C C C C C
___________
3 4 5 6 7 8  ← bit positions
___________
0 0 0 0 0 0  claim_token
0 0 0 0 0 1  solicit_successor_1  (has 1 response window)
0 0 0 0 1 0  solicit_successor_2  (has 2 response windows)
0 0 0 0 1 1  who_follows          (has 3 response windows)
0 0 0 1 0 0  resolve_contention   (has 4 response windows)
0 0 1 0 0 0  token
0 0 1 1 0 0  set_successor
```

4.1.3.2 Data frames

first MAC-symbol transmitted

```
      ┌──────── first MAC-symbol transmitted
      │
 ┌──────┬─────────┬─────────┐
 │ F F  │  M M M  │  P P P  │
 └──────┴─────────┴─────────┘
   1 2     3 4 5     6 7 8   ← bit positions
```

```
where
  F F = frame type:
 (1 2 ← bit positions)
  ____
  0 1 = LLC_data_frame
  1 0 = reserved (former management)
  1 1 = reserved (further study)
```

```
 M M M = MAC-action:
(3 4 5 ← bit positions)
 -----
 0 0 0 = request_with_no_response
 0 0 1 = request_with_response  (see 6.6.2)
 0 1 0 = response               (see 6.6.2)

 P P P = priority:
(6 7 8 ← bit positions)
 -----
 1 1 1 = highest priority
 1 1 0
 1 0 1
 1 0 0
 0 1 1
 0 1 0
 0 0 1
 0 0 0 = lowest priority
```

NOTE: The use of the "10" in the frame type field to indicate frames destined for the management entity is not recommended and is included only to retain backwards compatibility with previous versions of this standard.

Other bit patterns in the FC octet are reserved for future study. The action of a station upon receiving an FC value not defined in this standard is not specified.

4.1.4 Address Fields. Each frame shall contain two address fields: The destination address field and the source address field, in that order. Addresses shall be either 16 bits or 48 bits in length. All addresses on a given Local Area Network (LAN) shall be of the same length.

4.1.4.1 Destination Address Field. The following illustration shows the possible representations of destination addresses:

(1) 16-bit address form

```
┌──── first MAC-symbol transmitted

 I/G |   15-bit address   |

                       └──── most significant
                             bit of address
```

(2) 48-bit locally administered form

```
┌──── first MAC-symbol transmitted

 I/G |  1  |          46-bit address          |

                                         └──── most significant
                                               bit of address
```

(3) 48-bit universally administered form

```
┌──── first MAC-symbol transmitted
┌─────┬─────┬───────────────────────────────┐
│ I/G │  0  │        46-bit address         │
└─────┴─────┴───────────────────────────────┘
                                          └──── most significant
                                                bit of address
```

```
where
  I/G = individual/group address bit
```

The first MAC-symbol transmitted of the destination address (the I/G bit) distinguishes individual addresses from group addresses:

0 = individual address
1 = group address

For 48-bit addresses, the second MAC-symbol transmitted of source or destination addresses (the local/universal or L/U bit) distinguishes locally administered addresses from universally administered, unique addresses.

Individual Addresses. An individual address identifies a particular station on the LAN and shall be distinct from all other individual station addresses on the same LAN.

Group Addresses. A group address is used to address a frame to multiple destination stations. Group addresses may be associated with zero, one, or more stations on a given LAN. In particular, a group address is an address associated by convention with a group of logically related stations.

Broadcast Addresses. The group address consisting of all *ones* (that is, 16 or 48 *ones* for two- or six-octet addressing, respectively) shall constitute a broadcast address, denoting the set of all stations on the given LAN.

NOTE: For some of the frame types used by the token-passing bus MAC procedures, the contents of the destination address field are irrelevant. In such cases, the originating station's own address or any other individual addresses may be sent in this field.

Address Administration (48-bit addresses only). There are two methods of administering the set of 48-bit station addresses: locally or through a universal authority. The second bit transmitted of the destination address indicates whether the address has been assigned by a universal or local administrator.

0 = universally administered
1 = locally administered

Universal Administration. With this method, all individual addresses are distinct from the individual addresses of all other LAN stations on a universal basis. The procedure for administration of these addresses is not specified in this standard.

NOTE: Information concerning the Registration Authority and its procedures may be obtained on request to the Secretary General, ISO Central Secretariat, 1 rue de Varembé, Case postale 56, CH1211, Genève, Switzerland/Suisse, quoting the number of this International Standard.

For information on universal address administration contact the Secretary, IEEE Standards Board, Institute of Electrical and Electronics Engineers, Inc., P.O. Box 1331, 445 Hoes Lane, Piscataway, NJ 08855-1331, USA.

Local Administration. Individual station addresses are administered by a local (to the LAN) authority. (This is the only method allowed for 16-bit addresses.)

NOTE: Appendix 4.3 contains a suggested method for hierarchical structuring of locally administered addresses.

4.1.4.2 Source Address Field. The source address identifies the station originating the frame and has the same format and length as the destination address in a given frame, except that the I/G bit shall be set to 0; the significance of it being set to 1 is a subject for future study.

4.1.4.3 Numerical Interpretation of Addresses. Strictly speaking, addresses are bit strings which serve as unique station identifiers or group identifiers. For the purpose of the MAC address comparison within the token bus MAC sublayer, as used in ordering the logical ring and as expressed in the formal access control machine of 7.2.3, each MAC-address bit string is interpreted as if it were an unsigned integer value sent least significant bit first, and thus as if the last bit transmitted had the highest numeric significance.

NOTE: This interpretation does not extend beyond the logical ring ordering operations of the token-passing bus MAC sublayer.

Additionally, the address bits are used in determining delays in the contention process and transmission lengths in the token claiming process. These processes start with the most significant address bits, under the above interpretation, using two bits at a time. Thus the internal processing order is reversed from the serial transmission order on the medium.

4.1.5 MAC Data_Unit Field. Depending on the bit pattern specified in the frame's FC octet, the MAC data_unit field can contain either an LLC protocol data_unit as specified in ISO 8802-2 : 1989 (IEEE Std 802.2-1989) [15], or a value specific to one of the MAC control frames.

In the case where the MAC data_unit field contains a LLC PDU, the MAC data_unit field shall be transmitted to the Physical Layer in the same bit order as received from the LLC sublayer. Likewise, the MAC data_unit field shall be delivered to the LLC sublayer in the same bit order as received from the Physical Layer.

4.1.6 Frame Check Sequence (FCS) Field. The FCS is a 32-bit frame checking sequence, based upon the following standard generator polynomial of degree 32:

$$X^{32} + X^{26} + X^{23} + X^{22} + X^{16} + X^{12} + X^{11} + X^{10} + X^{8} + X^{7} + X^{5} + X^{4} + X^{2} + X + 1$$

The FCS is the one's complement of the sum (modulo 2) of

(1) The remainder of

$$X^{K} * (X^{31} + X^{30} + X^{29} + \ldots + X^{2} + X + 1)$$

divided (modulo 2) by the standard 32-bit generating polynomial, where K is the number of bits in the FC, address (SA and DA), and MAC data_unit fields, and

(2) The remainder of the division (modulo 2) by the standard generator polynomial of the product of X^{32} by the content of the FC, address (SA and DA), and MAC data_unit fields.

The FCS is transmitted commencing with the coefficient of the highest degree term.

As a typical implementation, at the transmitter, the initial content of the register of the device computing the remainder of the division is preset to all ones and is then modified by the generator polynomial (as described above) on the FC, address, and data_unit fields. The one's complement of the resulting remainder is transmitted as the 32-bit FCS.

At the receiver, the initial content of the register of the device computing the remainder is preset to all ones. The incoming serial data bits and the FCS, when divided by the generator polynomial, result, in the absence of transmission errors, in a unique non-zero remainder value. The unique remainder value for the 32-bit FCS is the polynomial:

$$X^{31} + X^{30} + X^{26} + X^{25} + X^{24} + X^{18} + X^{15} + X^{14} + X^{12}$$
$$+ X^{11} + X^{10} + X^{8} + X^{6} + X^{5} + X^{4} + X^{3} + X + 1$$

NOTES: (1) To test the FCS generation and checking logic in a station, an implementation should provide a means of bypassing the FCS generation circuitry and providing an FCS from an external source. The ability to pass frames that have FCS errors along with the received FCS value and an error indication, to higher levels of the protocol, is another desirable testability feature.

(2) The FCS polynomial provides a Hamming distance of four in the frame transmission and frame detection process, provided that the total frame length, between SD and ED exclusive, is less than 11454 octets. The maximum data frame length, 8191, is 11454 rounded down to the nearest power of two (minus one).

4.1.7 End Delimiter (ED). The frame structure requires an ED, which ends the frame and determines the position of the FCS. The data between the SD and the ED shall be an integral number of octets. All bits between the SDs and EDs are covered by the FCS.

The ED consists of signaling patterns that are always distinguishable from data. The ED also contains bits of information that are not error checked.

The ED is coded as follows:

```
                        ┌──── first MAC-symbol transmitted

end delimiter (ED):   | N N 1 N N 1 I E |

                        1 2 3 4 5 6 7 8 ← bit positions
where
  N = non_data MAC-symbol
  1 = one MAC-symbol
  I = intermediate bit  (1 = more to transmit,
                         0 = end of transmission)
  E = error bit         (0 = no error, 1 = error)
```

The seventh ED MAC-symbol is called the intermediate bit (or “I” bit). If *one*, it indicates that more transmissions from the station follow. If *zero*, it indicates that this is the last frame transmitted by the station and silence follows the ED. The I bit assists a repeater and possibly the Physical Layer in determining what follows the ED.

The eighth ED MAC-symbol is called the error bit (or “E” bit). When set to *one* by a repeater, the error bit indicates that this frame had an FCS error within the frame and the error did not occur on the communication path between the immediately preceding sending station/repeater and this receiving station. When the error bit is set to *one*, the receiving station may treat the frame as an invalid frame. The originating station shall always transmit the error bit as *zero*.

4.1.8 Abort Sequence. The abort sequence is used to prematurely terminate the transmission of a frame. The abort sequence shall be transmitted on an octet boundary relative to that part of the aborted frame already transmitted.

An abort sequence consists of the following pattern:

first MAC-symbol transmitted

N N 0 N N 0 0 0	N N 1 N N 1 I E

An abort sequence shall also be transmitted by a repeater upon receiving an invalid coding sequence.

4.2 Enumeration of Frame Types. This subsection shows how the components of the frames are arranged in the various frame types transmitted by the MAC sublayer. Section 5 discusses the frames and terminology used here.

4.2.1 MAC Control Frame Formats. The following frames are sent and received by the MAC sublayer and are not passed to higher layers.

4.2.1.1 Claim_token. The frame has a data_unit whose value is arbitrary and whose length in octets is 0, 2, 4, or 6 times the system’s slot_time also measured in octets.

PRE	SD	00000000	DA	SA	arbitrary value, length = (0,2,4,6) * slot_time octets	FCS	ED

4.2.1.2 Solicit_successor_1. The frame has a DA equal to the value of the station’s NS and a null data_unit. One response window always follows this frame.

PRE	SD	00000001	DA	SA	FCS	ED	...

one response window

4.2.1.3 Solicit_successor_2. The frame has DA equal to the value of the station's NS or TS and a null data_unit. Two response windows always follow this frame.

PRE	SD	00000010	DA	SA	FCS	ED

two response windows

4.2.1.4 Who_follows. The frame has a data_unit equal to the value of the station's NS. The format and length of the data_unit is the same as a source address. Three response windows always follow this frame. (This gives receivers two extra slot_times to make a comparison with an address other than TS.).

PRE	SD	00000011	DA	SA	value NS	FCS	ED

three response windows

4.2.1.5 Resolve_contention. The frame has a null data_unit. Four response windows always follow this frame.

PRE	SD	00000100	DA	SA	FCS	ED

four response windows

4.2.1.6 Token. The frame has DA equal to the value of the station's NS, and has a null data_unit.

PRE	SD	00001000	DA	SA	FCS	ED

4.2.1.7 Set_successor. The frame has data_unit equal to the value of the station's NS or TS. The format and length of the data_unit are the same as that of a source address.

PRE	SD	00001100	DA	SA	new value of NS	FCS	ED

4.2.2 LLC Data Frame Format. LLC data frames have a DA and data_unit specified by a station's LLC sublayer. A frame of this type with a non-null data_unit shall be passed to the receiving station's LLC sublayer (see 4.1.3.2).

PRE	SD	01MMMPPP	DA	SA	LLC_data_unit	FCS	ED

4.2.3 Invalid Frames. An invalid frame is defined as one which meets at least one of the following conditions:

(1) It is identified as such by the PLE. (For example, it contains *bad_signal* symbols.)
(2) It is not an integral number of octets in length.
(3) It does not consist of an SD, one FC field, two properly formed address fields, one data_unit field of appropriate length (dependent on the bit pattern specified in the FC field), one FCS field, and an ED, in that order.
(4) The FCS computation, when applied to all octets between the SD and the ED, fails to yield the unique remainder specified in 4.1.6.

An implementation may also include any of the following additional conditions for an invalid frame:

(5) The FC field contains an undefined bit pattern.
(6) The error bit within the ED is set to *one*.

Invalid frames shall be treated as noise. Their existence, as noise bursts, is relevant at some points in the token bus elements of procedure.

Appendix

(This Appendix is not a part of this International Standard or of IEEE Std 802.4-1990, but is included for information only.)

4.3 Appendix—Recommendation for a Hierarchical Structure for Locally Administered Addresses. The concepts introduced in this Appendix are under study for inclusion in a future revision of this standard.

This Appendix describes an addressing structure for a bus network divided into multiple logical segments, with one or more MAC bridge stations interconnecting the segments. Structuring MAC addresses in a hierarchical fashion can facilitate the operation of these bridge stations.

A logical segment is defined as the collection of all stations of a LAN that have the same segment address and that can exchange frames without any intermediary MAC bridge entity. Stations of a segment can communicate with stations with different segment addresses only through a MAC bridge or some other intermediary. For example, for a token bus implementation, the stations of a logical segment share a single token forming a single logical ring distinct from other logical segments.

A hierarchical address permits a MAC bridge station to recognize frames that require forwarding to other logical segments.

The source and destination address partitioning recommended for this purpose is

(1) 16-bit hierarchical form

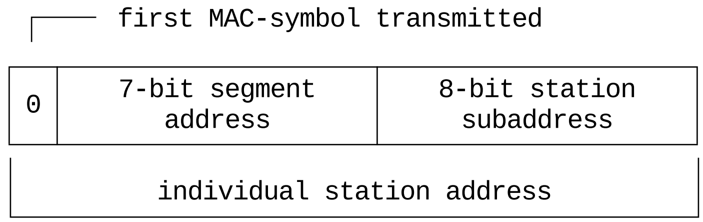

(2) 48-bit locally administered hierarchial form

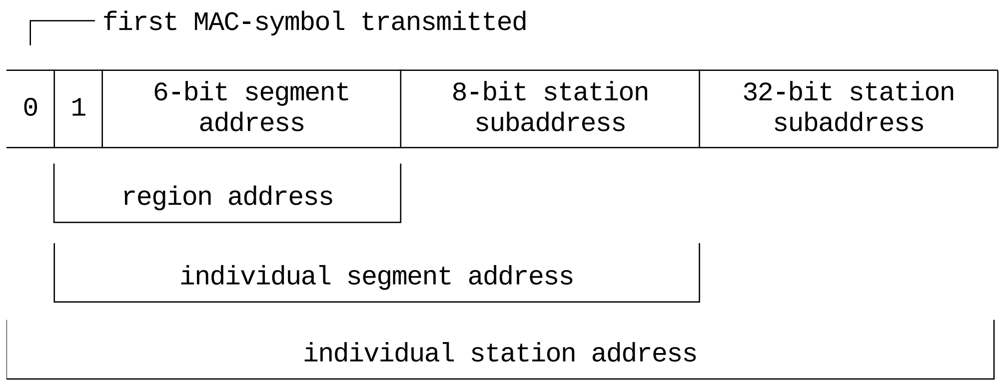

The 8-bit or 32-bit station subaddress distinguishes stations on a single logical segment. The 8-bit segment subaddress distinguishes logical segments within a single region defined by the 6-bit region address.

5. Elements of MAC Sublayer Operation

A description of the token-passing bus MAC mechanism is provided in this section. This section is intended to assist the reader in understanding the MAC sublayer and its operation.

Section 6 contains precise definitions of MAC-specific terms and mandatory aspects of the mechanism. Where statements included in this section conflict with those in Section 6, or are incomplete, those in Section 6 shall take precedence.

Section 7 describes the required behavior of the access control machine (ACM) of the MAC sublayer. Where statements in this section or Section 6 conflict with the formal description in Section 7, or are incomplete, the formal description in Section 7 shall take precedence.

This section describes the token-passing bus MAC sublayer's operational and exception recovery functions. The relationship of this section to other sections of this standard and to LAN specifications is illustrated in Fig 5-1.

Fig 5-1
Relation to LAN Model

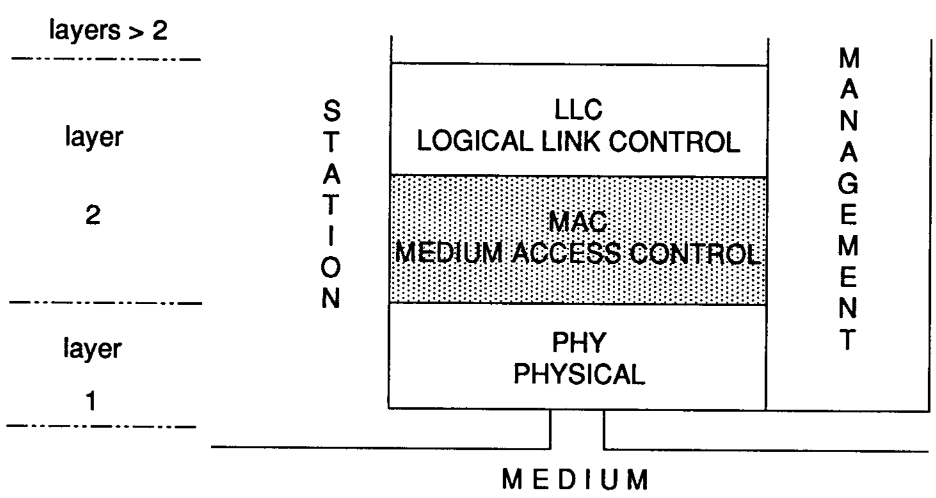

Specific responsibilities of the MAC sublayer for a broadcast medium involve managing ordered access to the medium, providing a means for admission and deletion of stations (adjustment of logical ring membership), and handling fault recovery.

The faults considered here are those caused by communications errors or station failures. These faults include

(1) Multiple tokens
(2) Lost token
(3) Token pass failure
(4) "Deaf" station (that is, a station with an inoperative receiver)
(5) Duplicate station addresses
(6) Faulty transmitter

This MAC protocol is intended to be robust, in the sense that it should tolerate and survive multiple concurrent errors.

Some basic observations are useful in understanding the operation of token passing on a broadcast medium:

(a) Stations are connected in parallel to the medium. Thus, when a station transmits, its signal is received (or "heard") by all stations on the medium. Other stations can interfere with the first station's transmission but can not predictably alter its contents.
(b) When a station transmits, it may assume that all other stations hear something (though not necessarily what was transmitted).
(c) When a station receives a valid frame (properly formed and delimited and containing a correct FCS), it may infer that some station transmitted the frame, and therefore, that all stations heard something.
(d) When a station receives something other than a valid frame (that is, noise), it may make no inference about what the other stations on the medium might have heard.
(e) Not all stations need be involved in token passing (only those which desire to initiate transmissions).
(f) Multiple tokens and lost tokens may be detected by any station. There are no special "monitor" stations required to perform token recovery functions.
(g) Due to spatial separation, stations cannot be guaranteed to have a common perception of the system state at any instant. (The MAC protocol described herein accounts for this.)

5.1 Basic Operation. Steady-state operation (the network condition where a logical ring has been established and no error conditions are present) simply requires the sending of the token to a specific successor station as each station is finished transmitting (see Fig 1-2).

Other essential and more difficult tasks are establishment of the logical ring (at either initialization or reestablishment in the case of a catastrophic error), and maintenance of the logical ring (allowing stations to enter and leave the logical ring without disrupting the other stations in the network).

The token passes the right to transmit among all stations in the logical ring. Each participating station knows the address of its predecessor (the station from which it received the token), referred to as Previous Station or PS. It knows its successor (which station the token should be sent to next), referred to as Next Station or NS. It knows its own address, referred to as This Station or TS. These predecessor and successor addresses are dynamically determined and maintained by the algorithms described. Whenever a station changes its successor (NS), an indication of this change is passed to the management entity (if enabled).

The following subsections introduce major elements and features of the token bus access protocol.

5.1.1 Slot_time. In describing the access operations, the term slot_time is used to refer to the maximum time any station need wait for an immediate MAC level response from another station. Slot_time is precisely defined in 6.1.9.

The slot_time (along with the station's address and several other management parameters) shall be known to the station before it attempts to transmit on the network. If all stations in a network are not using the same value for slot_time, the MAC protocol may not operate properly. The method of setting these parameters in each station is outside the scope of this standard.

5.1.2 Right to Transmit. The token (right to transmit) is passed from station to station in descending numerical order of station address. When a station which belongs to the logical ring hears a token frame addressed to itself, it "has the token" and may transmit data frames. When a station has completed transmitting data frames it passes the token to the next station in the logical ring, as discussed in 5.1.3.

When a station has the token it may temporarily delegate to another station its right to transmit by sending a request_with_response data frame. When a station hears a request_with_response data frame addressed to itself it shall respond with a response data frame, if the request_with_response option is implemented. The response data frame causes the right to transmit to revert back to the station which sent the request_with_response data frame.

5.1.3 Token Passing. After each station has completed transmitting any data frames it may have and has completed other maintenance functions (described in 5.1.4), the station passes the token to its successor by sending a "token" MAC_ control frame.

After sending the token frame the station listens for evidence that its successor has heard the token frame and is active. If the sender hears a valid frame following the token within one slot_time, it assumes that its successor has the token and is transmitting. Otherwise the token sending station attempts to assess the state of the network.

If the token sending station hears a noise_burst (e.g., an unidentifiable sequence or a frame with an incorrect FCS), it cannot be sure which station sent the transmission. The MAC protocol treats this condition in a way which minimizes the chance of the station causing a serious error.

Since a station on a broadband network should always hear its own frames, if the token sending station hears a single noise_burst without hearing its own token frame, the station assumes that it heard its own token that had been garbled and continues to listen.

If a second noise_burst is heard or if a noise_burst is heard after the station hears its own token frame, the token sending station continues to listen in the CHECK_TOKEN_PASS state for up to four more slot_times. If nothing more is heard, the station assumes that the noise_burst it heard was not a garbled frame from the successor station and so it repeats the token transmission. If anything is heard during the four slot_time delay, the station assumes its successor successfully received the token.

If the token holder does not hear a valid frame after sending the token the first time, it repeats the token pass operation once, performing the same monitoring as during the first attempt.

If the successor does not transmit after a second token frame, the sender assumes that its successor has failed. The sender then sends a who_follows frame with its successor's address in the data_unit field of the frame. All stations compare the value of the data_unit field of a who_follows frame with the address of their predecessor (the station that normally sends them the token). The station whose predecessor is the successor of the sending station responds to the who_follows frame by sending its address in a set_successor frame. The station holding the token thus establishes a new successor, bridging the failed station out of the logical ring.

If the sending station hears no response to a who_follows frame, it sends the frame a second time. If there is still no response, the station tries another strategy to reestablish the logical ring. The station now sends a solicit_successor_2 frame with its own address as both DA and SA, asking any station in the system to respond to it. Any operational station that hears the request and needs to be part of the logical ring responds, and the logical ring is reestablished using the response window process discussed in 5.1.4.

If all attempts at soliciting a successor fail, the station assumes that a fault may have occurred; either all other stations have failed, all stations have left the logical ring, the medium has broken, the station's own transmitter is malfunctioning, or the station's own receiver has failed so that it cannot hear other stations who have been responding to its requests. Under such conditions the station quits attempting to maintain the logical ring. If the station has no frames to send, it listens for some indication of activity from other stations. If the station has data frames to send, it sends its remaining data frames and then repeats the token pass process. Once the station has sent its frames and still cannot locate a successor, it becomes silent and listens for another station's transmissions.

In summary, the token is normally passed from station to station using a short token frame. If a station fails to pick up the token, the sending station uses a series of recovery procedures that grow increasingly more drastic as the station repeatedly fails to find a successor station.

5.1.4 Response Windows. New stations are added to the logical ring through a controlled contention process using "response windows." A response window is a controlled interval of time (equal to one slot_time) after transmission of a MAC_control frame in which the station sending the frame pauses and listens for a response. If the station hears a transmission start during the response window, the station continues to listen to the transmission, even after the response window time expires, until the transmission is complete. Thus, the response windows define the time interval during which a station will hear the beginning of a response from another station.

The two frame types, solicit_successor_1 and solicit_successor_2, indicate the opening of response windows for stations wishing to enter the logical ring. The solicit_successor frame specifies a range of station addresses between the frame source and destination addresses. Stations whose addresses fall within this range and who wish to enter the logical ring respond to the frame.

The sender of a solicit_successor frame transmits the frame and then waits, listening for a response in the response window(s) following the frame. Responding stations send the frame sender their requests to become the next station in the logical ring. If the frame sender hears a valid request, it allows the new station to enter the logical ring by changing the address of its successor to the new station and passing its new successor the token.

In any response window there exists the possibility that more than one station simultaneously will desire logical ring entry. To minimize contention when this happens, the ring entry sequence is limited by requiring that a station request admission only when a window is opened for an address range that spans its address.

There are two solicit_successor frames. Solicit_successor_1 has one response window. Solicit_successor_2 has two response windows. Solicit_successor_1 is sent when the station's successor's address is less than the station's address. This is the normal case when the token is being passed from higher to lower addressed stations. Solicit_successor_1 allows only stations whose address is in the range between the token sender and its expected next token holder to respond, thus limiting the possible contenders and preserving the descending order of the logical ring.

Exactly one station in the logical ring has its station's address below that of its successor (that is, the unique station having the lowest address which sends the token to the "top" of the address ordered logical ring). When soliciting successors, this station shall open two response windows, one for those stations with addresses below its own, and one for stations with addresses above its successor. The station with the lowest address sends a solicit_successor_2 frame when opening response windows. Stations having an address below the sender respond in the first response window, while stations having an address higher than the sender's successor respond in the second response window.

In any response window, when the soliciting station hears a valid set_successor frame, it has found a new successor. When multiple stations simultaneously respond, only unrecognizable noise may be heard during the response period. The soliciting station then sequences through an arbitration algorithm to identify a single responder by sending a resolve_contention frame. The stations that had responded to the earlier solicit_successor frame and that have not yet been eliminated by the algorithm, choose a two-bit *listen delay* value (usually from the station's address) and listen for 0, 1, 2, or 3 slot_times as determined by that listen delay value. (This listen delay value is further described later.) If these contending stations hear anything (that is, non-silence) while listening, they eliminate themselves from the arbitration. If they hear only silence, they continue to respond to further resolve_contention requests from the soliciting station. This process normally leads to the establishment of a single new successor.

5.1.5 Bound on Token Rotation Time. The worst-case token rotation time determines the maximum delay a station experiences in gaining access to the network. A very long token rotation occurs if many stations attempt to add new stations to the logical ring by performing the solicit successor procedure on the same rotation of the token.

The maximum token rotation time is bounded by deferring the solicit successor procedure whenever the token rotation time is longer than a defined maximum. The timer token_rotation_timer (ring_maintenance) in each station measures the rotation time of the token. If the time exceeds a value set by the management entity, target_rotation_time (ring_maintenance), the station does not perform the solicit successor procedure. On the next token rotation, the station determines if the token now is rotating fast enough to perform the solicit successor procedure.

In conjunction with the priority option (see 5.1.8), the value of target_rotation_time (ring_maintenance) determines the maximum token rotation time.

5.1.6 Initialization. Initialization is primarily a special case of adding new stations; it is triggered by the exhaustion of an inactivity timer (bus_idle_timer) in one station. If the inactivity timer expires, the station sends a claim_token frame. As in the response window algorithm, the initialization algorithm assumes that more than one station can try to initialize the network at a given instant. This is resolved by address sorting the initializers.

Each potential initializer sends a claim_token frame having an data_unit field length that is a multiple of the network slot_time (the multiple being 0, 2, 4, or 6 based on selected bits of the station address). Each initializing station then waits one slot_time for its own transmission, and those of other stations that chose the same frame length, to pass. The station then samples the state of the medium.

If a station senses non-silence, it knows that some other station(s) sent a longer length transmission. The station defers to those stations with the longer transmission and re-enters the IDLE state.

If silence was detected, and unused bits remain in the address string, the station attempting initialization repeats the process using the next two bits of its address to derive the length of the next transmitted frame. If all bits have been used and silence is still sensed, the station has "won" the initialization contest and now holds the token.

Once there is a unique token in the network, the logical ring builds by way of the response window process previously described.

NOTE: A random pair of bits is used at the end of the address sort algorithm to ensure that two stations with the same address (which is a fault condition) will not permanently bring down an entire system. If the two stations do not separate (random choices identical), they both attempt to form a logical ring and, at most, one of them should succeed. If they do separate (random choices are different), one will defer. In the latter case, the station that defers will hear a transmission from a station with an identical address and so will discover the error condition.

5.1.7 Exiting Ring. A station can remove itself from the logical ring at any time by simply choosing not to respond to a token passed to it, allowing the fault recovery mechanisms in the MAC protocol to patch it out. A more efficient method is the following: when the station has the token and desires to exit the logical ring, the station sends a set_successor frame to its predecessor (the station that transmitted the token to it) containing the address of its successor. The exiting station then simply sends the token as usual to its successor. Readmission to the logical ring requires one of the sequences described in 5.1.4 and 5.1.6.

5.1.8 Priority Option. The token-passing bus access method provides an optional priority mechanism by which higher layer data frames awaiting transmission are assigned to different "service classes," ranked or ordered by their desired transmission priority. This priority mechanism allows the MAC sublayer to provide eight service classes to the LLC sublayer, and higher level protocols. The priority of each frame is determined by the service class requested in the MA-UNITDATA request command (see 2.2.1).

The token-passing bus access method distinguishes only four levels of priority, called "access classes." Thus there are four request queues to store frames pending transmission. The access classes are named 0, 2, 4, and 6, with 6 being the highest priority and 0 the lowest.

The MAC sublayer maps the service class requested by the LLC sublayer into a three bit priority value, which is included in the FC field. The priority value is then mapped into the MAC access class by ignoring the least significant bit of the priority field. Thus service classes 0 and 1 correspond to access class 0, service classes 2 and 3 to access class 2, service classes 4 and 5 to access class 4, and service classes 6 and 7 to access class 6.

Any station not using the optional priority feature shall transmit every data frame with an access class of 6 (the highest priority). The service class value in the MA-UNITDATA request shall still be carried in the FC octet. For all stations, the rule governing the transmission of these highest priority frames is that a station shall not transmit consecutive frames for more than some maximum time set by the management entity. This time, called the hi_pri_token_hold_time, prevents any single station from monopolizing the network. If a station has more data frames to send at access class 6 than it can transmit in one hi_pri_token_hold_time period, it is prohibited from sending additional frames after that time has expired and shall pass the token. However, when a station that is using the optional priority feature has lower access class frames to send and has time available, it shall send these frames subject to the priority system rules described in these paragraphs.

The object of the priority system is to allocate network bandwidth to the higher priority frames and only send lower priority frames when there is sufficient band width. The network bandwidth is allocated by timing the rotation of the token around the logical ring. Each access class is assigned a "target" token rotation time. For each access class the station measures the time it takes the token to circulate around the logical ring. If the token returns to a station in less than the target rotation time, the station can send frames of that access class until the target rotation time has expired. If the token returns after the target rotation time has been reached, the station cannot send frames of that priority on this pass of the token.

Each station using the optional priority scheme shall have three token rotation timers for the three lower access classes. Each access class has a queue of frames to be transmitted. When a station receives the token it first services the highest access class queue, which uses the hi_pri_token_hold_time to control its operation. After having sent any frames of the highest priority, the station begins to service the token rotation timers and queues working from higher to lower access class.

Each access class acts as a virtual substation in that the right to transmit is passed internally from the highest access class downward, through all access classes, before being passed to the station's successor.

The access class service algorithm consists of loading the residual value from a token rotation timer into the token hold timer and reloading the same token rotation timer with the target_rotation_time for that access class. (Thus frames sent by a station, for this access class, are accounted for in the access class's next token rotation time computation.) If the token hold timer has a remaining positive value, the station can transmit frames at this access class until either the token hold timer times out or this access class's queue is empty. When either event occurs the station begins to service the next lower access class. When the lowest level is serviced the station performs any required logical ring maintenance and passes the token to its successor.

5.1.9 Token Rotation Timer Example. The following example describes a simplified system with only two active access classes, the highest and one other lower level. Assume there are four stations in the logical ring, with addresses 10, 8, 6, and 2. Assume stations 10 and 2 send high priority frames every token revolution; while 8 and 6 send as many lower priority frames as possible. (The example assumes that propagation delays are negligible compared with the time to send data frames.)

The example begins after a period during which no data frames have been sent so that the token has been rotating as rapidly as possible. In Fig 5-2, each station is represented by a box that contains the station's address and target token rotation time value for the station's active frame queue. The values in the left-hand column under each station, labeled TRTC, are the token rotation times observed at that station for the previous rotation of the token. The values in the right-hand column, labeled XMIT, are the number of frames transmitted by the station each time it holds the token. Each row of this table represents one rotation of the token. In this example, station 10 is assumed to have three high priority frames of 128 octets, stations 8 and 6 have only low priority frames of 400 and 356 octets, respectively, and station 2 has two high priority frame of 305 octets. The token rotation time is given in octet_times and assumes 19 octet_times per token pass per station.

Looking at the example we see that on the first token rotation, station 10 measures an effective rotation time of 76 (4*19) and sends three frames of high priority data. Station 8 thus measures an effective rotation time of 460 (3*128+76), leaving 1140 units of "time" to send data before having to pass the token to sta-

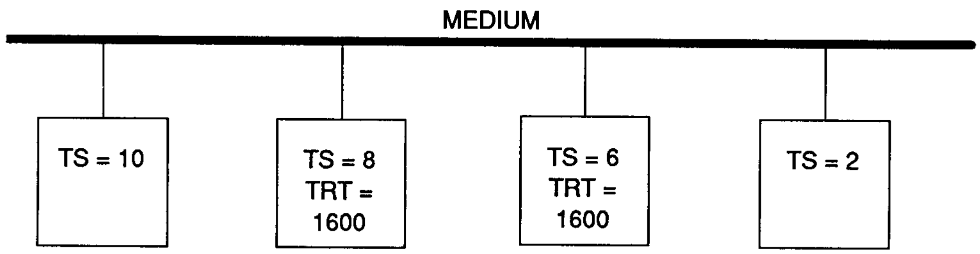

Token Rotation	TRTC	XMIT	TRTC	XMIT	TRTC	XMIT	TRTC	XMIT
1	76	3	460	3	1660	0	1660	2
2	2270	3	2270	0	1070	2	1782	2
3	1782	3	1782	0	1782	0	1070	2
4	1070	3	1070	2	1870	0	1870	2
5	1870	3	1870	0	1070	2	1782	2
6	1782	3	1782	0	1782	0	1070	2
7	1070	3	1070	2	1870	0	1870	2
8	1870	3	1870	0	1070	2	1782	1
9	1477	3	1477	1	1877	0	1165	1
10	1165	3	1165	2	1565	1	1921	1
11	1921	3	1921	0	1121	2	1477	1
12	1477	3	1477	1	1877	0	1165	1

Fig 5-2
Token Rotation Time "Priority" Example

tion 6. Station 6 measures an effective token rotation time of 1660 (since 1660 octets of data have been sent since it last received the token). Station 6 cannot send low priority data and immediately passes the token. (If station 8 had less data to send, it could have passed the token to station 6 sooner, giving station 6 an opportunity to send on this rotation.) Station 2 sends two frames of high priority data, unconstrained by the lower priority rotation timer.

On the second rotation of the token station 8 cannot send, but the token arrives early enough for station 6 to send two frames of low priority data.

On the third rotation of the token in the example, stations 10 and 2 continue to send high priority data frames. Neither stations 8 nor 6 can send low priority data since the token is rotating too slowly on this circuit.

On the fourth and fifth rotations, stations 10 and 2 continue to use most of the network bandwidth for high priority frames. Stations 8 and 6 share the remaining bandwidth by each sending two frames on alternate token rotations.

Rotations 5 through 7 repeat the usage pattern shown in rotations 2 through 4, showing a stable cyclic bandwidth allocation pattern: over any three rotations, stations 10 and 2 use 66% of the bandwidth and stations 6 and 8 share the remaining 34%. If all stations always offer the same traffic, this sharing continues indefinitely.

On the eighth rotation of the token, station 2 begins to send only one frame of high priority data, allowing the lower priority stations more of the available network bandwidth. That a stable, cyclic usage pattern is again established can be seen by comparing token rotations 9 and 12. This usage pattern again repeats over three rotations. Stations 10 and 2 together use 48% of the bandwidth, and stations 6 and 8 share the remaining 52%.

This example demonstrates how the token rotation timer method allows lower access class traffic to fill in bandwidth that is not used by higher access class traffic. Note how the network bandwidth not used by stations with higher access class traffic is more or less equally shared among stations with lower access class traffic. Although in this example each station only sends traffic of one access class, this standard allows any station to send at any access class.

5.1.10 MAC-Acknowledged Connectionless Class of Service. The request_with_response mechanism, coupled with appropriate MAC-user functions, provides a MAC-acknowledged connectionless class of service. When a higher layer entity requests a data transmission using the MAC-acknowledged connectionless class of service, the MAC-user entity issues to the MAC entity a MA-UNITDATA request with a quality-of-service parameter specifying request_with_response.

When the local MAC entity obtains the token, it sends the request_with_response frame and waits for another valid frame. If a timer expires without receiving a valid frame, the local MAC entity retransmits the original frame. This sequence is repeated until a valid frame is received or the allowed number of retries is exhausted.

When the remote (responding) MAC entity receives the frame specifying request_with_response, it passes the received frame to the remote MAC-user entity. That remote MAC-user entity generates an appropriate response and directs the remote MAC entity to send this response frame at once to the source of the request_with_response frame (that is, the originating station).

When the local (originating) MAC entity receives the response frame, or when the retry count is exhausted, it associates the frame (if available) with the request_with_response frame now being processed and notifies the originating MAC-user entity of the completion of its original request.

The retry mechanism of the request_with_response procedure prevents loss of frames at any arbitrary level of confidence determined by the max_retry_limit parameter (see 7.1.1). However, since the local station repeats the request_with_response frame when a response frame is not received, it is possible for the remote station to receive duplicates of the original request_with_response frame. The remote MAC-user entity should eliminate duplicate frames.

The local (originating) MAC entity engages in one request_with_response activity at a time. All retries and timeouts for that request are completed before the local MAC-entity processes another request or initiates the ring-maintenance/token-passing procedures.

5.1.11 Randomized Variables. Several of the variables used by the MAC protocol have two bit "random" values. Some of these randomized variables are used to improve error recovery probabilities under certain conditions. The randomization of max_inter_solicit count causes stations to operate "out of step" when opening response windows.

5.2 Access Control Machine (ACM) States. The MAC logic in a station is described here as a computation machine that sequences through a number of distinct phases, called states. These states are introduced in the following subsections. The states and the transitions between them are illustrated in Fig 5-3. (The dashed lines group states into functional areas.) Section 7 contains the complete state transition table which provides a formal description of the token-passing bus ACM.

5.2.0 Offline. OFFLINE is the state the ACM is in immediately following power-up or following the detection, by the MAC sublayer, of certain fault conditions. After powering up, a station tests itself and its connection to the medium without transmitting on the medium. This "internal" self-testing is station-implementation-dependent and does not affect other stations on the network. Thus, the self-test procedure is beyond the scope of this standard.

After completing any power-up procedures, the station remains in the OFFLINE state until it has had all necessary internal parameters initialized and been instructed to go online.

Fig 5-3
MAC Finite State Machine Diagram

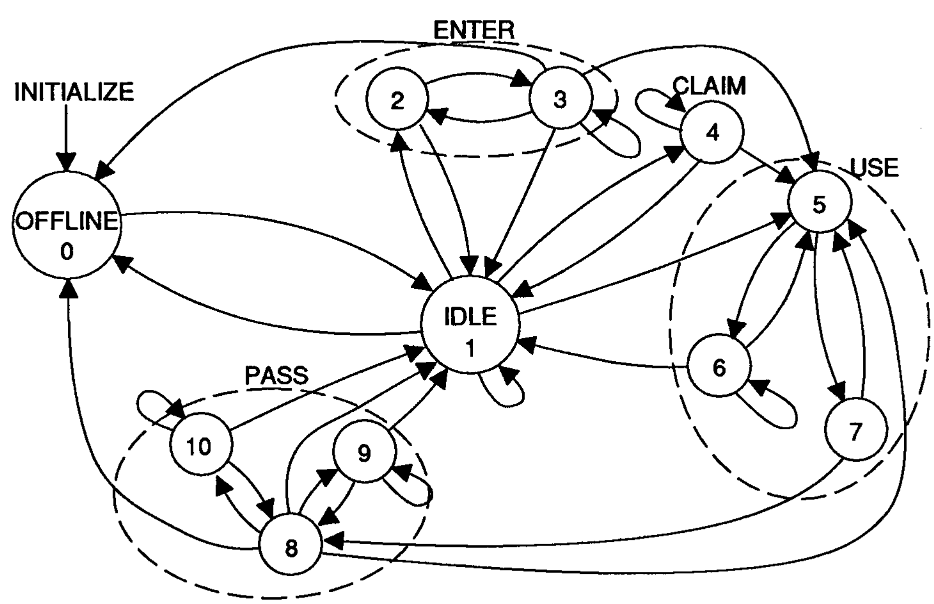

5.2.1 Idle. IDLE is the normal “quiescent” ACM state.

If a MAC_control frame is received for which the station shall take action, the appropriate state is entered. For example, if a token frame addressed to the station is received, the station enters the USE_TOKEN state.

If the station goes for a long period of time (a defined multiple of the slot_time) without hearing any activity on the medium, it may infer that recovery of the logical ring is necessary. The station attempts to claim the token (enters the CLAIM_TOKEN state) and (re)initialize the logical ring.

5.2.2 Demand In. The DEMAND_IN state is entered from the IDLE state if a solicit successor frame that spans the station's address is received by a station desiring logical ring entry. (The DEMAND_IN state is also entered from the DEMAND_DELAY state during the contention resolution process discussed in 5.2.3.) In the DEMAND_IN state the contending station sends the token holder a set_successor frame and goes to the DEMAND_DELAY state to await a response.

If a station intends to respond to a solicit_successor or a who_follows frame in the first response window, the station enters the DEMAND_IN state from the IDLE state with a zero delay and so immediately transmits a set successor response and goes to the DEMAND_DELAY state. If the station intends to respond in the second or later response window after a frame, the station delays in the DEMAND_IN state before transmitting a set_successor frame.

While delaying in the DEMAND_IN state, if the station hears any transmissions, it assumes that another station with a higher numbered address is requesting the token and so it shall return to the IDLE state.

5.2.3 Demand Delay. DEMAND_DELAY is the state the station enters after having sent a set_successor frame in the DEMAND_IN state. In the DEMAND_DELAY state a station can expect to hear

(1) A token from the token holder addressed to the station, indicating that the station's set_successor frame was heard,
(2) A token from the token holder addressed to another station, indicating that another station's set_successor frame was heard,
(3) A resolve_contention frame from the token holder indicating that all stations that are still demanding to be added to the logical ring should perform another step of the contention resolution process, or
(4) Set_successor frames from other stations, which the station ignores.

If the station either hears nothing or hears a frame other than one of (1), (3), or (4), the station shall leave the DEMAND_DELAY state. The station then abandons soliciting the token and returns to the IDLE state.

In (1), the token holder has heard the soliciting station and sent the token. The contention resolution process is over. The soliciting station, upon receiving the token, goes to the USE_TOKEN state and begins transmitting.

In (2), the token holder has heard another soliciting station and sent it the token. The contention resolution process is over. The station abandons soliciting the token and returns to the IDLE state.

In (3), the token holder has heard responses from multiple stations soliciting the token and sent a resolve_contention frame. All stations currently in the DEMAND_DELAY state respond to this frame. The responding stations first set a delay and return to the DEMAND_IN state where they listen for other demanders. If no other demanders have been heard by the time the delay period has expired, each station then sends another set_successor frame to the token holder.

The contention for the token by multiple stations is resolved by having each station delay a period of time in the DEMAND_IN state before transmitting another set_successor frame. The delay interval is chosen by taking the station's (unique) address and using two bits from that address to determine the delay interval. The first pass of the resolution process uses the most significant two address bits; the next pass the next two address bits; etc. Thus stations delay 0, 1, 2, or 3 slot_times when entering the DEMAND_IN state before transmitting.

When multiple stations request to enter the logical ring, the desired result is for the token holding station to pass the token to the highest addressed station. To select the highest addressed contending station from multiple contenders, the one's complement of the station's address is used to determine the delay in the DEMAND_IN state. Thus, stations with numerically higher addresses delay shorter intervals and send their set_successor messages sooner than stations with lower addresses. Stations with numerically lower addresses hear the transmissions of the higher addressed stations and drop out of the contention process.

If two contending stations have the same value for the selected two address bits, they delay the same amount of time and transmit more or less simultaneously. If the token holder hears multiple responses and does not hear a valid set_successor from any station, it sends another resolve_contention frame, starting another step of the contention resolution process.

The contention resolution process can take at most 25 passes (48/2+1, for 48-bit addresses) of the following cycle:

(a) All remaining demanding stations send set_successor frames to the token holder.
(b) They all listen for a response from the token holder and ignore other set_successor frames.
(c) They all hear a resolve_contention frame from the token holder.
(d) They all delay a number of slot_times based on the next two bits of their own addresses.
(e) If they hear another frame during the delay, they drop from contention.

The contention resolution process should resolve so that the contending station with the highest address is heard by the token holder and receives the token. However, if two stations are erroneously assigned the same station address, they will both sequence through the contention process using the same delays and may not resolve.

To permit eventual resolution in this error condition, a final resolution pass is taken, using a two-bit random number, after the station's address bits have all been used and the contention is still unresolved. If both stations choose the same random value or another error prevents resolution, then the token holder and the contending stations abandon the resolution process until the next response window is opened. (Thus, two stations with the same address that consistently choose the same random number value may never be able to enter the logical ring.)

5.2.4 Claim Token. CLAIM_TOKEN is entered from the IDLE state after the inactivity timer expires (and the station desires to be included in the logical ring). In this state, the station attempts to initialize or reinitialize the logical ring by sending claim_token frames.

To resolve multiple simultaneous stations sending claim_token frames, each station delays for one slot_time after sending the claim_token frame, and then monitors the medium as previously described. If after this delay the bus is quiet, the station sends another claim_token frame.

If the station sends max_pass_count (where max_pass_count is half the number of bits in the station's address plus one, since the address bits are used in pairs) claim_token frames without hearing other transmissions, the station has successfully "claimed" the token, and goes to the USE_TOKEN state.

5.2.5 Use Token. USE_TOKEN is entered after just receiving or claiming a token. This is the state in which a station can send data frames.

As the station enters this state, it starts the token hold timer, which limits how long the station may transmit before proceeding to the next access class or passing the token. The value initially loaded in the token holding timer, hi_pri_token_hold_time, is a system imposed parameter.

When a station's token holding timer has expired and any transmission in progress is complete, or when the station has no more to transmit, it enters the CHECK_ACCESS_CLASS state.

When a station sends a data frame it sets the response window timer to 3 slot_times and enters the AWAIT_IFM_RESPONSE state.

5.2.6 Await IFM Response. AWAIT_IFM_RESPONSE is entered when a data frame has been sent. The ACM may wait for the interface machine (IFM) to signal the reception of a response.

If the frame sent in the USE_TOKEN state was a request_with_no_response frame, no response is expected. The USE_TOKEN state is again entered to check for another frame or token holding timer timeout. If the frame sent was a request_with_response frame, the station waits in the AWAIT_IFM_RESPONSE state for one of the following:

(1) A response frame addressed to the requestor
(2) Any other valid frame
(3) A timeout

If a response frame addressed to the requestor is heard, the USE_TOKEN state is again entered to check for another frame or token holding timer timeout. (The IFM passes the response frame to the specified MAC-user entity as it does all other data frames addressed to the station. The IFM also associates the response frame with the request_with_response frame just previously transmitted.)

If any other valid frame is heard, an error has occurred. The station returns to the IDLE state and processes the received frame.

If a timeout occurs before a valid frame is heard, the station repeats sending the request_with_response data frame. If the station repeats sending that frame the number of times specified by the max_retry_limit parameter, then the request is abandoned. The IFM notifies the MAC-user entity that no response to the frame was received. The USE_TOKEN state is entered to check for another frame or token holding timer timeout.

5.2.7 Check Access Class. CHECK_ACCESS_CLASS controls the transmission of frames for different access classes. If the priority option is not implemented all frames are considered to be high priority and the CHECK_ACCESS_CLASS state serves to control entry to the PASS_TOKEN state.

If the priority option is implemented, a station may send frames for lower priority access classes before passing the token. Each access class other than the highest has a target token rotation time. At the time that the station has the token and begins to consider transmitting frames for that access class the residual time left in the target rotation timer is loaded into the token holding timer and the station returns to the USE_TOKEN state. At this time the target rotation timer is also reloaded to its initial value.

Thus, the station will alternate between USE_TOKEN and CHECK_ACCESS_CLASS states for each access class. If time is available, data frames will be sent in the USE_TOKEN state. When the lowest priority access class has been checked, the station will proceed to the pass token process, described next.

When a station has completed sending data frames, it shall enter the PASS_TOKEN state. Three conditions can occur:

(1) The station knows its successor, so it simply passes the token and enters the CHECK_TOKEN_PASS state.
(2) The station knows its successor but must first check if new stations desire to enter the logical ring. The station sends a solicit_successor frame and enters the AWAIT_RESPONSE state.
(3) The station does not know its successor. (This condition occurs after initialization and under error conditions.) The station sends a solicit_successor_2 frame opening response windows for all stations in the system and enters the AWAIT_RESPONSE state.

5.2.8 Pass Token. PASS_TOKEN is the state in which a station attempts to pass the token to its successor or solicit a new successor.

When its inter_solicit_count value is zero and time remains on the ring_maintenance_timer, the station allows new stations to enter the logical ring before passing the token. The token holding station does this by sending a solicit_successor_1 or solicit_successor_2 frame, as appropriate, and enters the AWAIT_RESPONSE state. (See 5.1.4 for the details of this operation.)

If the address of the successor, NS, is known, the station performs a simple token pass following any new successor solicitation. (See 5.1.3 for details of this operation.) If the successor responds and the station hears a valid frame, the station has completed its token passing obligations.

If NS is not known, the station sends a solicit_successor_2 frame to itself. Since this frame has two response windows and identical source and destination addresses, it forces all stations on the network that desire to be in the logical ring (whether or not they were previously) to respond. Those stations whose addresses are lower than that of the sender of the token frame transmit in the first window; those with addresses higher transmit in the second window.

During the response window(s), the token holding station listens for set_successor frames from potential successors. If no such frame is heard, the station stops trying to maintain the logical ring and listens for transmissions from any other station. (See 5.1.4 for details of this operation.)

5.2.9 Check Token Pass. CHECK_TOKEN_PASS is the state in which the station waits for a reaction from the station to which it just passed the token.

The station sending the token waits one slot_time for the station receiving the token to transmit. The one slot_time delay accounts for the time delay between sending the token frame and the arrival back at the sender of the corresponding response.

If a valid frame is heard that started during the response window, the station assumes the token pass is successful. The frame is processed as if it were received in the IDLE state.

If nothing is heard in one slot_time, the station sending the token assumes the token pass was unsuccessful and returns to the PASS_TOKEN state to either repeat the pass or try another strategy.

If noise or an invalid frame is heard, the station continues to listen for additional transmissions, as described in detail in 5.1.3.

5.2.10 Await Response. AWAIT_RESPONSE is the state in which the station attempts to sequence candidate successors through a distributed contention resolution algorithm until one of those successors' set_successor frames is correctly received or until no successors appear. The state is entered from the PASS_TOKEN state whenever the station determines it is time to open a response window or if the station does not know its successor (as in initialization or when a token pass fails).

The station waits in the AWAIT_RESPONSE state for a number of response window times. If nothing is heard for the entire duration of the response window(s), the station goes to the PASS_TOKEN state, either to pass the token to its known successor or to try a different token pass strategy.

If a set_successor frame is received, the station waits for the remainder of the response window time to expire. The station then enters the PASS_TOKEN state and sends the token to the new successor.

If the received frame is other than a set_successor frame, the station drops the token (since some other station is acting as if it also has a token, thus creating a multiple token situation) and reenters the IDLE state.

If noise is heard during the response windows, the station cycles through a procedure of sending resolve_contention frames that each open four response windows, and waits for a distinguishable response that began in a response window. The loop repeats a maximum of max_pass_count times, each time instructing contending stations to select a different two bits of their address to determine in which of the four opened windows to transmit.

5.3 Interface Machine (IFM) Description. The IFM acts as an intermediary between the other functional machines of the MAC sublayer and the MAC-user entities with which the MAC entity communicates. Its internal operation is largely unspecified, because the MAC sublayer will function correctly, no matter how the IFM functional requirements are met.

The IFM has eight primary functions, three of which are optional:

(1) Accept and generate the service primitives specified for the LLC-MAC interface.
(2) Queue pending service requests.
(3) Recognize addresses of data frames destined for this station.
(4) Map LLC quality of service requests from LLC terms (service_class) into MAC terms (access_class).
(5) Optionally, maintain multiple queues of send requests, separated by access_class (or by service_class).
(6) Maintain first-in, first-out ordering within each queue of pending service requests.
(7) Optionally, generate a response_received indication to the ACM when a data frame is received whose quality parameter specifies response.
(8) Optionally, accept a response from the responding MAC-user entity following receipt of a frame specifying request_with_response, and send a response frame.

Since the service primitives are treated in detail elsewhere in this standard, they are not discussed further here.

The IFM should also handle address recognition for data frames. While some address recognition might be done in the RxM, the need to check potentially large numbers of group addresses for relevance could place unreasonable demands on the RxM.

NOTE: The definition of the group address recognition mechanism is under study.

The notion of quality of service is similar to that of priority, though it is slightly less specific, and deliberately so. Higher protocol layers may rank their messages in some way that may be only approximately realizable by the lower layers. In the LAN case, the LLC sublayer may assign any one of eight service_class rankings to each frame it requests to be sent. The token bus MAC sublayer may not be able to provide different qualities of service, or MAC access_classes, so it shall be able to map all service_classes into one access_class. The MAC sublayer may, optionally, provide four distinct access_classes. When it does so, the IFM shall maintain multiple request queues, so that the requests may be handled according to both access_class (or service_class) and arrival order.

5.4 Receive Machine (RxM) Description. The RxM (see Figs 1-3 and 5-4) accepts MAC-symbols (see 6.1.2) from the Physical Layer and generates high-level data structures and signals for the MAC ACM and the MAC IFM. The interface between the Physical Layer and the MAC sublayer is the PHY-UNITDATA indication primitive specified in Section 8. This description of the RxM embodies the PHY-UNITDATA indication primitive as an encoded MAC-symbol and an associated clock. The other RxM interfaces are internal to MAC and are composed of the signals and data structures shown in Fig 5-4.

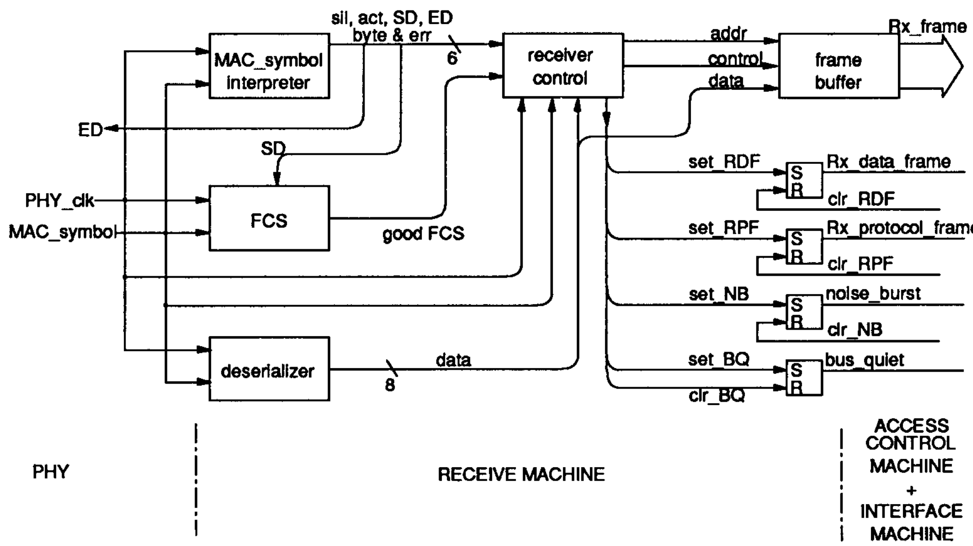

Fig 5-4
Receive Machine

bus_quiet. This signal is asserted when the medium is inactive. It is set and cleared by the RxM, and it is read by the ACM.

Rx_protocol_frame. A Boolean variable that is true when a valid MAC_control frame has been received. This variable is set by the RxM, and it is read and cleared by the ACM.

Rx_data_frame. A Boolean variable which is true when a valid data frame has been received. This variable is set by the RxM, read by both the ACM and the IFM, but cleared only by the IFM.

noise_burst. This signal indicates that there was activity on the medium that did not result in a valid frame (see 4.2.3 for the specification of invalid frames). This signal is set by the RxM, and it is read and cleared by the ACM.

frame_buffer. When a valid frame has been received, this data structure contains the frame control octet, destination_address, source_address, data_unit, and FCS fields of the received frame. The contents of this structure are written by the RxM. When the Rx_protocol_frame signal is set, the contents of this structure are used by the ACM. When the Rx_data_frame signal is set, the contents of this structure are used by the IFM and ACM.

There are four major functional blocks in the RxM: the MAC-symbol interpreter, the FCS block, the deserializer, and the receiver control block. (See Fig 5-4.) These blocks are described as synchronous machines clocked by PHY_clk. Where practical, these blocks are decomposed into finite state machines. See Figs 5-5, 5-6, and 5-7. In these figures, an unbracketed label on an arc is the condition for the arc; a label in brackets is an output of that arc.

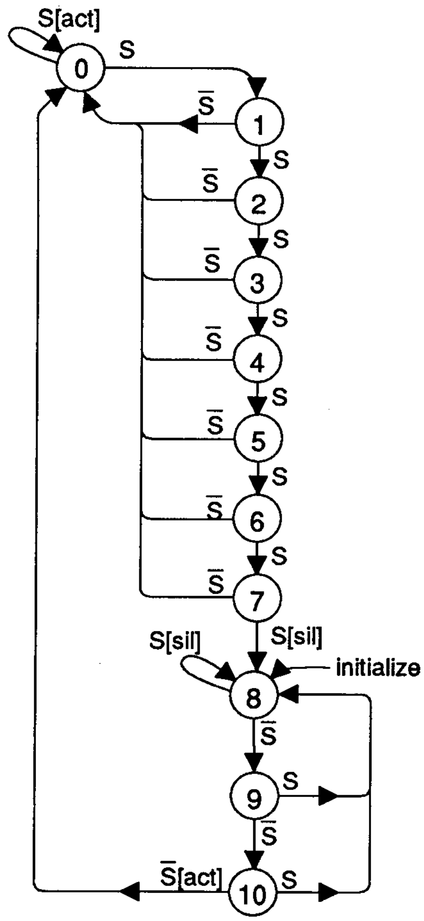

Fig 5-5
sil/act Detector Finite State Machine

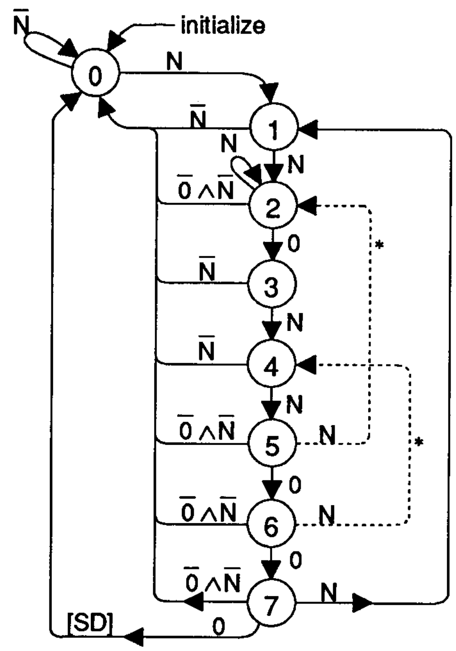

*** The destination for these arcs may be state 1 or 0 rather than the ones shown here.**

Fig 5-6
SD Detector Finite State Machine

The MAC-symbol interpreter accepts MAC-symbols and generates the following signals:

sil, act. These signals represent silence (*sil*) and activity (*act*) on the bus, respectively. The signals *sil* and *act* control set_BQ and clr_BQ (set and clear bus_quiet, respectively). The only reason they are not exactly equivalent is that the receiver control block shall maintain certain timing relationships among the bus_quiet and noise_burst variables. The state machine shown generates *sil* only after eight MAC-symbols of *silence* (S), and requires three non-*silence* MAC-symbols to generate *act*. In part, this is merely to provide a degree of hysteresis in the receiver and the exact counts may be varied by the implementer; however,

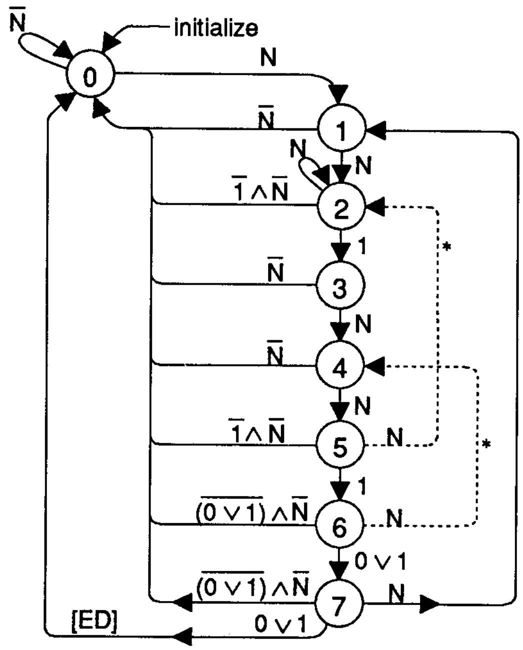

* The destination for these arcs may be state 1 or 0 rather than the ones shown here.

Fig 5-7
ED Detector Finite State Machine

when a valid frame is received, bus_quiet shall be set simultaneously or after the appropriate Rx_frame signal is set. Failure to do so can cause the ACM to malfunction.

SD, ED. Any valid frame delimiter in the received MAC-symbol sequence is detected, regardless of context. The ED signal is returned to PHY for use by conditioning and synchronizing circuits.

byte, err. (State machine not shown.) The byte/err detector is unconditionally initialized by SD, establishing byte alignment. While only 0 and 1 MAC-symbols are subsequently received, byte strobes true every eight PHY_clk periods. The reception of any MAC-symbol other than 0 or 1 starts an exit sequence. (See 6.1.2 for definitions of MAC-symbols.) Possible exit sequences are as follows:

(1) Reception of a P, B, or S MAC-symbol: report the error and terminate (wait for initialization).
(2) Reception of an N not on an octet boundary: report the error and terminate.
(3) Reception of an N on an octet boundary that does not prove to be a frame delimiter: report the error and terminate.

(4) Reception of an N on an octet boundary that introduces SD: report an error if currently receiving a frame and initialize.

(5) Reception of an N on an octet boundary that introduces ED: terminate.

The FCS block implements the FCS calculation specified in 4.1.6. It is initialized by any SD, and shifts in data until a nondata symbol is received.

The output of the deserializer is defined only when SD, ED, or a data byte is reported.

NOTE: Reporting of the value of the data bits in delimiters to higher layers is not required by this specification.

The receiver control block has three components. The noise_burst/bus_quiet (NB/BQ) detector sets and clears bus_quiet, and it sets noise_burst on the rising edge of bus_quiet if a valid frame has not been received (see Fig 5-8). The FC block transfers the deserializer output to the frame buffer, checks the FC octet for data frames, and determines whether the frame contained a legal number of octets. The valid frame detector further verifies that the FCS was correct and that the frame was composed of valid octets.

NOTE: The presentation of this section is for use by the ACM. The statistics recorded for the management entity in the RxM counters group (see 3.2.10) may use additional information from the received data to better report what is heard.

Fig 5-8
NB/BQ Finite State Machine

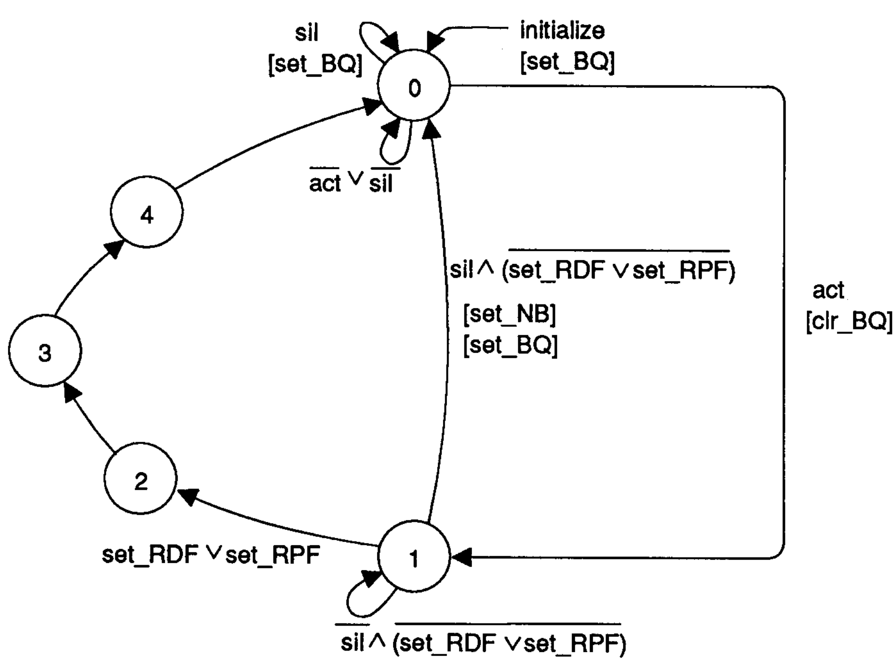

5.5 Transmit Machine (TxM) Description. The TxM has a fairly straightforward and simple operation. The ACM forwards frames for transmission to the TxM as a unit (at least for the purposes of this description). The TxM then passes the frame, along with appropriate delimitation, to the PLE, one MAC-symbol at a time, for transmission on the medium. The TxM is responsible for sending the proper amount of preamble, computing the FCS and including it within the transmitted frame, and delimiting the frame with SD and ED.

5.6 Regenerative Repeater Machine (RRM) Description. Bus repeaters are used to connect electrically, optically, or frequency separate bus segments together into an extended logical bus network. A repeater has a separate receiver for each segment to which it is connected. Each receiver separates clock and MAC-symbols, and has a receive switch that controls which receiver is the MAC-symbol source for the stuffer. A receiver is selected as the stuffer source by reporting non-*silence* when there is no other receiver acting as stuffer source. A stuffer source is not deselected until it reports *silence*.

Figure 5-9 shows a generalized repeater. Included in the figure are *n* distinct physical segments of a single bus. The simplest form of a repeater (a broadband remodulator) has only one transmitter/receiver pair. The stuffer shall

(1) Re-insert preamble symbols that are lost in the receiver [except as noted in 8.2.1.5(3)],

(2) Transmit the abort sequence when *bad_signal* is reported to it. The abort sequence always begins on an octet boundary relative to that part of the repeated data already transmitted, and

(3) Set the E bit in the ED if the repeated frame has an FCS error.

The output of the stuffer goes to one input of the originate/repeat switch; the other input to this switch comes from the TxM. When the TxM reports non-*silence*, the originate/repeat switch unconditionally selects that input. When the TxM reports *silence*, the switch selects the stuffer, after a short delay, to guarantee that symbols that originated in the local TxM will not be repeated.

While the originate/repeat switch is enabling the local TxM, all transmit switches are on, so the symbols are transmitted onto all segments. While the originate/ repeat switch is selecting the stuffer, all transmit switches are on, except the one corresponding to the segment whose receiver is the current symbol source for the stuffer.

A repeater may or may not be a station within the network. If the repeater also functions as a station, the repeater's RxM and TxM perform all the functions described in this standard. If the repeater does not function as a station, the RxM need only check the FCS of received frames, so as to compute the E-bit value for the retransmitted ED. The TxM of a nonstation repeater continuously sends *silence*.

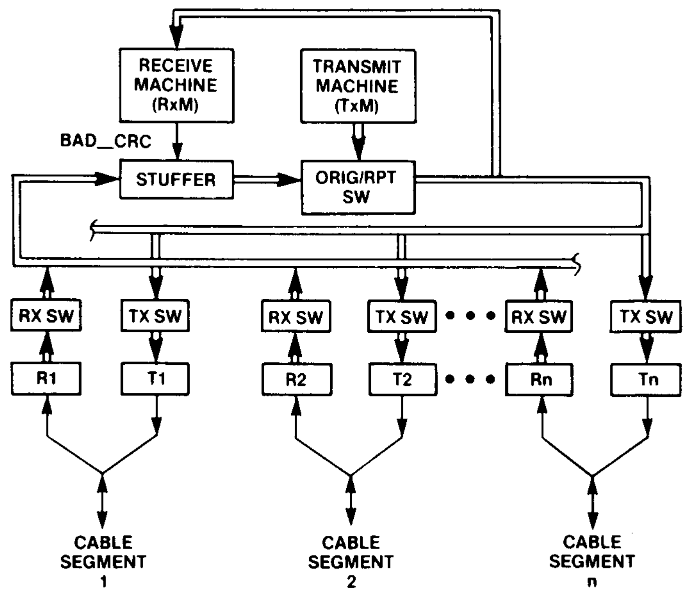

Fig 5-9
Bus Repeater

6. MAC Sublayer Definitions and Requirements

All of the aspects of the MAC sublayer operation and mechanism that are required for conformance to this standard and that are not specified in Sections 4 or 7 are specified in this section.

6.1 MAC Definitions. Critical MAC parameters that are constrained by this specification are defined as follows:

6.1.1 Immediate_response. The immediate transmission of a MAC-level response to a received frame. This assumes that no other transmissions or actions were intervening.

6.1.2 MAC-Symbols. The smallest unit of information exchanged between MAC sublayer entities. The six MAC-symbols are as follows:

Name	Abbreviation
zero	0
one	1
non_data	N
pad_idle	P
silence	S
bad_signal	B

Where binary 0 and 1 data bits are discussed, they are sent and received as zero and *one* MAC-symbols, respectively.

6.1.3 MAC-Symbol_time. The time required to send a single MAC-symbol. This is the inverse of the LAN data rate.

Nominal Data Rate	Nominal MAC-Symbol_time
1 Mb/s	1.0 µs
5 Mb/s	0.2 µs
10 Mb/s	0.1 µs
20 Mb/s	0.05 µs

6.1.4 Octet_time. Corresponds to the time interval required to transmit eight MAC-symbols.

6.1.5 PHY-Symbols (Physical Layer symbols). Correspond to the waveforms impressed on the physical medium. The PHY-symbol definitions are given in the following subsections:

(1) 12.7.2 for single-channel phase-coherent FSK
(2) 14.8.2 for multi-level duobinary AM/PSK (1 MAC-symbol/PHY-symbol)
(3) 14.11.1 for multi-level duobinary AM/PSK enhanced
(4) 16.7.2 for fiber optic media
(5) 18.7.2 for single-channel phase-continuous FSK

6.1.6 Transmission_path_delay. The maximum delay that transmissions experience going through the physical medium from a transmitter to a receiver. The following formula is a general definition of transmission_path_delay:

transmission_path_delay = (physical_medium_delay
+ amplifier_delay
+ repeater_delay)

The definition of transmission_path_delay for each signaling technique is given in the following subsections:
(1) 13.7 for single-channel phase-coherent FSK
(2) 15.7 for multi-level duobinary AM/PSK
(3) 17.7 for fiber optic media
(4) 19.7 for single-channel phase-continuous FSK

6.1.7 Station_delay. The time from the receipt of the PHY-symbols corresponding to the last MAC-symbol of the received ED at the receiving station's physical medium interface to the impression of the first immediate_response PHY-symbols onto the physical medium by that station's transmitter.

6.1.8 Safety_margin. A time interval no less than one MAC-symbol_time.

safety_margin ≥ MAC-symbol_time

6.1.9 Slot_time. The maximum time any station need wait for an immediate_response from another station. In the following formula the units of transmission_path_delay, station_delay, and safety_margin are in MAC-symbol_times. Slot_time is measured in octet_times, and is defined as

Slot _time = INTEGER ({[(2 * transmission_path_delay + station_delay}
+ safety_margin) / MAC-symbol_time] + 7} / 8)

6.1.10 Response_window. The basic time interval which the MAC protocol allows, following certain MAC_control frames, for an immediate_response from another station. This interval is one slot_time long.

response_window duration = slot_time

If a station, waiting for a response, hears a transmission start during a response_window, that station shall not transmit until at least after the received transmission has terminated.

6.2 Transmission Order. The frame formats used by the MAC sublayer and the detailed contents of the octets of those frames are specified in Section 4. The octets of a frame and the MAC-symbols of an octet shall be transmitted from the MAC sublayer to the Physical Layer, and vice versa, in the order specified in Fig 6-1; the first octet of the frame shall be transmitted first, and the first MAC-symbol of each octet shall be transmitted first. The first octet and first MAC-symbol correspond to the top octet and the left-most MAC-symbol shown in Fig 6-1. Those notations within octets of Fig 6-1 that are not MAC-symbols are casual descriptions that are specified in Section 4. Figure 6-1 describes a complete general frame or MAC protocol_data_unit (M_PDU).

Fig 6-1
MAC Protocol_data_unit Transmission Order

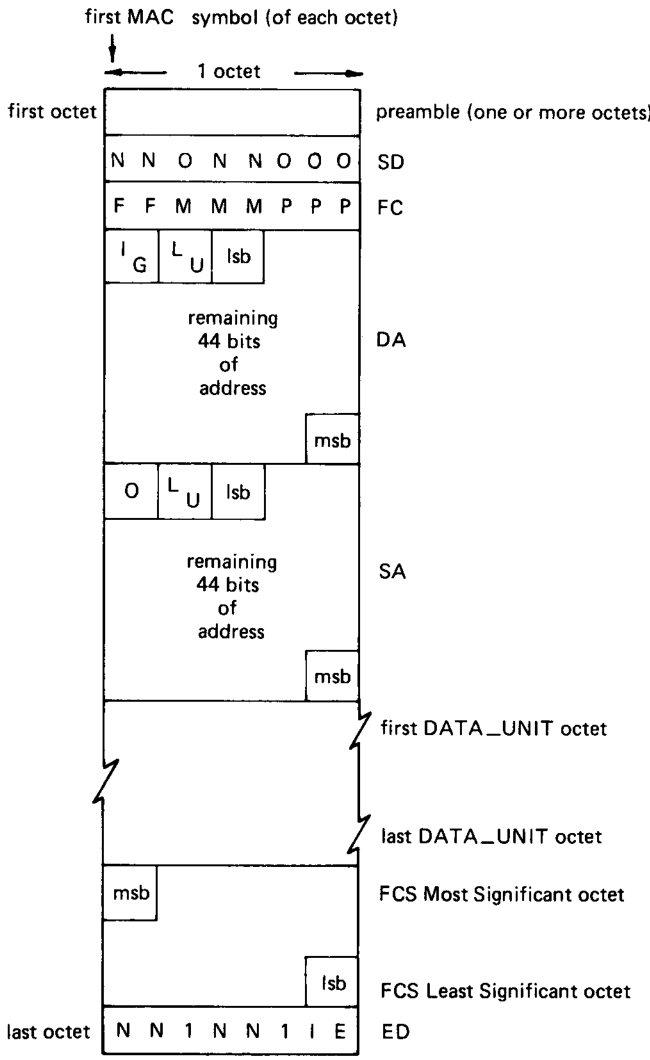

6.3 Delay Labeling. Vendors shall provide a worst-case value for equipment delay times. Vendors may optionally specify a minimum network slot_time when their equipment anticipates some minimum delay, in order to function correctly. When uncertain of the exact value of the equipment delay, vendors shall state an upper bound. It is recommended that vendors of equipment conforming to this standard label the equipment with that equipment's contribution to the station_ delay. A vendor of a complete station labels the station with the total station_ delay. A vendor of a component, intended to be assembled by an end-user into a station, labels the component or otherwise documents the delays that contribute to station_delay.

6.4 Miscellaneous Requirements

6.4.1 Station Initialization. On power-up, the station shall enter the OFFLINE state. While in the OFFLINE state, the station shall not impress any signaling on the LAN medium. The station shall progress from the OFFLINE state to the IDLE state only when it has been configured with the basic station operating parameters needed for correct operation of the MAC protocols. These operating parameters include at least

(1) init_TS (station address)
(2) address_length (implicit in TS)
(3) init_slot_time
(4) init_hi_pri_token_hold_time
(5) init_target_rotation_time(4) (if priority is implemented)
(6) init_target_rotation_time(2) (if priority is implemented)
(7) init_target_rotation_time(0) (if priority is implemented)
(8) init_target_rotation_time(ring_maintenance)
(9) init_ring_maintenance_timer_initial_value
(10) init_max_inter_solicit_count
(11) min_post_silence_preamble_length
(12) init_in_ring_desired
(13) init_max_retry_limit (if request_with_response is implemented)

6.4.2 Token Passing Order. The token shall be passed from station to station in station-address order, descending numerically, except that the station with the lowest address shall pass the token to the station with the highest address in order to close the logical ring. Figure 6-2 illustrates the address-ordered logical ring and shows the logical relationships that hold between addresses of adjacent stations in a logical ring with three or more members.

6.4.3 Station Receipt of Its Own Transmissions. In systems with significant transmission_path_delay, such as broadband systems, a transmitting station may receive its own transmissions after some small but significant delay. The MAC access mechanism of such a transmitting station shall not be misled by the receipt of its own transmissions. In Section 7 the state diagrams specify where a station should ignore its own transmissions.

6.4.4 Token Holding Time. The token holding station shall only begin transmitting a frame when there is time remaining on the token_hold_timer. A trans-

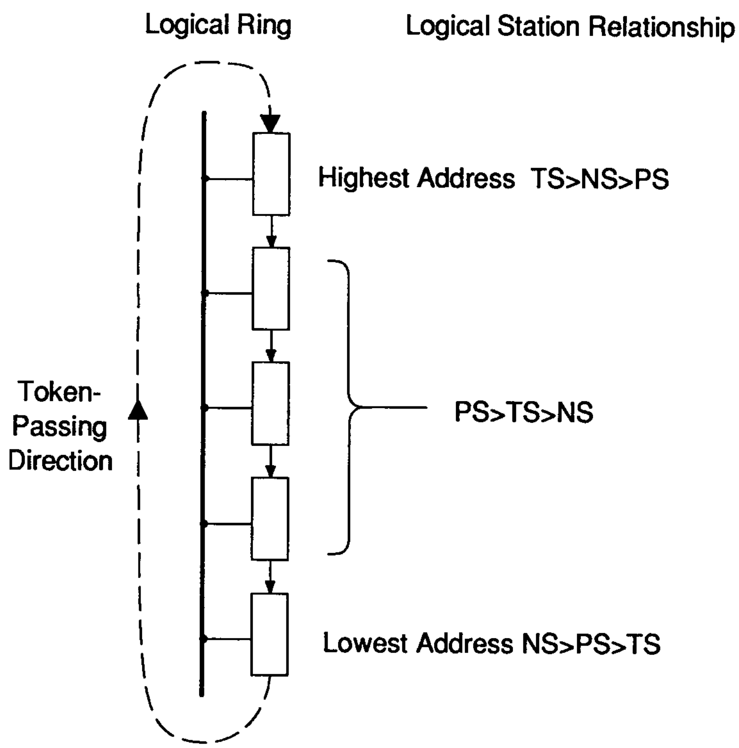

Fig 6-2
Logical Token-Passing Ring

mission may run past the expiration of the token hold timer by up to the time to transmit a maximum length data frame.

Use of the MAC-acknowledged connectionless class of service (request_with_response mechanism) may further exceed the token-holding time. The excess time may be caused by repeated retries with associated waiting periods.

6.4.5 Address Lengths. Addresses shall be either two octets (16 bits) or six octets (48 bits) long, including the individual/group (I/G) bit and local/universal (L/U) administration bit where appropriate. All stations operating on a single LAN shall be assigned and use addresses of the same length.

6.4.6 Randomized Variables. The station shall provide a two-bit (that is, four-valued) random variable for use in the MAC protocols. For the medium access protocol to benefit from the randomization, the technique used to create the random values shall be statistically independent between stations. Thus, random number generators tied in any way to the received data clock, for example, would not produce statistically independent variable values.

The variables shall be rerandomized periodically. “Periodically” shall be interpreted to mean either an interval not to exceed 50 ms, or every use of the random variable.

6.4.7 Contention Delay. If the station hears a solicit_successor or who_follows frame it determines in which response window to contend based on the station's address and the SA and DA addresses in the frame. If the station wants to contend in the first window it loads the contention timer with zero, so that the station proceeds immediately to the DEMAND_IN state. If the station wants to contend in the second window, the contention timer is loaded with one, so that the station listens during the first window.

Following receiving a resolve_contention frame, if the station is contending, it loads the contention timer with the one's complement of two bits selected from its own address as indexed by the resolution pass count. The station thus listens zero, one, two, or three slot_times before again contending.

6.4.8 Token Claiming. If the bus_idle_timer expires, a station may transmit a claim_token frame and set the claim_timer. When the claim_timer expires, if no transmissions are present at that instant, the station sends an additional claim_token frame and repeats the delay and transmission check. This procedure is repeated until either transmissions from another station are heard or the value of the claim_pass_count equals max_claim_pass_count.

The data_unit length of the claim_token frames are 0, 2, 4, or 6 slot_times as a function of two bits of the station's address. Indexing through the address permits the claim process to designate a single station as the token holder.

In the simplest case, the contending station with the highest address will always win the address sort. In practice, a different contending station can win the token. This can occur if differing delays among the contending stations cause them to drift “out of step” with each other during the contention process. Because each station's response delay shall not exceed its station_delay, the algorithm is still assured to designate a unique token holder.

6.5 Use of Address Bits in Contention Algorithms. The contention processes used to claim a new token or demand logical ring entry both use the bits of the station address to accomplish a sorting-like resolution in which the station having the numerically largest or highest address value wins. The following paragraphs treat the address as an array of binary (0/1) values or address_bits; for notational purposes, “address_bit(i)” indicates the ith binary bit of the station's address, with “address_bit(1)” the most significant bit. These address_bits are used two-at-a-time, starting with the most significant address_bits.

6.5.1 Claim_token Frame Length. A station that is attempting to claim a new token first determines that no other station is transmitting, then transmits a claim_token frame containing a data_unit with a length equal to 0, 2, 4, or 6 slot_times. It then waits or delays one slot_time before again listening for other transmissions. The token claiming contention process shall consist of N cycles of listening, transmitting, and delaying, where N is a function of the station's address length in bits:

N = (address_length / 2) + 1

For a two-octet address, N = 9; for a six-octet address, N = 25.

The length, L, of the nth claim_token frame's data_unit in octet_times (for the nth cycle of the token claiming process), shall be determined as follows:

for $1 \le n < N$
  bit_value := 2*address_bit((2*n)–1) + address_bit(2*n)
  L := 2*slot_time*bit_value

for $n = N$
  L := 2*slot_time*random_4

where

bit_value = 0, 1, 2, or 3 as a function of the two address bits used in cycle n, and
random_4 is a random number equal to 0, 1, 2, or 3

6.5.2 Demand Delay Time Interval. A station which is demanding entry into the logical ring first listens for other transmissions, delaying its next transmission for 0, 1, 2, or 3 slot_times. Then, if it has heard no other transmissions, it transmits a fixed length set_successor frame. This delay, preceding set_successor frame transmissions, is called the demand_delay. The contention process for engaging in logical ring entry shall consist of, at most, N cycles of listening demand_delays and transmission, where N is a function of the address length in bits:

N = (address_length / 2) + 1

The number of slot_times, D, to delay before the nth transmission (for the nth cycle of the contention resolution process) shall be determined as follows:

for $1 \le n < N$
  D : = 3 – bit_value

for $n = N$
  D : = random_4

where bit_value and random_4 are as specified in 6.5.1

6.6 Options within MAC Sublayer

6.6.1 Priority Mechanism. The implementation of multiple classes of service, or priorities, is an option. A station that does not implement the priority mechanism shall transmit all data frames with an access_class value of 6, corresponding to the highest priority.

Where the priority mechanism is implemented, it shall meet the requirements set forth in 6.6.1.1 through 6.6.1.3.

6.6.1.1 Access_classes. The priority mechanism shall provide four levels of service with respect to a frame's priority of access to the medium; these levels are called access_classes. The access_classes shall be identified as 0, 2, 4, and 6, and access_class 6 shall be the highest priority or most favored level of service.

6.6.1.2 Quality_of_service to Access_class Mapping. Where priority is implemented, the priority component of the MAC-user's quality_of_service parameter, contained in the MA-UNITDATA request, is eventually satisfied through use of the access_classes. The priority component of the quality_of_service request shall first be satisfied by assignment of the request to one of eight MAC service_classes. These MAC service_classes shall then be mapped into MAC access_classes as described in the following table:

service_class	access_class	priority
0, 1	0	lowest
2, 3	2	
4, 5	4	
6, 7	6	highest

6.6.1.3 Token Rotation Timers. A station that implements the priority mechanism shall provide three (actual or virtual) token_rotation_timers, one for each access_class. These timers shall all run concurrently, counting downward from an initial value to zero, at which point they shall stop counting and their status shall be "expired." These timers shall count in units of octet_times, and shall otherwise be managed as specified in Section 7.

6.6.2 Request_with_response Mechanism. The ability for the MAC sublayer to send request_with_response frames and to reply to such frames with a response frame is an option. A station that does not implement the request_with_response mechanism shall not respond to request_with_response frames.

NOTE: In previous versions of this standard, the request_with_response option was named "immediate response" option. The name was changed to reduce confusion; the function is the same.

6.7 Additional MAC Capabilities. The functions described in this subsection are recommended capabilities of the MAC sublayer that are not formally defined in the ACM state machine in Section 7. These capabilities, while not required, can enhance the functionality or robustness, or both, of the implementation. Since the detailed operation of these capabilities does not affect station interoperability, the description in this standard is intentionally vague. The functions are described in general terms to allow implementors to provide an optimal implementation.

6.7.1 Duplicate Address Detection. The access protocol described in this standard depends on all stations having unique station addresses. If two stations have an identical address, network operation is unpredictable.

The cause of two stations having the same address is outside the scope of this standard. The cause of most duplicate address errors is the failure of the network administrator to correctly configure stations. To decrease the probability of creating a duplicate address, vendors should take steps to either

(1) Require that the station address be entered into a station before it can enter the logical ring, or

(2) Preconfigure stations with universally administered addresses. (See 4.1.4.1 for information on universal address administration.)

The ACM in Section 7 specifies certain functions that perform duplicate address detection. These functions are a minimum set that shall be implemented by all conformant stations. However, analysis indicates that stations with duplicate addresses can exist undetected on certain networks which use only the required duplicate address detection algorithms.

Some applications may require additional duplicate address detection mechanisms beyond the minimum required by the standard. For example, duplicate addresses might be detected by mechanisms that

(a) Recognize MAC control frames with the station's source address that the station did not send, or

(b) Recognize MAC control frames with the station's source address received later than one slot time after the station finished transmitting.

As long as such mechanisms do not interfere with interoperability, they are beyond the scope of, but not precluded by, this standard.

6.7.2 "Aliased" Data Frames. Certain applications, such as MAC-level bridges, can require the MAC to send an aliased data frame, i.e., a data frame whose SA field does not contain the address of the station transmitting it. The following changes to the ACM ensure proper functioning of the ACM even when aliased request_with_no response frames are sent. (Aliasing is not permitted of request_with_response and response frames.)

6.7.2.1 Echo Detection Modifications. The ACM uses the test "Rx_frame.SA/ = TS" to ensure that a received data frame is not an echo of its own transmission. This test occurs in the exit conditions of the arcs:

(1) unexpected_frame_6,

(2) pass_ok, and

(3) unexpected_ frame_10.

A station transmitting aliased data frames must employ some other mechanism to ensure that its ACM will never traverse one of these arcs on hearing its own echo. What changes, if any, are required will depend on the implementation. A method that works for all implementations is for the station to pause, as necessary, before sending any request_with_response, solicit_successor, or token frames, so as to ensure that at least one slot_time has expired since the transmission of any aliased frame.

6.7.3 MAC "Heartbeat" Indicator. It may be important for higher layer functions to have confidence that the MAC sublayer is operating properly and connected to a functioning network. The MAC "heartbeat" indication provides a periodic indication internal to the station that the MAC sublayer is functioning properly and that other stations' MACs are working.

The heartbeat indication is signaled when the ACM traverses the no_response_10 arc with the new value of the pass_state variable equal to pass_token. The heartbeat indication occurs when the station opens a response window and hears nothing in the window. The heartbeat signals that the station is receiving the token from another station and the value of its successor is unchanging, which are the normal steady-state network conditions.

The heartbeat indication should occur periodically in a station where in_ring_desired remains true and sole_active_station is false; in other words, in a station that is a member of an active logical ring. If the heartbeat indication does not occur periodically, a monitor within the station may assume that either the station's MAC is malfunctioning or some other catastrophic network failure has occurred. After possibly a few retries, the station should be removed from the network to protect other stations from the effects of the possible MAC malfunction.

6.7.4 Ring Exit. If in_ring_desired is set to false and stations are entering and exiting the network, it is remotely possible for a station to re-enter the ring with NS_known true but NS not valid. This condition will be resolved through the normal ring recovery mechanism.

This situation occurs when the station takes the arc leave_ring, then in_ring_desired or any_send_pending become true before the station can take the arc exit_ring, and thus the station will re-enter the ring through the arc repair_demand_in and not update NS. This sequence can lead to token pass failures, who_follows queries, and in an extreme case, a solicit_any.

If desired, this situation can be avoided by preventing the station that is in the process of leaving the ring from reentering through the arc repair_demand_in.

6.8 Delegation of Right to Transmit. A station holding a token may request a second station to transmit a response without the second station holding the token. The first station, in effect, delegates the authority to transmit to another station. The secondary station shall conform to all of the requirements imposed by this standard on the token holder except for participating in the logical token-passing ring and associated protocol mechanisms (unless the secondary is or needs to be in ring). The secondary station need not be in the logical ring from an address sense. The secondary shall not transmit on the network unless

(1) Delegated as a transmitter by a token holder, or
(2) Transmitting is authorized by the procedures specified in Section 7.

7. Access Control Machine (ACM) Description

The token-passing bus MAC mechanism is defined in this section. The section begins with a description of the variables and functions used in the definition of the mechanism. The second part of this section is a formal state machine description of the access control mechanism using the variables and functions discussed in the first part.

7.1 Variables and Functions. The variables and functions of the state machine description are grouped into categories as follows:

(1) Variables defined by the management entity
(2) Variables defined by the IFM
(3) Timers
(4) Variables defined by the RxM
(5) Other ACM variables and functions

7.1.1 Management Variables. The management entity provides the MAC sublayer with the following parameters:

TS. This station's address. A bit-string variable set to the value of the station's 16-bit or 48-bit address. The value of TS implicitly defines the network address length.

slot_time. An integer in the range 1 to 2^{13}–1 octet_times. See 5.1.1 and 6.1.9.

min_post_silence_preamble_length. An integer in the range 1 to 15 equal to the minimum number of octets of preamble to be used at the beginning of a transmission after the station has been silent. The value of min_post_silence_preamble_length is determined by the type of Physical Layer entity (PLE) used in the station. See 9.3 for a description of how the value of this parameter can be obtained from PLE management at station initialization.

max_pass_count. An integer equal to half the station's address length in bits, plus one (thus, equal to 9 for a 16-bit address length and 25 for a 48-bit address length).

The value of max_pass_count limits loops in the ACM. The value is used to limit the token contention process. After cycling through max_pass_count contention cycles the process shall be stopped if, due to an error, a single contender cannot be resolved.

The value of max_pass_count is also used to stop the token claiming process. After sending max_pass_count_claim token frames, if no other station is heard, a station can claim the token.

max_inter_solicit_count. An integer number of token possessions within the range 2^4 to 2^8–1. This value, in addition to the ring_maintenance_timer, determines how often a station opens response windows. Normally, a station opens response windows prior to every *n*th pass of the token, where *n* is the value of max_inter_solicit_count.

If all stations in the ring use the same max_inter_solicit_count, they could all consistently open response windows on the same token rotation. This action would lead to rapid token rotations where no response windows are opened, and occasional rotations where every station opens a response window before passing the token.

To avoid all stations in the ring having the same value of max_inter_solicit_count, the least significant two bits of the value shall be chosen randomly. The actual value used for the max_inter_solicit_count shall be changed by each station by rerandomizing the least significant two bits of the variable at least every 50 ms or on every use.

max_retry_limit. An integer in the range 0 to 7 that determines how many times a station resends a request_with_response frame for which it does not get a response.

target_rotation_time(access_class). An array of integers in the range 0 to 2^{21}–1 octet_times, used with the priority option and with the ring_maintenance timer. (See 7.1.4.2 for a discussion of the variable's function.)

ring_maintenance_timer_initial_value. An integer in the range 0 to 2^{21}–1 octet_times, used to determine the initial value of the token_rotation_timer(ring_maintenance) upon entry to the ring. A large value will cause the station to solicit successors immediately upon entry to the ring; a value of zero will cause the station to defer this solicitation for at least one rotation of the token. An implementation may treat any non-zero value as infinite.

hi_pri_token_hold_time. An integer in the range 0 to 2^{16}–1 octet_times. Used to control the maximum time a station can transmit frames at access_class 6. If the priority option is not implemented, then hi_pri_token_hold_time determines how long a station can send frames of any access class.

in_ring_desired. A Boolean variable that determines the ACM's steady-state condition when it has no queued transmission requests. If the variable is true, the station should be in_ring (a participant in the token-passing logical ring). If false, the station should be out_of_ring (an observer of the token-passing logical ring).

7.1.2 Interface Machine Variables and Functions

get_pending_frame(access_class). A function provided by the IFM of MAC. This function removes the first frame from the pending frame queue for the indicated access class, and returns it to the ACM for transmission.

get_pending_response ( ). A function provided by the IFM of MAC. This function accepts the response frame created by the MAC user and returns it to the

ACM for transmission. There can be a maximum of one response frame at any one time.

send_pending (access_class). An array of Boolean functions reflecting the state of the pending frame queues. The value of a function will be true if the pending frame queue for the indicated access_class is non-empty, and false otherwise.

any_send_pending. A Boolean function reflecting the logical OR of all the Boolean functions send_pending (access_class). any_send_pending is true if at least one of the pending frame queues is nonempty. If all queues are empty, the function's value is false.

response_received. A Boolean function that returns '"true" when the IFM has received a valid response frame after the station has sent a request_with_ response frame.

power_ok. A Boolean variable indicating that the ACM may begin operation. The value is provided by the management entity.

7.1.3 Note on Logical Ring Membership Control. The Boolean variable in_ring desired and function any_send_pending determine the operation of the ACM with respect to contending for the token and being in the logical ring as follows:

Variables and Function Values		ACM Action
in_ring_desired	any_send_pending	
false	false	Do not contend for token. Drop out if currently in token-passing logical ring.
false	true	Contend for token. Send data, which may empty the pending frame queues and make any_send_pending false. Exit logical ring if any_send_pending becomes false.
true	false	Contend for token if not sole active station. Remain in token-passing logical ring even without data to send.
true	true	Contend for token. Remain in token-passing logical ring and send data.

7.1.4 Timers. A number of timers are used in the description of the state machine. A *timer* is expressed as a set of procedures and a Boolean variable. The procedures are named *xx_timer.start(value)*, where *xx_timer* is the timer name and *value* is an integer that sets the timer delay. *xx_timer.value* returns the current value of the counter. The Boolean variables are named *xx_timer.expired* and have a value of *false* while the timer is counting down and *true* when the timer has expired and reaches a value of *zero*.

For example, the bus_idle_timer is set to a value of one (slot_time) by executing *bus_idle_timer.start*(1). The variable *bus_idle timer.expired* is then false for one slot_time.

7.1.4.1 Slot_time Interval Timers. The five timers (bus_idle_timer, claim_timer, response_window_timer, contention_timer, and token_pass_timer) work in integral multiples of the network slot_time. (The five timers are not used concurrently; thus, they could be implemented in a single hardware timer.)

bus_idle_timer. Controls how long a station listens in the IDLE state for any data on the medium before entering the CLAIM_TOKEN state and re-initializing the network. Most stations wait seven slot_times. The one station in the network having lowest_address true waits six slot_times. The function max_bus_idle returns the value 6 or 7 depending on the state of lowest_address.

claim_timer. Controls how long a station listens between sending claim_token frames. The claim_timer is always loaded with the value 1.

response_window_timer. Controls how long a station that has opened a response window listens before transmitting its next frame.

When sending a solicit_successor, who_follows, or resolve_contention frame this timer controls the length of time a station solicits responses. After sending a solicit_successor frame, the sending station loads the response_window timer with the number of windows opened. The timer thus determines how long the station remains in the AWAIT_RESPONSE state listening for stations to respond. If the timer expires and nothing is heard, the station goes to the PASS_TOKEN state and sequences to the next pass_token substate.

When sending a request_with_response data frame, the response_window_timer controls how long a station waits for a response frame before repeating the request_with_response frame.

contention_timer. Controls how long a station listens in the DEMAND_IN state after hearing a resolve_contention, solicit_successor, or who_follows frame when the station wants to contend for the token. If the station hears a transmission while listening, it has lost the contention and shall return to the IDLE state.

token_pass_timer. Controls how long a station listens after passing the token to its successor.

If any frame is heard before the token_pass timer expires, the station assumes that its successor has accepted the token. If the timer expires and a frame is not heard, the station assumes its successor did not accept the token and sequences to the next stage of the pass token procedure.

7.1.4.2 Octet_time Interval Timers. The remaining timers have a granularity of one octet_time, rather than one network slot_time. They are used to implement the access class structure and limit the time during which a station may start to transmit frames for each access class.

token_rotation_timer(access_class). This is a set of four timers, one for each of the lower three access classes and one for ring maintenance. The first three of these timers only exist for stations implementing the priority option. The fourth timer, ring_maintenance, always exists.

When the station begins processing the token at a given access class the associated timer is reloaded with the value of the target_rotation_time for that access class. When the station again receives the token, it may send data of that access class until the residual time in the associated token_rotation_timer has expired.

Upon initial entry to the ring, the first three priority timers are set to a value of zero (expired), and the token_rotation_timer(ring_maintenance) is set to the value ring_maintenance_timer_initial_value.

token_hold_timer. The residual time from the current token_rotation_timer is loaded into the token_hold_timer just before the token_rotation_timer is reloaded. The station may send data frames of the corresponding access class as long as the token_hold_timer has not expired.

When the station is sending highest access_class messages, the value of hi_pri_token_hold_time is loaded into the token_hold_timer. Thus, highest access class messages are limited to only a fixed number of octets, plus the completion of the last message, regardless of current network loading.

7.1.5 Receive Machine Variables and Functions. The outputs of the RxM are described in the following paragraphs:

bus_quiet. A Boolean variable, which is true whenever the PLE is reporting that *silence* is being received and false when something other than *silence* is being received. The variable bus_quiet is set and reset by the RxM and is only read by the ACM.

Rx_frame. A record written by the RxM. The record is updated to reflect the contents of the most recently received valid frame. The major fields in the record are

FC. The one-octet frame control field
DA. The two- or six-octet destination address field
SA. The two- or six-octet source address field
data_unit. The multi-octet data_unit field
FCS. The four-octet frame check sequence field

Rx_protocol_frame. This signal indicates that a valid frame has been received, and that the frame type is one of the MAC control frame types. This signal is set by the RxM, and it is read and cleared by the ACM.

Rx_data_frame. This signal indicates that a valid frame has been received, and that the frame type is a data frame. This signal is set by the RxM, and read by both the ACM, and the ACM/IFM; it is cleared only by the IFM.

noise_burst. A Boolean variable set by the RxM when bus_quiet goes true (the bus goes from non-*silence* to *silence*) and neither Rx_protocol_frame nor Rx_ data_frame were set during the transmission (that is, no valid frame was heard). It is reset by the ACM when the noise burst condition has been processed.

7.1.6 Other Variables and Functions. The following variables and functions are local to the MAC ACM.

TH—token holder's address. The address of the current token holder. A temporary buffer loaded from the SA field of a solicit_successor or who_follows frame. If a set_successor frame is sent by the station as part of the contention process, the DA address is taken from TH.

NS—next station's address. The address of a station's successor in the logical ring. NS is set when a station that does not know its successor hears a solicit_successor frame and contends for the token. The station sets NS to the value of the destination (DA) address field of the frame. (If the station successfully contends in a response window, it will receive the token and eventually pass it to the station whose address was loaded into NS.) For example, suppose a station with address 25 is not in the logical ring, and wants to enter. If this station hears a solicit_successor frame sent by station 30 with a DA address of 20, it will set NS to 20, the DA address in the frame. If the station contends in the window and is heard by station 30, it will be passed a token. When the station has completed sending data frames it passes the token to its successor station 20.

The NS variable is also loaded whenever the station is in_ring and receives a set_successor frame addressed to it.

Whenever the value of NS is changed, an LM_EVENT notify primitive is invoked for the management entity, if enabled.

NOTE: Once a station believes it is a member of an active logical ring and knows the value of NS it no longer reloads NS when a contention window is opened spanning the station's address. The reason is that under error recovery conditions, stations will send solicit_successor_2 frames addressed to themselves that open response windows for all stations. If all stations reset their NS variables at this point, any logical ring that existed would collapse.

NS_known. A Boolean variable that indicates whether the station thinks it knows the address of its successor. NS_known is set true when the station wins the demand_in contention process or is in_ring and receives a set_successor frame addressed to it such that the data unit (the new NS) does not equal TS.

NS_known is set false whenever the station leaves the logical ring or when NS is set equal to TS by a set_successor frame.

If NS_known is false, the station is not a member of an active logical ring and the values of NS and PS are undefined.

PS—previous station's address. The variable is set to the value of the source address of the last token addressed to the station.

If a who_follows frame is heard, the contents of the data_unit field of the frame are compared with the contents of PS. If they are equal, the station responds to the who_follows request with a set_successor frame.

An example will clarify the use of PS. If a logical ring contains stations with addresses 30, 25, and 20, the station with address 20 will have 25 in its PS register, since this is the address of the station that sends it the token. If station 25 fails, when station 30 tries to send the token to station 25 it will get no response. After two tries at passing the token, station 30 sends a "who_follows 25?" frame. Station 20 responds by sending a "set your successor to 20" frame. In this manner the failed station, 25, is quickly patched out of the ring.

max_access_class. An integer constant used to initialize the sequencing of the processing of the pending frame queues. The value of max_access_class is 6, the highest priority access class.

access_class. An integer used to sequence through the access classes while transmitting data frames.

The first, or highest priority, access_class equals the value of max_access_class (that is, 6). The variable access_class is then decremented (by 2) through all classes until less than zero, and then the station performs its ring maintenance functions and passes the token.

in_ring. A Boolean variable set true when the station wins the demand_in token contention process or when the station successfully completes the claiming token process and set false when the station removes itself from the logical ring.

sole_active_station. A Boolean variable used to mute stations that have defective receivers. If a station's receiver becomes inoperative in an undetected manner, the station would otherwise disrupt the operation of the system by continually claiming the token and then soliciting a successor station.

If the sole_active_station variable is true, a station is prevented from entering the claiming token process unless it has data to send. Thus, a station with an inoperative receiver and no data to send will remain passively out of the ring.

If a station is a member of the ring and its receiver fails, it will be unable to hear its successor claiming the token. The station will cycle through the token-passing recovery algorithm, quickly reaching the point where it has sent a solicit_successor_2 frame addressed to itself and received no intelligible response. At this point, the station sets sole_active_station true and becomes passive.

Sole_active_station is set false whenever the station hears a valid frame from another station.

lowest_address. A Boolean variable that is set true if the station's successor address is greater than the station's address and set false otherwise.

At any one time there should be only one station in the logical ring with lowest_address set true. This is the station with the lowest_address of all those currently in the logical ring. When this station opens response windows during a token pass, it shall open two windows. The first window is used by stations that have an even lower address that wish to enter the ring. The second window is used by stations that have an address higher than the next normal token holder (the station currently with the highest address in the ring).

Lowest_address is computed and set by a station whenever NS is changed.

Lowest_address is used for a second purpose unrelated to token passing. If the token-holding station fails, another station recovers the token. The bus_idle_timer is a "watchdog" timer. If nothing is heard by a station for an interval greater than this timer, the claim token process is started.

In an effort to minimize interference during the claiming process, one station is selected to use a shorter bus_idle_timer value than the other stations. This station recovers all lost token failures, except ones it causes. The station with lowest_address true is presumed unique, so it is assigned this role.

just_transmitted. A Boolean variable that is set true when the station passes the token and set false if the station hears a valid frame from another station.

Just_transmitted is used to detect duplicate addressing failures on the network. If a station hears a valid MAC control frame with a source address equal to its own address and just_transmitted is false, the station cannot have sent the frame. If such a frame is heard, the MAC sublayer notifies the management entity of the detection of another station on the network using the same MAC address; the ACM then enters the OFFLINE state.

heard_successor. A Boolean variable used by a soliciting station to record the reception of a valid set_successor frame during a demand_in sequence. If heard_successor is true when the response_window_timer expires, the resolution was successful.

claim_pass_count. An integer with a range from 0 to max_pass_count, used as an index to TS to select two bits from the station's address. The value of the selected bits (times twice the slot_time) determines the length of the data_unit field of the claim_token frame to be sent. After each claim_token frame the value of the variable claim_pass_count is incremented by one.

contend_pass_count. An integer with a range from 0 to max_pass_count, used as an index to TS to select two bits from the station's address. The one's complement of the selected bits (times the slot_time) determines the length of time a station delays in the DEMAND_IN state after receiving a resolve_contention frame. If no other frames are heard before the contention timer expires, the station sends a set_successor frame to the token holder, increments the value of contend_pass_count, goes to the DEMAND_DELAY state, and waits for the token or for another resolve_contention frame.

contention_delay(cycle). An integer function that returns the value 0, 1, 2, or 3. This value is based upon the one's complement of a pair of address bits that are indicated by the cycle in the address sort. The value is used to control the number of slot_times the station delays before transmitting when demanding entry to the logical ring.

resolution_pass_count. An integer with a range from 0 to max_pass_count, used to count the number of resolve_contention passes the token-holding station makes. If the counter reaches the value of max_pass_count, the resolution process is abandoned and the ACM sequences to the next pass_token substate.

inter_solicit_count. An integer in the range 0 to 2^8–1, which determines when a station shall open a response window. Before passing the token, the value of inter_solicit_count is checked. If the value is zero, a new successor is solicited by opening a response window. If the value is not zero, the counter is decremented and the token is passed. Whenever anything is heard during the response windows following the solicit_successor frame, the counter value is set to zero so that it will again be zero when the station next receives and passes the token. Thus receipt of a set_successor frame during a response window causes the station to reopen the response window before that next token pass (ring maintenance timer permitting).

remaining_retries. An integer in the range 0 to max_retry_limit, used to count the number of retransmissions upon timeout of a frame specifying request_with_response.

suppress_FCS. A Boolean variable used within the ACM to indicate that the current frame should be transmitted without having the TxM append a FCS.

transmitter_fault_count. An integer in the range of 0 to 7, used to infer that the station's transmitter has probably failed and thus that the station's transmissions cannot be heard correctly by other stations on the network.

The value of transmitter_fault_count is incremented each time the station sequences to the end of the token contention process or fails to pass the token to any successor. Neither of these failures occurs during normal operation. The value of transmitter_fault_count is reset to zero if the station either wins the demand_in token contention process or successfully solicits a new_successor, since such an event indicates that another station correctly heard a transmission from the station.

If the value of transmitter_fault_count is incremented to max_transmitter_fault_count, the station reports a faulty_transmitter to Management and enters the OFFLINE state. The value of max transmitter_fault_count is 7, allowing for an occasional protocol sequencing impasse due to noise. If the station cannot enter the ring or pass the token (if already in the logical ring) seven times in a row, the inference made is that something has failed in the station, probably in the transmitter, and so the station removes itself from the logical ring.

first_time. A Boolean variable that controls processing of noise bursts in the CHECK_TOKEN_PASS_state. The variable is set true upon entry to the CHECK_TOKEN_PASS_state and is set false when echoes_expected is false and a noise burst is heard. If a noise burst is heard when first_time is false, the station returns to the IDLE state.

echoes_expected. A Boolean variable used when attempting to pass the token. It aids the station in distinguishing between hearing its own token transmission and hearing an immediately transmitted, garbled transmission from the new token holder. Echoes_expected is set true when the token is passed and set false when the station hears, or presumes it hears, its own token echoed.

pass_state. A multistate variable used to control the operation of the PASS_TOKEN state substates. The action taken in the state depends on the value of the variable pass_state as follows. (The actions are listed in the order taken by a station soliciting successors and failing.)

pass_state value	action
solicit_successor	Send solicit_successor frame. Enter AWAIT_RESPONSE state.
pass_token	Send token to successor. Enter CHECK_TOKEN_PASS_ state.
repeat_pass_token	Same action as pass_token substate.
who_follows	Send who follows frame. Enter AWAIT_RESPONSE state.
repeat_who_follows	Same action as who_follows substate.
solicit_any	Send solicit_successor_2 frame with DA = TS, opening 2 response windows that span all other stations. Enter AWAIT_RESPONSE state.
total_failure	Set sole_active_station true and either silently pass the token back to itself (if the station has more data to send) or enter the IDLE state. This station will not transmit again unless it has data to send or it hears a valid frame from another station.

IFM_EVENT notify. A function provided by the ACM of MAC. This function alerts the IFM to a significant occurrence within the ACM.

LM_EVENT notify. A function provided by the MAC sublayer that alerts the management entity to a significant occurrence within the MAC sublayer.

7.1.7 Initialization Variables. The value of the variables described in this section are copied into corresponding operational variables when the station is initialized. Each initialization variable has the same range as the correspondingly named operational variable.

init_TS. Copied to TS at initialization.

init_slot_time. Copied to slot_time at initialization.

init_max_inter_solicit_count. Copied to max_inter_solicit_count at initialization.

init_max_retry_limit. Copied to max_retry_limit at initialization.

init_hi_pri_token_hold_time. Copied to hi_pri_token_hold_time at initialization.

init_target_rotation_time(access class). Copied to target_rotation_time (access class) at initialization.

init_ring_maintenance_timer_initial_value. Copied to ring_maintenance_timer_initial_value at initialization.

init_in_ring_desired. Copied to in_ring_desired at initialization.

7.2 Access Control Machine (ACM) Formal Description. The ACM is described formally in this subsection. The descriptive model used is a hybrid of finite state machine and procedural programming language notations and concepts. The ACM may be in any one of 11 major states at any given time. Its actions and its transition into another state are determined by its evaluation of the MAC state variables discussed in the previous subsection.

The ACM in these tables is specified to operate properly in the full duplex environment. That is, the machine will not be mislead by the reception of its own transmissions. Furthermore, the ACM will not work properly in a half duplex environment. It is up to the implementor to make an implementation that is not full duplex act as if it were full duplex.

The conditions evaluated and the actions taken for each state machine transition are expressed using syntax based upon the Ada®[14] Programming Language.[15] The ACM variables and other information units used by the ACM are described using Ada-like data declarations. Where an information unit cannot be expressed using Ada syntax, a descriptive statement, enclosed in angle brackets (<>), is used. The relaxed Ada-like style is adopted for readability and is not intended to represent compilable statements.

The following are presented: programming language style declarations of most of the variables, functions, and procedures used in the state machine description; a summary of the ACM state transitions (arcs); and the ACM state transition tables.

The state transition tables present the current state, transition name, and the next state on one line at the beginning of each transition description. Below this line are

(1) The conditions that must be true for the particular transition to be taken, and
(2) The actions to be taken before going to the prescribed next state.

NOTE: For the purpose of description, the ACM is presumed to be instantaneous with respect to external events. Thus, all conditions are evaluated simultaneously.

[14] Ada is a registered trademark of the United States Government Department of Defense, Ada Joint Program Office.

[15] The definition of the Ada language is found in ANSI/MIL STD 1851A-1983, Reference Manual for the Ada Programming Language. ANSI documents are available from the Sales Department, American National Standards Institute, 1430 Broadway, New York, NY 10018, USA.

7.2.1 ACM Data Declarations

NOTE: Constants are represented as numeric values rather than the transmission-order bit strings of Section 4.

-- constants and parameters

```
-- These are all treated as constants in this declaration,
--    with the understanding that many of them are actually
--    station parameters set by Management.

symbol_time: constant duration := 1 / <network data rate>;
octet_time: constant duration := 8 * symbol_time;
address_length: constant (16,48);      -- network parameter
max_pass_count: constant ( 9,25) := (address_length/2) + 1;
max_address: constant (2**16 - 1, 2**48 - 1)
                      := (2**address_length - 1);
timer_size: constant integer := (2**21 1);
max_frame_length: constant integer := (2**13 - 1); -- 8191
max_data_unit_length: constant integer :=
                (max_frame_length - (5 + 2*address_length/8) );
               -- This gives 8174 for 48 bit addresses,
               --    and 8182 for 16 bit addresses.
max_access_class: constant integer :=  6;
ring_maintenance: constant integer := -2;
                                 -- min value of access_class
max_transmitter_fault_count: constant integer := 7;
init_TS: constant integer range 0 .. max_address
                                     := <station specific>;
init_slot_time: constant duration range octet_time ..
                (2**13 - 1) * octet_time := <network specific>;
init_in_ring_desired: constant Boolean := <station specific>;
init_max_inter_solicit_count: constant integer
            range 16 .. 255 := <station specific>;
init_max_retry_limit: constant integer
            range 0..7 := 3; -- typ value
init_hi_pri_token_hold_time: constant integer
            range 0 .. (2**16 - 1):= <station specific>;
init_target_rotation_time: array ( -2,0,2,4 ) constant integer
            range 0 .. timer_size := <station specific>;
init_ring_maintenance_timer_initial_value constant integer
            range 0 .. timer_size := <station specific>;
```

-- type definitions

```
type slot_times is integer range 0..7;
type address is integer range 0 .. max_address;
type pass_count is integer range 0 .. max_pass_count;
type LLC_priority is integer range 0 .. 7;
type MAC_priority is (0, 2, 4, 6);

type MAC_actions is ( request_with_no_response,
                      request_with_response,
                      response );
```

```
type MAC_symbol is (zero, one, non_data,
                    pad_idle, bad_signal, silence);
       -- Also referred to as (0, 1, N, P, B, S).

subtype data_symbol is MAC_symbol range zero .. one;
       -- Also referred to as D.

type data_octet is array (0 .. 7) of data_symbol;

type data_unit_type is
     record
         length: integer range 0 .. max_data_unit_length:=0;
         data: array (0 .. max_data_unit_length) of data_octet;
     end record;

type frame_identifier is ( MAC_control,      -- 2#00#
                           data );      -- 2#10# or 2#01# or 2#11#

type frame_ctrl_typ is (fid: frame_identifier)
     record
         case fid is
            when MAC_control => mftyp: integer range 0 .. 63;
            when data =>
                     record
                           MAC_action: MAC_actions :=
                                       request_with_no_response;
                           priority: LLC_priority := 7;
                     end record;
         end case;
     end record;

for frame_ctrl_typ use                 -- Define subfield structure.
    record at mod 8
        fid        at 0 range 0 .. 1;
        mftyp      at 0 range 2 .. 7;
        MAC_action at 0 range 2 .. 4;
        priority   at 0 range 5 .. 7;
    end record;

type FCS_type is array ( 0 .. 3 ) of data_octet;

type general_frame is
     record
         FC: frame_ctrl_typ;
         DA: address;
         SA: address;
         data_unit: data_unit_type;
         FCS: FCS type;
     end record;
-- State 8 ( PASS_TOKEN ) has seven substates, each
--   representing a different aspect of the token-passing
--   process. These substates are kept track of by the state
--   variable pass_state, which is of type tok_pass_substate.
```

```
   type tok_pass_substate is
          ( solicit_successor,
            pass_token
            repeat_pass_token,
            send_who_follows,
            repeat_who_follows,
            solicit_any,
            total_failure );
```

-- constant definitions

```
     null_data_unit: data_unit type; -- has a length of zero
        claim_token: frame_ctrl_typ := ( MAC_control,
                                         mftyp => 2#000000# );
solicit_successor_1: frame_ctrl_typ := ( MAC_control,
                                         mftyp => 2#0100000# );
solicit_successor_2: frame_ctrl_typ := ( MAC_control,
                                         mftyp => 2#010000# );
        who_follows: frame_ctrl_typ := ( MAC_control,
                                         mftyp => 2#110000# );
 resolve_contention: frame_ctrl_typ := ( MAC_control,
                                         mftyp => 2#001000# );
              token: frame_ctrl_typ := ( MAC_control,
                                         mftyp => 2#000100# );
      set_successor: frame_ctrl_typ := ( MAC_control,
                                         mftyp => 2#001100# );
```

-- procedures and packages

```
-- generic timer package
generic
     type tick is integer range 0 .. timer_size;
     type time_unit is (octet_time, slot_time);
package timer( time_int: time_unit := octet_time )
     procedure start( init: in tick );
     function value return tick;
     function expired return Boolean;
end timer;
    -- Timers count from a “start” value down to zero.
    -- Timers are “expired” when their value is zero.
    -- Time_int defines the timer granularity or counting interval.
    -- Each count corresponds to one “time unit,” which cor-
    --    responds either to one slot_time or one octet_time
    --    on the medium.
```

```
function claim_data_unit( cycle: in pass_count )
                                 return data_unit_type;
    -- Returns a data unit with arbitrary contents
    --    and with its length (and thus its transmission time)
    --    equal to 0, 2, 4 or 6 slot_times.

    -- This length depends upon the two bits of the
    --    address indicated by the cycle of the address sort.

function max_bus_idle return slot_times;
    -- Returns a 6 or 7 slot_time time interval
    --    if lowest_address then return 6
    --                      else return 7.

function response_delay( below_TH: in Boolean ) return slot_times;
    -- Returns a 0 or 1 slot_time time interval
    --     if below_TH then return 0
    --                 else return 1.

function contention_delay( cycle: in pass_count ) return slot_times;
    -- Returns a 0, 1, 2, or 3 slot_time time interval
    --    based on the one's complement of two bits of the
    --    address indicated by the cycle of the address sort.

function send_pending( queue: in MAC_priority ) return Boolean;
    -- Returns TRUE if there is a pending frame
    --    in the indicated pending frame queue.

function any_send_pending return Boolean;
    -- Returns TRUE if there is a pending frame in
    --    any of the pending frame queues.

procedure get_pending_frame( queue: in MAC_priority;
                             FC: out frame_ctrl_typ;
                             DA: out address;
                             SA: out address;
                             data_unit: out data_unit_type;
                             FCS_suppression: out Boolean );
    -- Returns the first frame from the indicated pending
    --    frame queue.

procedure get_pending_response( FC: out frame_ctrl_typ;
                                DA: out address;
                                SA: out address;
                                data_unit: out data_unit_type
                                FCS_suppression: out Boolean );
    -- Returns the response frame created by the MAC-user after
    --    receiving a request_with_response frame. This "queue"
    --    can have a maximum of one entry.

procedure send( FC: in frame_ctrl_typ;
                DA: in address := <TS OR any individual address>;
                SA: in address := <TS>;
                data_unit: in data_unit type := null_data_unit;
                FCS suppression: in Boolean := false );

    -- Sends the indicated frame by passing it symbol by symbol
    --    to the Physical Layer for transmission.
    -- If FCS_suppression is requested, an FCS is not appended and
    --    the last four octets of the data_unit serve as the FCS.
```

```
function response_received return Boolean;
    -- Returns TRUE when IFM has received response to outstanding
    --    request_with_response frame.
```

-- Boolean state variables

```
          bus_quiet: Boolean;
     echoes_expected: Boolean;
          first_time: Boolean;
     heard_successor: Boolean;
             in_ring: Boolean;
     in_ring_desired: Boolean;
    just_transmitted: Boolean;
      lowest_address: Boolean;
         noise_burst: Boolean;
            NS_known: Boolean;
            power_ok: Boolean;
       Rx_data_frame: Boolean;
   Rx_protocol_frame: Boolean;
 sole_active_station: Boolean;
        suppress_FCS: Boolean;
```

-- address variables

```
         TS: address;             -- This Station
         NS: address;             -- Next Station
         PS: address;             -- Previous Station
         TH: address;             -- Token Holder
destination: address;             -- destination address
     source: address;             -- source address
```

-- counters

```
       claim_pass_count: pass_count;
     contend_pass_count: pass_count;
  resolution_pass_count: pass_count;
    inter_solicit_count: integer range 0 .. 255;
transmitter_fault_count: integer
                           range 0 .. max transmitter_fault_count;
```

-- protocol or frame data units

```
cdu: data_unit_type;
sdu: data_unit_type;
```

-- operational variables

```
access_class: integer range -2 .. max_access_class;
max_inter_solicit_count: integer range 16 .. 255;
max_retry_limit: integer range 0 .. 7;
min_post_silence_preamble_length: integer range 1 .. 15;
pass_state: tok_pass_substate;
remaining_retries: integer range 0 .. max_retry_limit;
Rx_frame: general_frame;
slot_time: duration range octet_time .. (2**13 - 1) * octet_time;
hi_pri_token_hold_time: integer range 0 .. (2**16 - 1);
target_rotation_time: array ( -2,0,2,4 )
                      integer range 0 .. timer_size;
ring_maintenance_timer_initial_value: integer range 0 .. timer_size;
```

7.2.2 ACM State Transition Summary

Current State	Transition Name	Next State
0. OFFLINE	initialize	1. IDLE
1. IDLE	new_successor	1. IDLE
1. IDLE	no_successor_1	1. IDLE
1. IDLE	receive_token	5. USE_TOKEN
1. IDLE	entry_demand_in	2. DEMAND_IN
1. IDLE	repair_demand_in	2. DEMAND_IN
1. IDLE	own_frame_1	1. IDLE
1. IDLE	duplicate_address_1	0. OFFLINE
1. IDLE	ring_patch	2. DEMAND_IN
1. IDLE	non_idle_bus	1. IDLE
1. IDLE	no_token	4. CLAIM_TOKEN
1. IDLE	exit_ring	1. IDLE
1. IDLE	idle_station	1. IDLE
1. IDLE	transmit_response	1. IDLE
1. IDLE	other_heard	1. IDLE
2. DEMAND_IN	continue_contention	3. DEMAND_DELAY
2. DEMAND_IN	lost_contention_2	1. IDLE
3. DEMAND_DELAY	contention_delay	2. DEMAND_IN
3. DEMAND_DELAY	won_contention	5. USE_TOKEN
3. DEMAND_DELAY	won_patch	5. USE_TOKEN
3. DEMAND_DELAY	lost_contention_3	1. IDLE
3. DEMAND_DELAY	end_contention	1. IDLE
3. DEMAND_DELAY	end_all_contention	0. OFFLINE
3. DEMAND_DELAY	ignore_contenders	3. DEMAND_DELAY
3. DEMAND_DELAY	ignore noise	3. DEMAND_DELAY
3. DEMAND_DELAY	long_bus_idle	1. IDLE
4. CLAIM_TOKEN	lose_address_sort	1. IDLE
4. CLAIM_TOKEN	continue_address_sort	4. CLAIM_TOKEN
4. CLAIM_TOKEN	win_address_sort	5. USE_TOKEN
5. USE_TOKEN	send_frame	6. AWAIT_IFM_RESPONSE
5. USE_TOKEN	no_send	7. CHECK ACCESS_CLASS

7.2.2 ACM State Transition Summary *(Continued)*

Current State	Transition Name	Next State
6. AWAIT_IFM_RESPONSE	no_timeout	5. USE_TOKEN
6. AWAIT_IFM_RESPONSE	retry	6. AWAIT_IFM_RESPONSE
6. AWAIT_IFM_RESPONSE	no_response_6	5. USE_TOKEN
6. AWAIT_IFM_RESPONSE	unexpected_frame_6	1. IDLE
6. AWAIT_IFM_RESPONSE	own_frame_6	6. AWAIT_IFM_RESPONSE
7. CHECK_ACCESS_CLASS	next_access_class	5. USE_TOKEN
7. CHECK_ACCESS_CLASS	leave_ring	8. PASS_TOKEN
7. CHECK_ACCESS_CLASS	do solicit_successor	8. PASS_TOKEN
7. CHECK_ACCESS_CLASS	do_pass_token	8. PASS_TOKEN
7. CHECK_ACCESS_CLASS	do_solicit_any	8. PASS_TOKEN
8. PASS_TOKEN	open_1_response_window	10. AWAIT_RESPONSE
8. PASS_TOKEN	open_2_response_windows	10. AWAIT_RESPONSE
8. PASS_TOKEN	pass_token	9. CHECK_TOKEN_PASS
8. PASS_TOKEN	who_follows query	10. AWAIT_RESPONSE
8. PASS_TOKEN	solicit_any	10. AWAIT_RESPONSE
8. PASS_TOKEN	silent_pass_to_self	5. USE_TOKEN
8. PASS_TOKEN	sole_station	1. IDLE
8. PASS_TOKEN	no_successor_8	1. IDLE
8. PASS_TOKEN	no_future	0. OFFLINE
9. CHECK_TOKEN_PASS	token_pass_failed	8. PASS_TOKEN
9. CHECK_TOKEN_PASS	own_frame_9	9. CHECK_TOKEN_PASS
9. CHECK_TOKEN_PASS	own_token	9. CHECK_TOKEN_PASS
9. CHECK_TOKEN_PASS	probably_own_token	9. CHECK_TOKEN_PASS
9. CHECK_TOKEN_PASS	pass_ok	1. IDLE
9. CHECK_TOKEN_PASS	not_sure	9. CHECK_TOKEN_PASS
9. CHECK_TOKEN_PASS	probably_ok	1. IDLE
10. AWAIT_RESPONSE	no_response_10	8. PASS_TOKEN
10. AWAIT_RESPONSE	resolution_succeeded	8. PASS_TOKEN
10. AWAIT_RESPONSE	own_address_10	10. AWAIT_RESPONSE
10. AWAIT_RESPONSE	hear_successor	10. AWAIT_RESPONSE
10. AWAIT_RESPONSE	ignore response	10. AWAIT_RESPONSE
10. AWAIT_RESPONSE	unexpected_frame_10	1. IDLE
10. AWAIT_RESPONSE	send_resolve	10. AWAIT_RESPONSE
10. AWAIT_RESPONSE	resolution_failed	8. PASS_TOKEN

7.2.3 State Transition Tables

7.2.3.0 Offline State

```
Current State          Transition Name      Next State
    Exit Condition
         Action Taken
-----------------------------------------------------------------------
0. OFFLINE             initialize           1. IDLE

    power_ok
    AND LM_ACTION invoke ( initialize )      -- see 3.2.11.2
         -- MAC configuration information supplied by Management.
         -- It is Management's responsibility to initiate the
         --   MAC sublayer on station power-up.
         -- This includes values for
         --      init_slot_time      (as octet-time multiple)
         --      init_TS      (This Station's address)
         --      init_address_length      (implicit in TS)
         --      init_target_rotation_time         (one each for
         --           access_class 4, 2, 0, and ring_maintenance, -2)
         --      init_hi_pri_token_hold_time
         --      init_max_inter_solicit_count
         --      init_ring_maintenance_timer_initial_value
         --   and other variables that may need to be initialized.

         -- Instantiate timers used in the ACM:

         package bus_idle_timer        is new timer( time_int
                                              => slot_time );
         package contention_timer      is new timer( time_int
                                              => slot_time );
         package claim_timer           is new timer( time_int
                                              => slot_time );
         package response_window_timer is new timer( time_int
                                              => slot_time );
         package token_pass_timer      is new timer( time_int
                                              => slot_time );
         package token_hold_timer      is new timer( time_int
                                              => octet_time );
         package token_rotation_timer
                 is array( ring_maintenance .. 4 ) of new timer
                    ( time_int => octet_time ); -- only -2,0,2,4 used

-- Copy initialization parameters to current variables.
     TS := init_TS;
     slot_time := init_slot_time;
     max_inter_solicit_count := init_max_inter_solicit_count;
     max_retry_limit := init_max_retry_limit;
     hi_pri_token_hold_time := init_hi_pri_token_hold_time;
     target_rotation_time := init_target_rotation_time;
     ring_maintenance_timer_initial_value
       := init_ring_maintenance_timer_initial_value;
     in_ring desired := init_in_ring_desired;
     min_post_silence_preamble_length
       := <corresponding physical parameter>;
           -- Get preamble length from Physical Layer.
```

```
-- Set initial value of internal variables.
     in_ring := false;
     NS_known := false;
     lowest_address := false;
     sole_active_station := false;
     just_transmitted := false;
     transmitter_fault_count := 0;

     LM_EVENT notify( no_successor );
     bus_idle_timer.start( max_bus_idle );
```

7.2.3.1 Idle State

```
Current State          Transition Name      Next State
    Exit Condition
         Action Taken
----------------------------------------------------------------
1. IDLE                receive_token        5. USE_TOKEN

     Rx_protocol_frame
     AND Rx_frame.FC = token
     AND Rx_frame.SA /= TS
     AND Rx_frame.DA = TS
     AND in_ring

         -- token receipt
         sole_active_station := false;
         PS := Rx_frame.SA;      -- set predecessor
         -- set access_class to highest level
         access_class := max_access_class;
         token_hold_timer.start( hi_pri_token_hold_time );
         Rx_protocol frame := false;
----------------------------------------------------------------
1. IDLE                new_successor        1. IDLE

     Rx_protocol_frame
     AND Rx_frame.FC = set_successor
     AND Rx_frame.SA /= TS
     AND Rx_frame.DA = TS
     AND Rx_frame.data_unit /= TS
     AND in_ring

         -- new successor heard
         NS := Rx_frame.data_unit;
         lowest_address := ( NS . TS );
         NS_known := true;
         sole_active_station := false;
         inter_solicit_count := 0;
         just_transmitted := false;
         LM_EVENT notify( new_successor, NS );
         bus_idle_timer.start( max_bus_idle );
         Rx_protocol_frame := false;
----------------------------------------------------------------
```

```
Current State          Transition Name       Next State
    Exit Condition
         Action Taken
-------------------------------------------------------------------
1. IDLE                no_successor_1        1. IDLE

    Rx_protocol_frame
    AND Rx_frame.FC = set_successor
    AND Rx_frame.SA /= TS
    AND Rx_frame.DA = TS
    AND Rx_frame.data_unit = TS

         -- successor is myself - one station ring
         NS := Rx_frame.data_unit;
         NS_known := false; -- no one station rings
         sole_active_station := false;
         inter_solicit_count := 0;
         just_transmitted := false;
         LM_EVENT notify( no_successor );
         bus_idle_timer.start( max_bus_idle );
         Rx_protocol_frame := false;
-------------------------------------------------------------------
1. IDLE                entry_demand_in       2. DEMAND_IN

    Rx_protocol_frame
    AND Rx_frame.SA /= TS
    AND NOT ( in_ring AND NS_known )   -- not part of functioning ring
    AND ( in_ring desired OR any_send_pending )

    AND ( ( Rx_frame.FC=solicit_successor_1
            AND Rx_frame.SA > TS
            AND Rx_frame.DA < TS )

        OR ( Rx_frame.FC = solicit_successor_2
             AND ( Rx_frame.SA > TS
                   OR Rx_frame.DA < TS ) ) )

         -- Solicit_successor heard which allows this station
         --   to respond as a potential successor.
         -- Contention is anticipated.
         contend_pass_count := 0;
         inter_solicit_count := 0;
         just_transmitted := false;
         TH := Rx_frame.SA;   -- save address of current token holder
         NS := Rxframe.DA;
         lowest_address := ( TS < NS );
         sole_active_station := false;
         -- Initialize token_rotation timers to “expired.”
         token_rotation_timer( 4 ).start( 0 );
         token_rotation_timer( 2 ).start( 0 );
         token_rotation_timer( 0 ).start( 0 );
         token_rotation_timer( ring_maintenance ).start(
                             ring_maintenance_timer_initial_value );
         contention_timer.start( response_delay( TS < TH ) );
         Rx_protocol_frame := false;
-------------------------------------------------------------------
```

```
Current State          Transition Name      Next State
    Exit Condition
         Action Taken
---------------------------------------------------------------
1. IDLE                repair_demand_in     2. DEMAND_IN

    Rx_protocol_frame
    AND Rx_frame.SA /= TS
    AND in_ring
    AND NS_known
    AND ( in_ring_desired OR any_send_pending )

    AND ( ( Rx_frame.FC = solicit_successor_1
            AND Rx_frame.SA > TS
            AND Rx_frame.DA < TS )

        OR ( Rx_frame.FC = solicit_successor_2
            AND ( Rx_frame.SA > TS
                  OR Rx_frame.DA < TS ) ) )

         -- This station was in_ring, but is now being skipped
         -- for some reason.
         -- Solicit_successor heard which allows this station to
         --   respond as a potential successor.
         -- Contention is anticipated.
         contend_pass_count := 0;
         inter_solicit_count := 0;
         just_transmitted := false;
         TH := Rx_frame.SA;
         sole_active_station := false;
         contention_timer.start( response_delay( TS < TH ) );
         Rx_protocol_frame := false;
---------------------------------------------------------------
1. IDLE                own_frame_1          1. IDLE

    ( ( Rx_protocol_frame AND just_transmitted)
      OR Rx_data_frame )
    AND Rx_frame.SA = TS

         -- Ignore own frame.
         bus_idle_timer.start( max_bus_idle );
         Rx_protocol_frame := false;
```

```
Current State          Transition Name      Next State
    Exit Condition
         Action Taken
-----------------------------------------------------------------
1. IDLE                duplicate_address_1  0. OFFLINE

    Rx_protocol_frame
    AND Rx_frame.SA = TS
    AND NOT just_transmitted

         -- Detect a duplicate station address report to
         --  Management LM_EVENT notify( duplicate_address );
-----------------------------------------------------------------
1. IDLE                ring_patch           2. DEMAND_IN

    Rx_protocol_frame
    AND Rx_frame.FC = who_follows
    AND Rx_frame.SA /= TS
    AND Rx_frame.data_unit = PS
    AND in_ring

         -- Recognize who_follows message from predecessor's
         --   predecessor. Respond with a set_successor message to
         --   patch a failed station (TS's predecessor) out of the
         --   ring. Contention is possible.
         contend_pass_count := 0;
         inter_solicit_count := 0;
         just_transmitted := false;
         TH := Rx_frame.SA;
         sole_active_station := false;
         contention_timer.start( 0 );
         Rx_protocol_frame := false;
-----------------------------------------------------------------
1. IDLE                non_idle_bus         1. IDLE

noise burst
         -- Bus not idle.
         -- Reset bus_idle_timer.
         bus_idle_timer.start( max_bus_idle );
         noise_burst := false;
-----------------------------------------------------------------
```

```
Current State          Transition Name      Next State
    Exit Condition
         Action Taken
-----------------------------------------------------------------------
1. IDLE                no_token             4. CLAIM_TOKEN

    bus_idle_timer.expired
    AND bus_quiet
    AND NOT ( Rx_protocol_frame OR Rx_data_frame OR noise burst )
    AND ( any send pending
          OR in_ring
          OR ( in_ring desired AND NOT sole_active_station ) )
                                -- in case TS's receiver has failed

         -- No_token, so claim one.
         claim_pass_count := 1;
         cdu := claim data_unit( claim_pass_count );
         send( FC => claim_token,
               data_unit => cdu );
         claim timer.start ( 1 );
-----------------------------------------------------------------------
1. IDLE                exit_ring            1. IDLE
-----------------------------------------------------------------------

    Rx_protocol_frame
    AND Rx_frame.SA /= TS
    AND in_ring
    AND NOT (in_ring desired OR any_send_pending)

    AND (  ( Rx_frame.FC = solicit_successor_1
           AND Rx_frame.SA > TS
           AND Rx_frame.DA < TS )

        OR ( Rx_frame.FC = solicit_successor_2
           AND ( Rx_frame.SA > TS
               OR Rx_frame.DA < TS ) ) )

       -- This station has now left the ring.
       just_transmitted := false;
       sole_active_station := false;
       in_ring := false;
       NS_known := false
       LM_EVENT notify( no_successor );
       bus_idle_timer.start( max_bus_idle );
       Rx_protocol_frame := false;
```

```
Current State         Transition Name      Next State
    Exit Condition
         Action Taken
------------------------------------------------------------------------
1. IDLE               idle_station         1. IDLE

    bus_idle_timer.expired
    AND bus quiet
    AND NOT ( Rx_protocol_frame OR Rx_data_frame OR noise burst )
    AND NOT any send pending
    AND NOT in_ring
    AND ( NOT in_ring desired OR sole_active_station )

        -- Ring has collapsed.
        -- Reset bus_idle_timer just to keep it running.
        bus_idle_timer.start( max bus idle );
------------------------------------------------------------------------
1. IDLE               transmit_response    1. IDLE

    Rx_data_frame
    AND Rx_frame.SA /= TS
    AND Rx_frame.DA = TS
    AND Rx_frame.FC.MAC_action = request_with_response

      -- A request_with_response was received.
      -- The station shall begin a response within
      --    two slot_times if this option is supported.
      sole_active_station := false;
      get pending response(
                          FC => frame control,
                          DA => destination,
                          SA => source,
                          data_unit => sdu,
                          FCS suppression => suppress_FCS );
      send( FC => frame control,
            DA => destination,
            SA => source,
            data_unit => sdu,
            FCS suppression => suppress_FCS );
      just_transmitted := true;
      bus_idle_timer.start( max_bus_idle );
------------------------------------------------------------------------
```

```
Current State          Transition Name       Next State
    Exit Condition
         Action Taken
----------------------------------------------------------------
1. IDLE                other_heard           1. IDLE

    -- NOTE: The following codifies the converse of the union
    -- of all of the above frame received conditions;
    --   i.e., this is supposed to have the same effect as an
    -- “otherwise” or a “none_of_the_above” clause.
    ( Rx_data_frame
        AND NOT ( Rx_frame.FC.MAC_action = request_with_response
                  AND Rx_frame.DA = TS )
        AND NOT Rx_frame.SA = TS )

    OR ( Rx_protocol_frame
      AND NOT        -- any of the following

         ( Rx_frame.SA = TS

           OR ( Rx_frame.FC = set_successor
                AND Rx_frame.DA = TS
                AND in_ring )

           OR ( Rx_frame.FC = token
                AND Rx_frame.DA = TS
                AND in_ring )

           OR ( Rx_frame.FC = solicit_successor 1
                AND ( in_ring OR in_ring desired
                      OR any send pending )
                AND Rx_frame.DA < TS
                AND Rx_frame.SA > TS )

                      OR ( Rx_frame.FC = solicit_successor_2
                AND ( in_ring
                      OR in_ring desired OR any send pending )
                          AND ( Rx_frame.DA < TS
                      OR Rx_frame.SA > TS )

           OR ( Rx_frame.FC = who_follows
                AND Rx_frame.data_unit = PS
                AND in_ring ) ) )

      just_transmitted := false;
      sole_active_station := false;
      bus_idle_timer.start( max bus idle );
      Rx_protocol_frame := false;
```

7.2.3.2 Demand In State

```
Current State          Transition Name      Next State
    Exit Condition
         Action Taken
-----------------------------------------------------------------------
2. DEMAND_IN           continue_contention  3. DEMAND_DELAY

    contention_timer.expired
    AND bus_quiet
    AND NOT ( Rx_protocol_frame OR Rx_data_frame OR noise_burst)

         send( FC => set_successor,
               DA => TH, -- most recent reliable SA
               data_unit => TS );
         bus_idle_timer.start( max_bus_idle );
-----------------------------------------------------------------------
2. DEMAND_IN           lost_contention_2    1. IDLE

   ( Rx_protocol_frame OR Rx_data_frame OR noise_burst )
         -- Drop contention and reprocess event in state 1.
         -- NO ACTION.
```

7.2.3.3 Demand In State

```
Current State          Transition Name      Next State
    Exit Condition
         Action Taken
-------------------------------------------------------------------
3. DEMAND_DELAY        contention_delay     2. DEMAND_IN

    Rx_protocol_frame
    AND Rx_frame.FC = resolve_contention
    AND contend_pass_count> max_pass_count

         -- Delay till next contention transmission.
         -- Delay time depends on bits of address,
         --   indexed into by contend_pass_count.
         contend_pass_count := contend_pass_count + 1;
         contention timer.start( contention_delay
                                    ( contend_pass_count ) );
         Rx_protocol_frame := false;
-------------------------------------------------------------------
3. DEMAND_DELAY        won_contention       5. USE_TOKEN

    Rx_protocol_frame
    AND Rx_frame.FC = token
    AND Rx_frame.DA = TS
    AND NOT NS_known

         -- Token being passed here, so this station is in the ring.
         in_ring := true;
         NS_known := true;
         PS := Rx_frame.SA;
         transmitter_fault_count := 0;
         -- Set access_class to highest level.
         access_class := max_access_class;
         LM_EVENT notify( new_successor, NS );
         token_hold_timer.start( hi_pri_token_hold_time );
         Rx_protocol_frame := false;
-------------------------------------------------------------------
3. DEMAND_DELAY        won_patch            5. USE_TOKEN

    Rx_protocol_frame
    AND Rx_frame.FC = token
    AND Rx_frame.DA = TS
    AND NS_known

         -- Token being passed here, so this station back in_ring.
         PS := Rx_frame.SA;
         transmitter_fault_count := 0;
         -- Set access_class to highest level.
         access_class := max_access_class;
         token_hold_timer.start( hi_pri_token_hold_time );
         Rx_protocol_frame := false;
-------------------------------------------------------------------
```

```
Current State          Transition Name       Next State
    Exit Condition
         Action Taken
---------------------------------------------------------------------
3. DEMAND_DELAY        lost_contention_3     1. IDLE

    ( contend_pass_count> max_pass_count
    AND ( Rx_data_frame

        OR ( Rx_protocol_frame
             AND Rx_frame.FC = token
             AND Rx_frame.DA /= TS )

        OR ( Rx_protocol_frame
             AND Rx_frame.FC /= token
             AND Rx_frame.FC /= set_successor
             AND Rx_frame.FC /= resolve_contention ) ) )

    OR ( Rx_protocol_frame
         AND Rx_frame.FC = resolve_contention
         AND contend_pass_count>= max_pass_count )

         -- Something else going on, so contention must be
         --   over. Go back to IDLE state and reprocess the event.
         -- NO ACTION.
---------------------------------------------------------------------
3. DEMAND_DELAY        end_contention        1. IDLE

    ( Rx_data_frame

      OR ( Rx_protocol_frame
           AND Rx_frame.FC = token
           AND Rx_frame.DA /= TS )
      OR ( Rx_protocol_frame
           AND Rx_frame.FC /= token
           AND Rx_frame.FC /= set_successor
           AND Rx_frame.FC /= resolve_contention ) )

    AND contend_pass_count>= max_pass_count
    AND transmitter_fault_count> max transmitter_fault_count

         -- Ran to end of contention process, so count
         --   consecutive occurrences to catch a faulty transmitter.
         -- Process the frame in IDLE.
         transmitter_fault_count := transmitter_fault_count + 1;
```

```
Current State          Transition Name       Next State
    Exit Condition
         Action Taken
-------------------------------------------------------------------
3. DEMAND_DELAY        end_all_contention    0. OFFLINE

    ( Rx_data_frame
      OR ( Rx_protocol_frame
           AND Rx_frame.FC = token
            AND Rx_frame.DA /= TS )

      OR ( Rx_protocol_frame
           AND Rx_frame.FC /= token
           AND Rx_frame.FC /= set_successor
           AND Rx_frame.FC /= resolve_contention ) )

    AND contend_pass_count>= max_pass_count
    AND transmitter_fault_count>= max transmitter_fault_count

        -- Probable “faulty transmitter” condition.
        -- Assume worst case.
        LM_EVENT notify( faulty_transmitter );
-------------------------------------------------------------------
3. DEMAND_DELAY        ignore_contenders     3. DEMAND_DELAY

    Rx_protocol_frame
    AND Rx_frame.FC = set_successor

         Rx_protocol_frame := false;
-------------------------------------------------------------------
3. DEMAND_DELAY        ignore_noise          3. DEMAND_DELAY

    noise_burst

         -- Ignore any unrecognizable signal streams.
         noise_burst := false;
-------------------------------------------------------------------
3. DEMAND_DELAY        long_bus_idle         1. IDLE

    bus_idle_timer.expired
    AND bus_quiet
    AND NOT ( Rx_protocol_frame OR Rx_data_frame OR noise_burst )
        -- Ring may have collapsed. restart the bus_idle_timer
        -- and go back to IDLE and wait to be sure.
        bus_idle_timer.start( max_bus_idle );
```

7.2.3.4 Claim Token State

```
Current State          Transition Name      Next State
    Exit Condition
         Action Taken
--------------------------------------------------------------------
4. CLAIM_TOKEN         lose_address_sort     1. IDLE

    claim_timer.expired
    AND NOT bus_quiet

         -- Other stations heard so drop from contention.
         bus_idle_timer.start( max_bus_idle );
--------------------------------------------------------------------
4. CLAIM_TOKEN         continue_address_sort 4. CLAIM_TOKEN

    claim_timer.expired
    AND bus_quiet
    AND claim_pass_count < max_pass_count

         -- Do an iteration of multipass address_sort algorithm.
         claim_pass_count := claim_pass_count + 1;
         -- Create a data_unit 0, 2, 4, or 6 slot_times long.
         cdu := claim_data_unit( claim_pass_count );
         send( FC => claim_token,
               data_unit => cdu );
         claim_timer.start( 1 );     -- one slot_time delay
--------------------------------------------------------------------
4. CLAIM_TOKEN         win_address_sort      5. USE_TOKEN

    claim_timer.expired
    AND bus_quiet
    AND claim_pass_count >= max_pass_count

         -- Token now claimed.
         in_ring := true;
         -- Initialize token rotation timers to “expired.”
         token_rotation_timer( 4 ).start( 0 );
         token_rotation_timer( 2 ).start( 0 );
         token_rotation_timer( 0 ).start( 0 );
         token_rotation_timer( ring_maintenance ).start(
                            ring_maintenance_timer_initial_value );
         inter_solicit_count := 0;
         -- Set access_class to highest level.
         access_class := max_access_class;
         token_hold_timer.start( hi_pri_token_hold_time );
```

7.2.3.5 Use Token State

```
Current State          Transition Name      Next State
    Exit Condition
         Action Taken
-----------------------------------------------------------------
5. USE_TOKEN           send_frame           6. AWAIT_IFM_RESPONSE

    send_pending( access_class )
    AND NOT token_hold_timer.expired

         -- Send one MAC-user frame of current access_class level.
         get_pending_frame(
                       queue => access_class,
                       FC => frame_control,
                       DA => destination,
                       SA => source,
                          -- SA shall equal TS except in relay nodes.
                       data_unit => sdu,
                       FCS suppression => suppress_FCS );
         remaining_retries := max_retry_limit;
         send( FC => frame_control,
               DA => destination,
               SA => source, data_unit => sdu,
               FCS suppression => suppress_FCS );
         response_window_timer.start( 3 );
-----------------------------------------------------------------
5. USE_TOKEN           no_send              7.CHECK_ACCESS_CLASS

    NOT send_pending( access_class )
    OR token_hold_timer.expired
         -- Time to relinquish the token.
         -- Set access_class to next lower level.
         access_class := access_class - 2;
```

7.2.3.6 Await IFM Response State

```
Current State          Transition Name      Next State
    Exit Condition
         Action Taken
-------------------------------------------------------------------
6. AWAIT_IFM_RESPONSE  no_timeout           5. USE_TOKEN

    frame control.MAC_action = request_with_no_response
    OR response_received
                        -- Generated by the IFM
                        --   upon receiving the response.
         -- Check for another serviceable request.
         -- NO ACTION.
-------------------------------------------------------------------
6. AWAIT_IFM_RESPONSE  retry                6. AWAIT_IFM_RESPONSE

    response_window_timer.expired
    AND bus_quiet
    AND remaining_retries> 0
    AND NOT frame_control.MAC_action = request_with_no_response
    AND NOT response_received
    AND NOT ( Rx_protocol_frame OR Rx_data_frame )

         remaining_retries := remaining_retries - 1;
         send( FC => frame_control,
               DA => destination,
               SA => source,
               data_unit => sdu,
               FCS suppression => suppress_FCS );
         response_window_timer.start( 3 );
-------------------------------------------------------------------
6. AWAIT_IFM_RESPONSE  no_response_6        5. USE_TOKEN

    response_window_timer.expired
    AND bus_quiet
    AND remaining_retries = 0
    AND NOT frame_control.MAC_action = request_with_no_response
    AND NOT response_received
    AND NOT ( Rx_protocol_frame OR Rx_data_frame )

         -- Notify IFM of non-response.
         IFM_EVENT notify( response_overdue );
         -- Check for another serviceable request.
-------------------------------------------------------------------
```

```
Current State          Transition Name       Next State
    Exit Condition
         Action Taken
-----------------------------------------------------------------------
6. AWAIT_IFM_RESPONSE  unexpected_frame_6    1. IDLE

    NOT frame_control.MAC_action = request_with_no_response
    AND NOT response_received
    AND ( Rx_protocol_frame
          OR ( Rx_data_frame
               AND Rx_frame.DA /= TS ) )
    AND Rx_frame.SA /= TS
         -- Some other station thinks it has the token, so defer
         -- by notifying IFM of the situation and dropping out.
         IFM_EVENT notify( unexpected frame );
-----------------------------------------------------------------------
6. AWAIT_IFM_RESPONSE  own_frame_6           6. AWAIT_IFM_RESPONSE

    NOT frame_control.MAC_action = request_with_no_response
    AND NOT response_received
    AND (Rx_protocol_frame OR Rx_data_frame)
    AND Rx_frame.SA = TS
         -- Ignore own frame.
         Rx_protocol_frame := false;
```

7.2.3.7 Check Access Class State

```
Current State         Transition Name      Next State
    Exit Condition
         Action Taken
--------------------------------------------------------------------------
7. CHECK ACCESS_CLASS next_access_class    5. USE_TOKEN

    access_class > ring_maintenance
         -- Load token_hold_timer with residual from the
         -- token_rotation_timer for this access_class.
         token_hold_timer.start(
                       token_rotation_timer( access_class ).value);
         -- Restart token rotation timer.
         token_rotation_timer( access_class ).start(
                            target_rotation_time( access_class ) );
--------------------------------------------------------------------------
7. CHECK ACCESS_CLASS leave_ring           8. PASS_TOKEN

    access_class = ring_maintenance
    AND NS_known
    AND NOT in_ring_desired
    AND NOT any_send_pending
    AND inter_solicit_count > 0

         -- All access levels checked and all queues emptied.
         -- Request to be removed from ring
         -- and pass_token to next station.
         send( FC => set_successor,
               DA => PS,
               data_unit => NS );
         pass_state := pass_token;
         inter_solicit_count := inter_solicit_count - 1;
         -- Restart ring maintenance token rotation timer.
         token_rotation_timer( ring_maintenance ).start(
                        target_rotation_time( ring_maintenance ) );
--------------------------------------------------------------------------
7. CHECK ACCESS_CLASS do_solicit_successor 8. PASS_TOKEN

    access_class = ring_maintenance
    AND NS_known
    AND inter_solicit_count = 0
    AND NOT token_rotation_timer( ring_maintenance ).expired

         -- All access_class levels checked and accessed,
         --   when possible.
         -- Offer entry to new stations now.
         pass_state := solicit_successor;
         -- restart ring maintenance token rotation timer
         token_rotation_timer( ring_maintenance ).start(
                        target_rotation_time( ring_maintenance ) );
--------------------------------------------------------------------------
```

```
Current State              Transition Name     Next State
    Exit Condition
         Action Taken

7. CHECK_ACCESS_CLASS      do_pass_token       8. PASS_TOKEN

    access_class = ring_maintenance
    AND NS_known
    AND ( ( ( in_ring desired OR any_send_pending )
            AND inter_solicit_count> 0 )
        OR ( inter_solicit_count = 0
           AND token_rotation_timer( ring_maintenance ).expired ) )

         -- All access_class levels checked and accessed,
         --   when possible.
         -- Pass the token to next_station.
         pass_state := pass_token;
         inter_solicit_count := max( inter_solicit_count - 1 , 0 );
         -- Restart ring maintenance token rotation timer.
         token_rotation_timer( ring_maintenance ).start(
                        target_rotation_time( ring_maintenance ) );

7. CHECK ACCESS_CLASS      do_solicit_any      8. PASS_TOKEN

    access_class = ring_maintenance
    AND NOT NS_known

         -- All access_class levels checked and accessed,
         --   when possible.
         -- Need to find some successor.
         pass_state := solicit_any;
         -- Restart ring maintenance token rotation timer.
         token_rotation_timer( ring_aintenance ).start(
                        target_rotation_time( ring_maintenance ) );
```

7.2.3.8 Pass Token State

```
Current State         Transition Name       Next State
    Exit Condition
         Action Taken
-----------------------------------------------------------------
-- NOTE on State 8. PASS_TOKEN substates and sequencing:
--      State 8 has seven substates, each representing a different
--      aspect of the token passing process. These substates are
--      kept track of by the state variable “pass_state,” which is of
--      type “tok_pass_substate.” The type definition is repeated
--      below. The function “tok_pass_substate’succ( )” is the Ada
--      built-in successor function - it allows sequencing to the
--      next substate without knowing the current substate name.
--
--      type tok_pass_substate is (solicit_successor,
--                                 pass_token,
--                                 repeat_pass_token,
--                                 send_who_follows,
--                                 repeat_who_follows,
--                                 solicit_any,
--                                 total_failure );
-----------------------------------------------------------------
8. PASS_TOKEN         open_1_response_window 10. AWAIT_RESPONSE

    pass_state = solicit_successor
    AND NOT lowest_address
         -- Open 1 response window
         --   for stations with addresses between TS and NS.
         resolution_pass_count := 0;
         heard_successor := false;
         Rx_protocol_frame := false;
         noise_burst := false;
         send( FC => solicit_successor 1,
               DA => NS );
         response_window_timer.start( 1 );
```

```
Current State         Transition Name        Next State
    Exit Condition
         Action Taken
--------------------------------------------------------------------
8. PASS_TOKEN         open_2_response_windows 10. AWAIT_RESPONSE

    pass_state = solicit_successor
    AND lowest_address

         -- Open 2 response windows
         --   for stations with addresses smaller or larger
         --   than any now in ring.
         resolution_pass_count := 0;
         heard_successor := false;
         Rx_protocol_frame := false;
         noise_burst := false;
         send( FC => solicit_successor_2,
               DA => NS );
         response_window_timer.start( 2 );
--------------------------------------------------------------------
8. PASS_TOKEN         pass_token             9. CHECK_TOKEN_PASS

    pass_state = pass_token
    OR pass_state = repeat_pass_token

         -- Pass token.
         send( FC => token,
               DA => NS );
         just_transmitted := true;
         -- just transmitted does not need to be set in this arc.
         -- Analysis has shown that it is not possible to go to the
         --   IDLE state without “flushing” the network of the
         --   station’s frames.  Implementation of this action is
         --   optional and its elimination may enhance duplicate
         --   address detection.

         -- The following 2 variables are used in CHECK_TOKEN_PASS.
         first_time := true;
         echoes_expected := true;

         Rx_protocol_frame := false;
         noise_burst := false;
         token_pass_timer.start( 1 );
--------------------------------------------------------------------
```

```
Current State          Transition Name       Next State
    Exit Condition
         Action Taken
-----------------------------------------------------------------
8. PASS_TOKEN          who_follows_query     10. AWAIT_RESPONSE

    pass_state = send_who_follows
    OR pass_state = repeat_who_follows

         resolution_pass_count := 0;
         heard_successor := false;
         Rx_protocol_frame := false;
         noise_burst := false;
         send( FC => who_follows,
               data_unit = NS );
         response_window_timer.start( 3 ); -- long delay
-----------------------------------------------------------------
8. PASS_TOKEN          solicit_any           10. AWAIT_RESPONSE

    pass_state = solicit_any

         -- Open 2 response windows soliciting all
         --  potential successors.
         resolution_pass_count := 0;
         heard_successor := false;
         Rx_protocol_frame := false;
         noise_burst := false;
         send( FC => solicit_successor_2,
               DA => TS );
         response_window_timer.start( 2 );
-----------------------------------------------------------------
8. PASS_TOKEN          silent_pass_to_self   5. USE_TOKEN

    pass_state = total_failure
    AND any_send_pending
         -- No one else is out there, so pass token to self
         --   until transmit queues are empty.
         -- Set access_class to highest level.
         access_class := max_access_class;
         token_hold_timer.start( hi_pri_token_hold_time );
-----------------------------------------------------------------
```

```
Current State          Transition Name      Next State
    Exit Condition
         Action Taken
-----------------------------------------------------------------
8. PASS_TOKEN          sole_station         1. IDLE

    pass_state = total_failure
    AND NOT any_send_pending
    AND sole_active_station

         -- Reached this state last time.
         -- This is probably the only active station,
         --   so just idle.
         in_ring := false;
         NS_known := false;
         lowest_address := false;
         bus_idle_timer.start( max_bus_idle );
-----------------------------------------------------------------
8. PASS_TOKEN          no_successor_8       1. IDLE

    pass_state = total_failure
    AND NOT any_send_pending
    AND NOT sole_active_station
    AND transmitter_fault_count < max_transmitter_fault_count

         -- No successor found. Possible "deaf receiver" or
         --   "faulty transmitter" condition.
         -- Assume fault present; don't pass token to self.
         in_ring := false;
         NS_known := false;
         lowest_address := false;
         sole_active_station := true;
         transmitter_fault_count := transmitter_fault_count + 1;
         LM_EVENT notify( no_successor );
         bus_idle_timer.start( max_bus_idle );
-----------------------------------------------------------------
8. PASS_TOKEN          no_future            0. OFFLINE

    pass_state = total_failure
    AND NOT any_send_pending
    AND NOT sole_active_station
    AND transmitter_fault_count>= max_transmitter_fault_count

         -- No successor found.
         -- Probable "faulty transmitter" condition.
         -- Assume worst case.
         LM_EVENT notify( faulty_transmitter );
=================================================================
```

7.2.3.9 Check Token Pass State

```
Current State          Transition Name       Next State
    Exit Condition
         Action Taken
-------------------------------------------------------------------------
9. CHECK_TOKEN_PASS    token_pass_failed     8. PASS_TOKEN

_ timer.expired
    AND bus_quiet
    AND NOT ( Rx_protocol_frame OR Rx_data_frame OR noise_burst )
         -- Successor failed to accept and use_token
         --   so proceed to next PASS_TOKEN substate.
         pass_state := tok_pass_substate'succ( pass_state );
-------------------------------------------------------------------------
9. CHECK_TOKEN_PASS    own_frame_9           9. CHECK_TOKEN_PASS

    ( ( Rx_protocol_frame AND Rx_frame.FC /= token )
      OR Rx_data_frame )
    AND Rx_frame.SA = TS

         -- Ignore own transmissions.
         -- Also ignore prior noise.
         first_time := true;
         echoes_expected := true;
         Rx_protocol_frame := false;
-------------------------------------------------------------------------
9. CHECK_TOKEN_PASS    own_token             9. CHECK_TOKEN_PASS

    Rx_protocol_frame
    AND Rx_frame.FC = token
    AND Rx_frame.SA = TS

         -- Ignore prior noise.
         first_time := true;
         echoes_expected := false;
         Rx_protocol_frame := false;
-------------------------------------------------------------------------
9. CHECK_TOKEN_PASS    probably_own_token    9. CHECK_TOKEN_PASS

    noise_burst
    AND echoes_expected

       -- Assume we have heard our own token.
       echoes_expected := false;
       noise_burst := false;
-------------------------------------------------------------------------
```

```
Current State              Transition Name     Next State
    Exit Condition
         Action Taken
---------------------------------------------------------------------
9. CHECK_TOKEN_PASS        pass_ok             1. IDLE

    ( Rx_protocol_frame OR Rx_data_frame )
    AND Rx_frame.SA /= TS

         -- Some station is using the token.
         -- Reprocess the frame in state 1 (IDLE).
         -- NO ACTION.
---------------------------------------------------------------------
9. CHECK_TOKEN_PASS        not_sure            9. CHECK_TOKEN_PASS

    noise_burst
    AND NOT echoes_expected
    AND first_time

         -- Something heard - possibly a frame
         -- sent by the token recipient.
         -- Watch for another frame.
         first_time := false;    --Allow only one pass on this arc.
         noise_burst := false;
         token_pass_timer.start( 4 );
---------------------------------------------------------------------
9. CHECK_TOKEN_PASS        probably_ok         1. IDLE

    token_pass_timer.expired
    AND bus_quiet
    AND NOT ( Rx_protocol_frame OR Rx_data_frame )
    AND noise_burst
    AND NOT first_time

         -- Heard a second noise burst - cannot be my token
         --   so successor probably transmitting.
         -- Reprocess event in state 1.
         -- NO ACTION.
=====================================================================
```

7.2.3.10 Await Response State

```
Current State               Transition Name        Next State
    Exit Condition
         Action Taken
-----------------------------------------------------------------
10. AWAIT_RESPONSE          no_response_10         8. PASS_TOKEN

     response_window_timer.expired
     AND bus_quiet
     AND NOT ( Rx_protocol_frame OR Rx_data_frame OR noise_burst )
     AND resolution_pass_count = 0
     AND NOT heard_successor

          -- Response windows expired, unused
          --   ( bus was quiet for entire time ).
          inter_solicit_count := max_inter_solicit_count;
          pass_state := tok_pass_substate'succ( pass_state );
-----------------------------------------------------------------
10. AWAIT_RESPONSE          resolution_succeeded   8. PASS_TOKEN

     response_window_timer.expired
     AND bus_quiet
     AND NOT ( Rx_protocol_frame OR Rx_data_frame )
     AND heard_successor

          -- new successor found, so resolution process is complete.
          -- PASS TOKEN to new successor
          --   and open response window again next pass.
          pass_state := pass_token;
-----------------------------------------------------------------
10. AWAIT_RESPONSE          own_address_10         10. AWAIT_RESPONSE

     ( Rx_protocol_frame OR Rx_data_frame )
     AND Rx_frame.SA = TS

          -- Ignore own transmissions.
          -- Also ignore prior noise bursts.
          noise_burst := false;
          Rx_protocol_frame := false;
```

```
Current State          Transition Name       Next State
    Exit Condition
         Action Taken
-------------------------------------------------------------------
10. AWAIT_RESPONSE     hear_successor        10. AWAIT_RESPONSE

     Rx_protocol_frame
     AND Rx_frame.FC = set_successor
     AND Rx_frame.SA /= TS
     AND Rx_frame.DA = TS
     AND Rx_frame.data_unit /= TS
     AND NOT heard_successor

          -- Some successor heard "in the clear."
          NS := Rx_frame.data_unit;
          lowest_address := ( NS> TS )
           NS_known := true;
          sole_active_station := false;
          heard_successor := true;
          inter_solicit_count := 0;
          Rx_protocol_frame := false;
          transmitter_fault_count := 0;
          LM_EVENT notify( new_successor, NS )
          -- NOTE: The way this specification is written, the first
          --       valid set_successor message heard will be used to
          --       determine the successor's address.
          --       A conformant station is not required to use the
          --       first set_successor message received. Any such
          --       message, correctly received, may be used.
-------------------------------------------------------------------
10. AWAIT_RESPONSE     ignore_response       10. AWAIT_RESPONSE

     Rx_protocol_frame
     AND Rx_frame.FC = set_successor
     AND Rx_frame.SA /= TS
     AND Rx_frame.DA = TS
     AND ( Rx_frame.data_unit = TS OR heard_successor )

        Rx_protocol_frame := false;
-------------------------------------------------------------------
```

```
Current State          Transition Name      Next State
    Exit Condition
         Action Taken
---------------------------------------------------------------------
10. AWAIT_RESPONSE     unexpected_frame_10  1. IDLE

     Rx_frame.SA /= TS
     AND ( Rx_data_frame

         OR ( Rx_protocol_frame
                AND ( Rx_frame.FC /= set_successor
                      OR Rx_frame.DA /= TS ) ) )

          -- Some other station thinks it has the token, so ignore
          --   any previous noise and defer by dropping out.
          -- Reprocess the frame in state 1 (IDLE).
          noise_burst := false;
---------------------------------------------------------------------
10. AWAIT_RESPONSE     send_resolve         10. AWAIT_RESPONSE

     response_window_timer.expired
     AND bus_quiet
     AND NOT ( Rx_protocol_frame OR Rx_data_frame )
     AND noise_burst
     AND NOT heard_successor
     AND resolution_pass_count < max_pass_count

         -- Infer a collision from presence of signal without a
         --   valid frame.
         -- Another 4 response_window resolution pass is called for.
         noise_burst := false;
         resolution_pass_count := resolution_pass_count +1;
         send( FC => resolve_contention );
         response_window_timer.start( 4 );
---------------------------------------------------------------------
10. AWAIT_RESPONSE     resolution_failed    8. PASS_TOKEN

     response_window_timer.expired
     AND bus_quiet
     AND NOT ( Rx_protocol_frame OR Rx_data_frame )
     AND NOT heard_successor
     AND ( ( NOT noise_burst AND resolution_pass_count> 0)
         OR resolution_pass_count>= max_pass_count )

          -- No resolution reached, and no more demands heard or
          --   possible, so pass token to known successor.
          inter_solicit_count := 0;
                               -- Solicit successors on next pass.
          pass_state := tok_pass_substate'succ( pass_state );
```

8. MAC Sublayer-Physical Layer Interface Service Specification

The services provided to the MAC sublayer by the Physical Layer of this standard are specified in this section. This section specifies these services in an abstract way. It does not specify or constrain the implementation entities and interfaces within a computer system. The relationship of this section to other sections of this standard and to Local Area Network (LAN) specifications is illustrated in Fig 8-1.

8.1 Overview of the LAN Physical Layer Service

8.1.1 General Description of Services Provided. The services provided by the Physical Layer are informally described here. These services provide for the trans mission and reception of MAC-symbols, each with a duration of one MAC-symbol_time. Jointly, they provide the means by which cooperating MAC entities can coordinate their transmissions and exchange information by way of a shared communications medium.

8.1.2 Model Used for the Service Specification. The model and descriptive method are detailed in Appendix A.

Fig 8-1
Relation to LAN Model

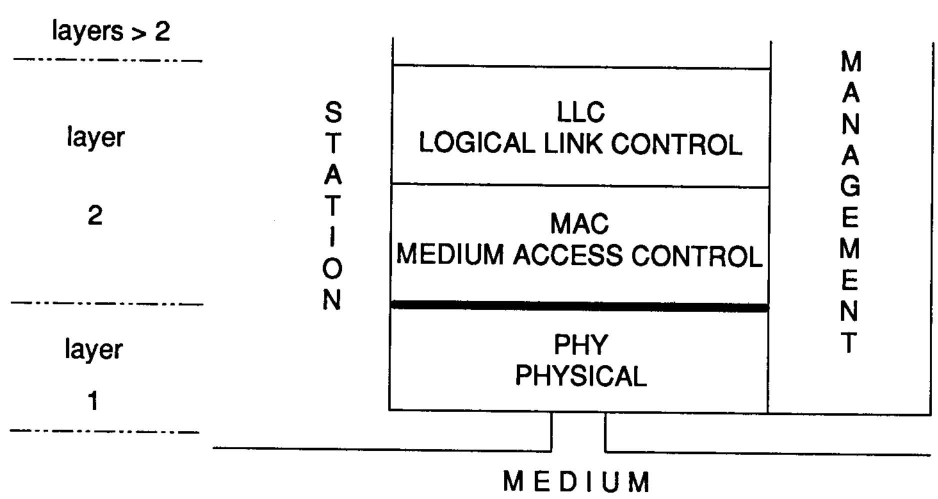

8.1.3 Overview of Interactions. The primitives associated with symbol trans mission and reception are

PHY-UNITDATA request
PHY-UNITDATA indication

PHY-MODE invoke
PHY-NOTIFY invoke

8.1.3.1 PHY-UNITDATA request and indication. The PHY-UNITDATA request primitive is passed to the Physical Layer to request that a symbol be impressed on the LAN's communications medium. Only one such request is accepted per MAC-symbol_time. The PHY-UNITDATA indication primitive is passed from the Physical Layer to indicate the reception of a MAC-symbol from the medium.

8.1.3.2 PHY_MODE invoke. The PHY-MODE invoke primitive is passed to the Physical Layer to establish the transmission mode for subsequent PHY-UNITDATA request primitives, either transmission to all connected media segments or, when the station is functioning as a bus repeater, transmission to all connected media segments other than the one which is the current source of symbols for the PHY-UNITDATA indication primitive. This primitive also affects the source of the station's MAC-symbol timing.

8.1.3.3 PHY-NOTIFY invoke. The PHY-NOTIFY invoke primitive is passed to the Physical Layer to notify the Physical Layer that an end-of-frame delimiter has just been detected, and that the following received symbols should be either *silence* or the results of a properly transmitted *pad_idle* sequence.

8.1.4 Basic Services and Options. All PHY-UNITDATA services are required in all implementations, and both of the PHY-UNITDATA primitives are mandatory. The PHY-MODE primitive is mandatory only in stations that can function both as originating and as repeating stations in the network. The PHY-NOTIFY primitive is mandatory only in stations that can function as repeating stations in the network.

8.2 Detailed Specifications. This subsection describes in detail the primitives and parameters associated with the Physical Layer services. The parameters are specified in an abstract sense and specify the information that shall be available to the receiving PLE. A specific implementation is not constrained in the method of making this information available.

8.2.1 PHY-UNITDATA request

8.2.1.1 Function. This primitive is the service request primitive for the symbol transfer service.

8.2.1.2 Semantics. This primitive shall provide parameters as follows:

PHY-UNITDATA request (
symbol
)

symbol shall specify one of the following:

(1) *Zero*—corresponds to a binary 0.
(2) *One*—corresponds to a binary 1.
Data is the collective name for *zero* and *one*.
(3) *Non_data*—used in delimiters, always sent in pairs, and always in octets with the form.

non_data non_data data non_data non_data data data data

(4) *Pad_idle*—send one symbol of preamble (*preamble* is a physical-entity-specific sequence of *data* symbols).
(5) *Silence*—send silence (or pseudo-silence) for a duration of one MAC-symbol_time.

8.2.1.3 Generated. This primitive is passed from the MAC sublayer to the Physical Layer to request that the specified symbol be transmitted on the LAN medium.

When so required by the appropriate Physical Layer standard this primitive shall be passed to the Physical Layer once for each PHY-UNITDATA indication that the MAC sublayer receives from the Physical Layer. In such cases, there shall also be an implementation-dependent constant phase relationship, determined by the MAC, between a PHY-UNITDATA indication and the next subsequent PHY-UNITDATA request.

8.2.1.4 Effect of Receipt. Receipt of this primitive causes the Physical Layer to attempt to encode and transmit the symbol using signaling appropriate to the LAN medium.

8.2.1.5 Constraints

Pad_idle symbols, which are referred to collectively as *preamble*, are transmit ted at the start of each MAC frame, both to provide a training signal for receivers and to provide a nonzero minimum separation between consecutive frames. The following constraints apply:

(1) An originating station shall transmit a minimum number of octet multiples of *pad_idle* so that their duration is at least 2 μs, and after completing transmission of the last required octet, it may (but need not) transmit more octets of *pad_idle* symbols before the first frame-delimiter.
(2) A repeating station on a single-channel bus (for example, an FSK repeater) shall retransmit at least the same minimum number of octets of *pad_idle*, and may transmit more *pad_idle* symbols, before repeating the first frame-delimiter and the remainder of the transmission.
(3) A repeating station on a dual-channel (for example, broadband) system, which transmits continuous signaling in the forward channel as a training and synchronization signal for all of the other stations in the LAN, shall retransmit at least one-half octet of *pad_idle* before repeating the first frame-delimiter and the remainder of a transmission. It need not retransmit the entire *pad_idle* signal as originated, since the primary purpose of the initial *pad_idle* sequence, that of a synchronizing sequence for stations in a switched-carrier environment, does not apply to the above-described for ward channel.

Non_data symbols shall be used only within frame delimiters, where they always shall be requested in pairs. The symbol sequences of those frame delimiters shall be

non_data non_data data non_data non_data data data data

where each *data* symbol is either the symbol *zero* or the symbol *one*.

When *data* symbols are transmitted between frame delimiters, the number of *data* symbols transmitted, not including those data symbols within the eight-symbol frame delimiter sequences, shall always be a multiple of eight. (That is, only complete octets of *data* symbols shall be transmitted between frame delimiters.) When *pad_idle* symbols are transmitted between frame delimiters, the number of *pad_idle* symbols transmitted shall always be a multiple of eight. Octets of *pad_idle* symbols and octets of *data* symbols shall always be separated by frame delimiter octets, or a sequence of *silence* symbols, or both. (That is, *pad_idle* octets and *data* octets cannot be intermixed.)

When a Physical Layer standard constrains the timing relationship between a PHY-UNITDATA request and the subsequent PHY-UNITDATA indication, as permitted by 8.2.1.3, then the jitter in the implementation-dependent constant phase relationship between consecutive PHY-UNITDATA indication and PHY-UNITDATA request primitives shall not be greater than 2%.

8.2.1.6 Additional Comments. The confirmation of this request is a timed confirmation, which can only be made once per transmit MAC-symbol_time. Consequently, this request shall only be sent once per transmit MAC-symbol_time.

8.2.2 PHY-UNITDATA indication

8.2.2.1 Function. This primitive is the service indication primitive for the symbol transfer service.

8.2.2.2 Semantics. This primitive shall provide parameters as follows:

PHY-UNITDATA indication (
symbol
)

symbol shall specify one of the following:

(1) *Zero*—corresponds to a binary 0.
(2) *One*—corresponds to a binary 1.
(3) *Non_data*—used in delimiters, always sent in pairs.
(4) *Pad_idle*—corresponds to one MAC-symbol_time during which preamble was received and not reported as either *zero* or *one*.
(5) *Silence*—corresponds to one MAC-symbol_time of received silence (or pseudo-silence).
(6) *Bad_signal*—corresponds to one MAC-symbol_time during which inappropriate signaling was received.

8.2.2.3 When Generated. This primitive is passed from the Physical Layer to the MAC sublayer to indicate that the specified symbol was received from the LAN medium.

8.2.2.4 Effect of Receipt. The effect of receipt of this primitive by the MAC sublayer entity is unspecified.

8.2.2.5 Additional Comments. This indication is a timed indication, which can only be made once per received MAC-symbol_time. Consequently, this indication shall only be sent once per received MAC-symbol_time.

Each transmission begins with *pad_idle* symbols, and it is expected that some, but not all, of these initial symbols may be "lost in transit" between the transmitting station and the receiving stations, and consequently reported as *silence* or *bad-signal*.

Where the Physical Layer encoding for successive *pad_idle*s is a modulation-specific sequence of the encodings for *zero* and *one*, receivers are permitted to decode such a transmitted sequence of *pad_idle*s as a sequence of *zero*s and *ones* and report them as such to the MAC entity. In other words, a receiver need not have the ability to detect and report *pad_idle* as such; rather it may report the corresponding signaling as *data*.

In the absence of errors or colliding transmissions, and with the above two exceptions for symbols transmitted as *pad_idle*, the sequence of symbols reported is identical to the sequence of symbols transmitted by the associated PHY-UNITDATA requests.

For some PLEs, it may take a few MAC-symbol_times to detect silence on the medium after the end of transmission. As a consequence, up to four *silence* symbols transmitted immediately after any symbol other than *silence* may be reported by the receiving Physical Layer as any MAC-symbol.

The PHY-UNITDATA indication shall be correctly generated independent of the current PHY-UNITDATA request. That is, the indication shall be valid even when the station is transmitting. This precludes the use of half-duplex Physical Layers.

8.2.3 PHY-MODE invoke

8.2.3.1 Function. This primitive is the service request primitive for the symbol transfer mode-setting service.

8.2.3.2 Semantics. This primitive shall provide parameters as follows:

PHY-MODE invoke (
mode
)

mode shall specify one of the following:

(1) *Originating*—MAC entity is originating symbols to be transmitted on all segments of the connected LAN medium.

(2) *Repeating*—MAC entity is functioning as a repeater, interpreting and repeating the symbols that the PHY entity is receiving from one segment to all the other segments of the connected LAN medium.

8.2.3.3 When Generated. This primitive is passed from the MAC sublayer to the Physical Layer to establish the specified operational mode of the PHY entity.

8.2.3.4 Effect of Receipt. Receipt of this primitive causes the Physical Layer to select the appropriate mode of operation for succeeding PHY-UNIT-

DATA request primitives. The *originating* mode causes the PHY entity to transmit symbols at a locally determined rate onto all segments of the medium connected to the station. The *repeating* mode causes the PHY entity to monitor all attached segments, selecting a segment on which symbols other than *silence* are being signaled (where possible) as the source of the received symbols reported by the PHY-UNITDATA indication primitive and as the source of both the receive and transmit symbol timing. The *repeating* mode also inhibits transmission onto the selected medium segment from which data and timing are being received.

8.2.3.5 Additional Comments. This mode selection shall be performed dynamically within a bus repeater that can also operate as an originating station; it may be performed statically (for example, at design time) in all other cases. Additional information on the use of this primitive can be found in the various PLE specifications (see 12.7, 14.8, 16.7, 18.7).

8.2.4 PHY-NOTIFY invoke

8.2.4.1 Function. This primitive is the service request primitive for the potential-end-of-reception notification service.

8.2.4.2 Semantics. This primitive shall be parameterless, as follows:

PHY-NOTIFY invoke

8.2.4.3 When Generated. This primitive is passed from the MAC sublayer to the Physical Layer to notify the Physical Layer that an end delimiter has just been detected, and that the following received symbols should be either *silence* or the results of a properly transmitted *pad_idle* sequence.

8.2.4.4 Effect of Receipt. The effect of receipt of this primitive by the PLE is determined by the implementation. (It may cause the PLE to make special provisions for detecting silence [for example, by checking for the known *pad_idle* sequence] or for switching its automatic gain control (AGC) and clock-recovery circuitry to a high-speed acquisition mode.)

9. Physical Layer Entity (PLE) Management

This section specifies the abstract entities (parameters, events, and actions) used in managing the PLE. The relationship of this section to other sections of this standard and to LAN specifications is illustrated in Fig 9-1.

There are two ways to manage a PLE: locally or remotely.

Local management is required to initially configure the PLE and to operate it when remote management is not available. Since the operation of local PLE management, other than the need to manage the entities described in this section, does not affect station interoperability, the specification of local management is outside the scope of this standard.

This standard uses the concept of remote management as defined in the Open Systems Interconnection (OSI) Management Framework described in ISO/IEC 7498-4 : 1989 [14]. Remote management uses a communications protocol to manipulate and observe the managed entities. This section is designed to be used with the companion remote management standard described by IEEE P802.1B [10]. Other protocols, such as described by ISO/IEC 9595 [18] and ISO/IEC 9596 [19], may also be used.

Fig 9-1
Relation to LAN Model

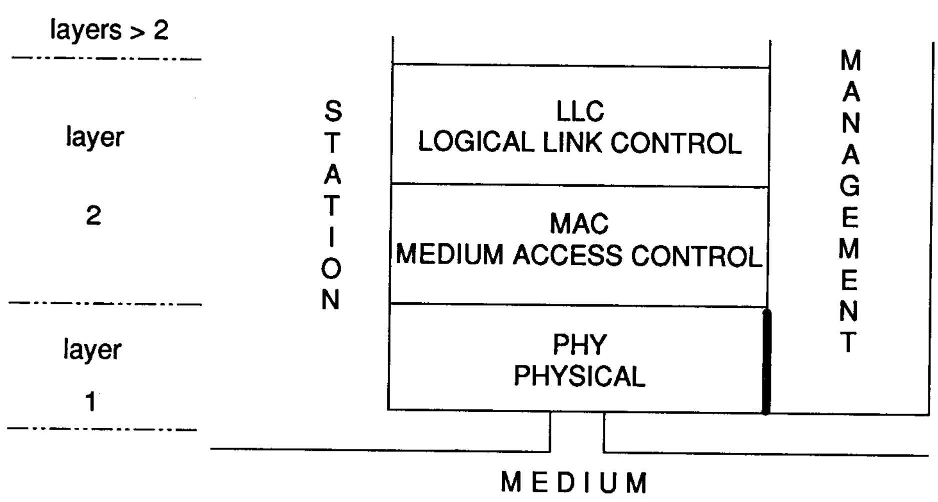

To implement a remote management protocol, the identification and value representation of the managed entities must be defined. Such a definition is given by a layer management interface (LMI) protocol.

NOTE: The definition of a Physical Layer LMI protocol is under study. An addendum to this standard is planned that will contain this protocol.

9.1 Overview. This section describes the entities communicable via the LMI, shown as a thick line in Fig 9-1. The PLEs described in this section consist of

(1) Parameters within the PLE written and read by the management entity,
(2) Events within the PLE that are reported to the management entity,
(3) Actions initiated by the management entity that cause changes within the PLE.

9.2 Physical Management Facilities. This subsection contains a description of use of the management parameters, actions, and events within the token-passing bus PLE.

9.2.1 Organization. There are three types of management entities:

(1) Parameters
(2) Actions
(3) Events

The PLE parameters accessible by the management entity are described first. These parameters are organized into groups. The groups not only collect the variables; they also define collections of parameters that can be accessed in a common manner.

NOTE: In implementing control of access to parameters, the intent of this standard is to implement the granularity of access control only to the group level. Thus, all parameters within a group share a common access level.

Any or all of the parameters in a group may be passed in a single remote management request. The semantics and grouping of the parameters so communicated are specified in the following subsection.

9.2.2 ResourceTypeID. This subsection discusses the parameters of the Physical LMI that are common to any LMI. These parameters create a "name-plate" which identifies the basic attributes of the PLE.

9.2.2.1 resourceType. This parameter identifies the resource as a PLE.

9.2.2.2 standardRevision. This parameter identifies the revision of this standard implemented by the PLE.

9.2.2.3 lmeOptions. Implementations that conform to this standard shall return a value of *none*.

9.2.3 Capabilities Group. These parameters describe the operational capabilities of the PLE. All parameters are optional. All of the parameters shall be read-only.

9.2.3.1 sectionConformance. This parameter identifies the section of this standard to which the PLE conforms.

9.2.3.2 operationalModes. This parameter specifies the modes of operation (originating, repeating, or both) that the PLE supports.

9.2.3.3 dataRates. This parameter specifies the nominal data rate (or rates in order of preference) at which the PLE can transmit and receive.

9.2.3.4 mediaSegments. This parameter enumerates the media line segments to which the PLE is connected, and specifies the PLE's capabilities (transmit, receive, or both) with respect to each segment. In enumerating the media segments first the *trunkId* and then the *copyId* segment identifiers are given.

TrunkId identifiers are used for PLEs which are multi-port repeaters.

CopyId identifiers are used for one of the redundant copies of a trunk; it is only needed in systems employing redundant media.

9.2.3.5 transmitPowerAdjustable. This parameter specifies whether the transmitter's output level is adjustable via the LMI.

9.2.3.6 transmitChannels. This parameter specifies the sets of nominal carrier frequencies (and thus channels) on which the PLE can transmit. If these frequencies are unknown, the PLE may return a null set.

9.2.3.7 receiveChannels. This parameter specifies the sets of nominal carrier frequencies (and thus channels) on which the PLE can receive. If these frequencies are unknown, the PLE may return a null set.

9.2.3.8 receivedSignalLevelMeasurable. This parameter specifies whether the PLE can provide the measure of received signal level required for the off-normal-power-level peer report event.

9.2.4 Operational State Group. These parameters describe the current state of the PLE. The values of these parameters are constrained to be logically consistent with the values of the parameter values specified in the capabilities group. All parameters are optional. Access by the management entity to these parameters shall be read-only.

9.2.4.1 operationalMode. This parameter specifies the mode of operation of the PLE.

9.2.4.2 dataRate. This parameter specifies the nominal data rate at which the PLE is operating.

9.2.4.3 receiveSegments. This parameter specifies the media line segment(s) which the PLE is monitoring for received signaling.

9.2.4.4 transmitSegments. This parameter specifies the media line segment(s) which the PLE is enabled to transmit (whether or not the PLE is "actively" transmitting at the moment).

9.2.4.5 receiveChannel. This parameter specifies the PLE's receive carrier frequency.

9.2.4.6 transmitChannel. This parameter specifies the PLE's transmit carrier frequency.

9.2.4.7 reportReceivedSignalLevel. This parameter specifies whether the PLE is enabled to signal an off-normal-power-level peer report event when receiving a signal meeting the event's trigger criteria.

9.2.5 Initialization State Group. The values of these parameters are copied to the corresponding operational state group (see 9.2.4) parameters when the PLE is initialized. The values of these parameters are constrained to be logically consistent with the parameter values specified in the capabilities group. All parameters are optional. Access by the management entity to these parameters shall be read-write.

9.2.5.1 initOperationalMode

9.2.5.2 initDataRate

9.2.5.3 initReceiveSegments

9.2.5.4 initTransmitSegments

9.2.5.5 initReceiveChannel

9.2.5.6 initTransmitChannel

9.2.5.7 initReportReceivedSignalLevel

9.2.6 Counters Group. These are 32-bit thresholded event counters associated with the PLE. The implementation of the counters and access by the management entity shall be optional.

Implementation of a threshold facility associated with each counter shall be treated as a separate option. Thus three conditions exist: No implementation, 32-bit counters without threshold, or 32-bit counters with threshold. If implemented, access by the management entity to these parameters shall be read-only, except for the portions associated with the threshold mechanism.

9.2.6.1 uncorrectedErrorEvents. The PLE increments this counter each time it receives a signaling sequence which is reported as one or more *bad_signal* symbols. The exact correspondence between the number of *bad_signal* symbols reported and the amount by which the counter is incremented is implementation-specific. (For example, the counter may be incremented once for each *bad_signal* symbol, or it may be incremented once for each string of such symbols.)

9.2.6.2 correctedErrorEvents. The PLE increments this counter each time it receives a signaling sequence that is detectably erroneous, but which the PLE can and does correct (for example, by using the forward-error-correcting potential of the modulation of Section 14). The exact correspondence between the number of symbols corrected and the amount by which the counter is incremented is implementation-specific.

NOTE: A close approximation of the total number of symbols received during a particular interval can be computed from the PLE's nominal data rate. (See 9.2.5.2.)

9.2.7 Action Definitions. Actions are requests by the management entity for the layer to perform some activity. All actions, unless otherwise defined in the specific layer management sections, are optional.

9.2.7.1 adjustTransmitPower. This action permits the management entity to request that the PLE adjust its transmitter output power level on a specified medium line segment. The action provides a segment identifier and a relative power level. The relative power level enumerates the qualitative power level of the transmitted signal relative to the nominal signal level specified in the appropriate PLE section of this standard (see 14.7).

The PLE shall derive a direction (increase or decrease) and a magnitude (small or large) for the desired adjustment from the relativePowerLevel parameter. The PLE shall respond to each such action request with an action report that indicates the success or reason for failure of the request.

After receiving an adjust transmit power action request, the PLE should reject all similar requests, by responding if necessary with a *duplicate adjustment rejected* status, for a period of time on the order of two seconds. After this request aging period, the PLE should again accept similar action requests.

9.2.7.2 selfTest. This action permits the management entity to request that the PLE perform a self-test and report the results of that testing. The action provides a value giving the duration of the test. The PLE shall respond to each such action request by initiating a self-test for up to the specified period of time and returning an action report to the requestor indicating the success, failure, or incomplete (due to lack of time) status of the self-test.

The action also specifies the loopback point for the test as one of the following:

(1) **loop3A.** Loopback shall occur at the PLE's MAC connector.
(2) **loop3B.** Loopback shall occur internally within the PLE's digital logic.
(3) **loop3C.** Loopback shall occur internally within the PLE's analog logic.
(4) **loop3D.** Loopback shall occur as close as possible to the PLE's connection to the medium.

See CCITT Recommendation X.150 [2] for further discussion of the meaning of the terminology used to define the loopback points in this action.

9.2.8 Event Definitions. Events are spontaneously reported by the PLE to the management entity. Reporting of events shall be optional.

NOTE: In implementing routing of events, the intent of this standard is to provide a unique routing path for each event.

9.2.8.1 thresholdReached. A thresholded counter has passed the threshold value. The event value shall be the counter's parameter identifier.

9.2.8.2 abnormalEventReport. This event, together with an appropriate identifying code, is indicated via the LMI upon the occurrence of one of the following abnormal conditions within the PLE:

(1) **unspecified.** A fault has been detected within the PLE and either the type of fault is not distinguishable. Any fault may be reported as *unspecified* in lieu of a more-specific fault code.
(2) **pleJabber.** The mandatory jabber timer within the PLE has detected that a PLE transmitter has been actively transmitting for more than the maximum permitted time period.
(3) **otherPLEFault.** A fault other than jabber has been detected in the operation of the PLE.
(4) **invalidSymbolStreamFromMAC.** The symbol stream presented to the PLE by the MAC is erroneous.
(5) **dte-dceInterfaceFault.** The optional DTE-DCE interface within the PLE is malfunctioning (see Section 10).
(6) **mediumFault.** The PLE has detected a true fault on the medium (not to be confused with lossOfReceivedSignal).
(7) **lossOfReceivedSignal.** A PLE that normally receives a continuous-carrier signal has detected loss of that signal (see Section 14). Since this is an expected but abnormal occurrence and is not inherently indicative of a fault, generation of this event is optional.

NOTE: In such situations, the PLE is still required to report any received signaling above the minimum receiver off level to the MAC as *bad_signal* symbols.

9.2.8.3 peerReport. This event, together with the appropriate cause, is indicated via the LMI when reporting of off-normal received signal levels has been enabled in the PLE (see 9.2.3.8). A relative power level parameter specifies the power level of a received signal relative to the nominal signal level specified in the appropriate PLE section of this standard. (Typical relative power level values are tooLow, low, ok, high, or tooHigh.) Local management should correlate this PLE report with the MAC source address (see 4.1.4.2) field of a received frame, or the received segment identifier. As a practical matter, this event should be reported at the LMI in a timely manner.

The event service request, used to convey this information by management, specifies both the specific cause of the event as reported by the PLE, and the MAC address of the peer station whose received signal was so assessed as reported by the MAC sublayer entity.

9.3 Additional Management. Additional management parameters are needed for proper configuration of the local MAC sublayer entity. The parameter minimumPostSilencePreambleLength specifies the modulation and data-rate-dependent number of octets of preamble that the MAC shall send when commencing a trans mission, as specified in 4.1.1, 12.6, 14.7, 16.6, and 18.6. These parameters are communicated to the local MAC sublayer entity for correct MAC operation. Support of minPostSilencePreambleLength is required.

10. Exposed DTE-DCE Interface

This section defines the functional, electrical, and mechanical characteristics of the interface between data terminal equipment (DTE) and data computer equipment (DCE) of token bus Local Area Networks (LANs). The purpose of this interface is to provide a standard means of connecting DCEs to DTEs in a token bus station. The functional part of the DTE-DCE interface is mandatory for all applications. Specific functions, however, are specified as an option. Both the electrical and mechanical parts are optional.

The relationship of the DTE-DCE interface to other sections of this standard is illustrated in Fig 10-1.

Fig 10-1
Relation to LAN Model

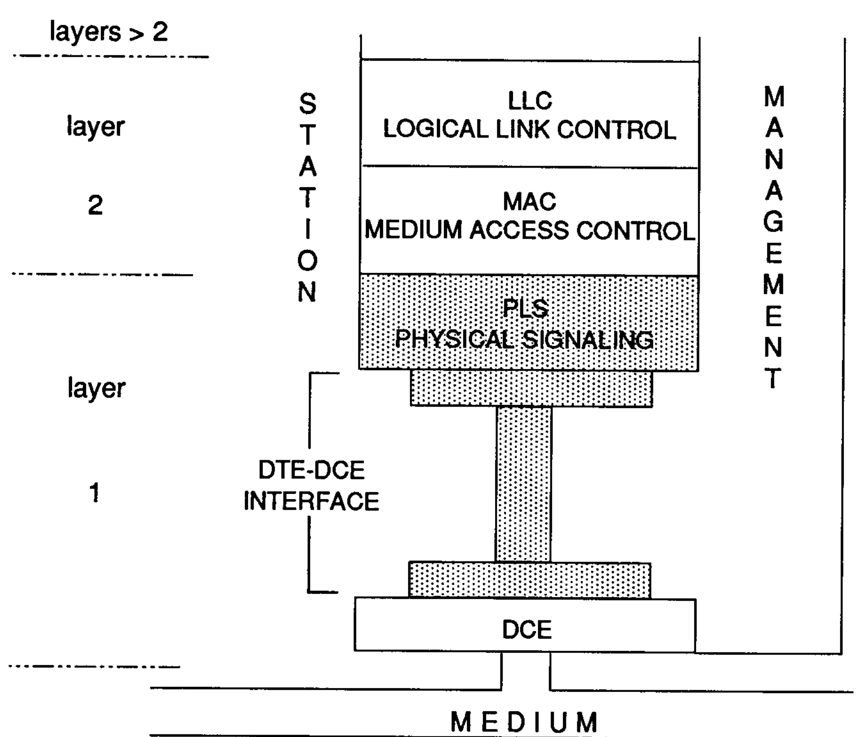

10.1 Overview of DTE-DCE Interface. The DTE-DCE interface provides the means for transferring requests for data unit transmission and for indicating data unit reception. The interface also serves as a communication channel for passing management information between the sublayer entities to which it is connected. The protocol defined in this section supports the communications required to implement the services defined in Section 8 (MAC sublayer-Physical Layer interface) and Section 9 (PLE management). The PLE signaling sublayer (PLS) shown in Fig 10-1 provides translation between the service primitives defined in Section 8 and the channel protocol defined in this section. The PLS is not formally defined in this section; however, the services required of the PLS are effectively defined by describing the interactions at its interfaces (MAC sublayer-Physical Layer interface and DTE-DCE interface). These interactions are defined in this section and in Section 8.

10.1.1 PHY-UNITDATA Request Channel. The PHY-UNITDATA request channel provides the communication facility required to support the PHY-UNITDATA request primitive as defined in Section 8. The request channel also serves as a communication channel for passing management requests from the DTE to the DCE. Timing circuits used to synchronize the exchange of messages across the interface are also included as part of the request channel. Timing information is always derived from the DCE.

10.1.2 PHY-UNITDATA Indication Channel. The PHY-UNITDATA indication channel provides the communication facility required to support the PHY-UNITDATA indication primitive defined in Section 8. To support this primitive the indication channel transfers indications of MAC-symbol reception to the DTE (PHY-UNITDATA indication). The indication channel also serves as a communication channel for passing management indications from the DCE to the DTE. Timing circuits used to synchronize the exchange of messages across the interface are also included as part of the indication channel. The DCE is always the source of timing information.

10.1.3 Physical Layer Management. The DTE-DCE interface supports two modes of operation. In the MAC mode, the request and indication channels transfer data. In management mode, management information is passed between the DCE and upper layer entities using the request and indication channels.

10.2 PHY-UNITDATA Request and Indication. The following is a detailed description of the request and indication channels. Each channel consists of a set of circuits that carry encoded information in the form of MAC symbols. The channel encoding is arranged to allow the most common subscriber nodes to be implemented with a minimum of hardware.

10.2.1 Nomenclature. Request and indication symbols and associated control and timing elements uniquely used within this section are defined below.

10.2.1.1 Request Channel

(1) Symbols

silence. Transmit silence (or pseudo-silence) for one MAC-symbol_time.

non_data. Transmit non-data for one MAC-symbol_time.

pad_idle. Transmit a binary one for one MAC-symbol_time.

one. Transmit a binary one for one MAC-symbol_time.

zero. Transmit a binary zero for one MAC-symbol_time.

(2) Timing Elements

transmit clock. Used to synchronize requests. Transmit clock is supplied by the DCE. The frequency of the transmit clock is equal to the data rate.

return clock. Derived from transmit clock by the DTE to synchronize the transmission of requests to the DCE.

10.2.1.2 Indication Channel

(1) Symbols

silence. Received silence (or pseudo-silence) for one MAC-symbol_time.

non_data. Received non-data for one MAC-symbol_time.

bad_signal. Received bad signal for one MAC-symbol_time.

one. Received a binary one for one MAC-symbol_time.

zero. Received a binary zero for one MAC-symbol_time.

(2) Timing Elements

receive clock. Used to synchronize indications. Receive clock is supplied by the DCE.

10.2.2 Protocol

10.2.2.1 Request Channel Protocol. In MAC mode, this channel is used by the PLS to provide requests (as MAC symbols) for data transmission or other action as specified here. The request channel consists of four symbol circuits and two timing circuits (one of which is optional). Transmission of symbols on these circuits is synchronous with transmit clock (TXCLK or RETURNCLK).

(1) **Request Channel Circuit Definition.** The request channel circuits carry signals from the DTE to DCE unless otherwise specified. Two circuit types (single-ended and differential) are specified by this standard; their electrical characteristics are specified in 10.4.

The request channel consists of the following circuits:

TXCLK	Transmit clock supplied by the DCE. (The frequency is equal to the data rate.)
RETURNCLK	Derived from TXCLK. (Optional—for use with differential signaling only).
TXSYM0 TXSYM1 TXSYM2	Provide encoded transmit symbols or management requests.
TXSYM3	Selects either MAC mode or management mode for the request channel.

(2) **Request Channel Mode Selection.** The request channel provides two modes of operation, MAC mode for transmitting data, and management mode for sending management requests. The TXSYM3 circuit is used to select the operating mode.

(a) **Request Channel MAC Mode.** The MAC mode is selected when TXSYM3 = "1" and is considered the "normal" (i.e., transmit) mode. The DTE sends physical data units encoded as MAC symbols (*silence*, *non_data*, *pad_idle*, *one*, and *zero*) on circuits TXSYM0, TXSYM1, and TXSYM2 to the DCE as requests for data unit transmission. As previously stated, these symbols are synchronized to TXCLK (or return clock for differential signaling).

(b) **Request Channel Management Mode.** The management mode is selected when TXSYM3 = "0". The request channel is used by PLS management entities to send management commands and data to the DCE. Physical Layer management is discussed in 10.3.

When management mode is selected the DTE sends physical data units corresponding to MAC *silence* symbols.

10.2.2.2 Request Channel Encoding. The encoding of the TXSYM circuits in MAC mode is summarized in Table 10-1.

Table 10-1
Request Channel Encoding—MAC Mode

	TXSYM3	TXSYM2	TXSYM1	TXSYM0
silence	1	1	1	X
non_data	1	1	0	X
pad_idle	1	0	1	X
one	1	0	0	1
zero	1	0	0	0

NOTE: X—Indicates don't care.

The electrical representation of a "1" and a "0" is specified in 10.4. Encoding of TXSYM circuits when in management mode is described in 10.3.2.4.

10.2.2.3 Indication Channel Protocol. When in MAC mode, this channel is used by the DCE to provide encoded indications of data unit reception to the DTE. The indication channel consists of five circuits. Transmission of symbols on these circuits is synchronous with receive clock, RXCLK. All circuits originate in the DCE.

(1) **Indication Channel Circuit Definition**. The indication channel circuits carry signals from the DCE to DTE. Two circuit types (single-ended and differential) are specified by this standard, their electrical characteristics are specified in 10.4. The indication channel consists of the following circuits:

RXCLK	Receive clock supplied by the DCE. (The frequency is equal to the data rate.)
RXSYM0 RXSYM1 RXSYM2	Provide encoded receive symbols or management indication.
RXSYM3	Indicates either MAC (receive) mode or management mode.

(2) **Indication Channel Mode Selection.** The indication channel has two modes of operation, MAC mode for reporting received data, and management mode for sending management information. The RXSYM3 circuit indicates which operating mode the DCE is in.

(a) **Indication Channel MAC Mode.** The MAC mode is indicated when RXSYM3 = "1" and is considered the "normal" (i.e., receive) mode. The DCE sends physical data units encoded as MAC symbols (*silence*, *non_data*, *bad_signal*, *one*, and *zero*) on circuits RXSYM0, RXSYM1, and RXSYM2 to the DTE. These are indications of data unit reception. As previously stated, these symbols are synchronized to RXCLK.

(b) **Indication Channel Management Mode.** Management mode is indicated when RXSYM3 = "0". The indication channel is used by the DCE sublayer management entity to send management indications and data to the DTE. Physical Layer management is discussed in 10.3.

10.2.2.4 Indication Channel Encoding. The encoding of the RXSYM circuits in MAC mode is summarized in Table 10-2.

Table 10-2
Indication Channel Encoding—MAC Mode

	RXSYM3	RXSYM2	RXSYM1	RXSYM0
silence	1	1	1	X
non_data	1	1	0	X
pad_idle	1	0	1	X
one	1	0	0	1
zero	1	0	0	0

NOTE: X—Indicates don't care.

The electrical representation of a "1" and a "0" is specified in 10.4. Encoding of RXSYM circuits when in management mode is described in 10.3.2.5.

10.3 DCE Management. In addition to serving as a communication link between sublayer protocol entities, the DTE-DCE interface provides communication between sublayer management entities. In this capacity, the request and indication channels pass management information.

10.3.1 Overview of Management Mode. The following two sections, 10.3.1.1 and 10.3.1.2, describe informally the functions of the request channel and the indication channel when in management mode. A detailed description of the communication protocol follows.

10.3.1.1 Management Mode Request Channel. The request channel enters management mode by impressing a "0" on the TXSYM3 circuit as described in previous sections. When in management mode, the request channel can communicate the following messages:

(1) **Reset**—Commands the DCE to initialize the DCE sublayer entity, disable the transmitter, and enable loopback mode (if supported). The reset command will override the execution of any previous commands. A NAK (non acknowledgment) to a reset command is an error indication.

(2) **Disable Loopback**—Commands the DCE to disable loopback. If loopback is not implemented, the DCE shall respond to a loopback disable command with either an ACK (acknowledgment) or a NAK.

(3) **Enable Transmitter**—Commands the DCE to enable the selected transmitter.

NOTE: Following a reset of the DCE, both Disable Loopback (if implemented), and Enable Transmitter shall be commanded before the station may transmit onto the medium.

(4) **Serial Station Management (SM) Data**—In management mode, management data can be sent serially to the DCE. This is intended for use in the management of sophisticated DCEs (primarily broadband modems). The serial SM data mode is described in more detail in 10.3.2.6 and is optional.

(5) **IDLE**—The IDLE command is a state in which SM data mode is selected but the serial data circuit is idle (marking).

10.3.1.2 Management Mode Indication Channel. The indication channel indicates DCE management mode by impressing a "0" on the RXSYM3 circuit as described previously. There are two cases in which the DCE will indicate management mode:

(1) In response to a request via TXSYM3 = "0".

(2) At the initiative of the DCE to report an error event.

When in management mode, the indication channel communicates the following messages to the DTE through encoding of the interface circuits:

(a) ACK (acknowledgment)—Indicates an acknowledgment to a DTE request.

(b) NAK (nonacknowledgment)—Indicates a negative acknowledgment to a DTE request.

(c) IDLE—Indicates nonactivity of the indication channel, or that the DCE is busy processing a request.

(d) Physical Layer Error—Indicates a DCE fault.

10.3.2 Management Protocol. The communication protocol for management mode is specified in this subsection.

10.3.2.1 Management Mode Selection. In response to DTE management mode request (TXSYM3 = "0"), the DCE shall enter management mode and indicate confirmation (RXSYM3 = "0"). The DCE shall remain in management mode until the DTE removes the request (TXSYM3 = "1"). Management commands can occur at any time. The DCE shall respond to a command with an ACK, NAK, or IDLE within the time period specified in 10.3.2.2, independent of the previous state of RXSYM3. An exception to this is Physical Layer error reporting (see 10.3.2.3).

10.3.2.2 Management Mode Commands. The following rules define the handshaking protocol to be followed by the DTE (when issuing commands) and the DCE (when responding to commands):

(1) **DTE Commands.** The DTE shall issue commands as follows. Commands other than IDLE shall be held until receiving the response (ACK or NAK) and may be held indefinitely. The IDLE command may be exited without waiting for the IDLE indication.

(2) **DCE Responses.** In response to commands from the DTE the DCE shall indicate the following:

(a) In response to a reset, a disable loopback, or an enable transmitter command, the DCE shall respond with an ACK, a NAK, or an IDLE within 32 MAC-symbol_times (32 clock periods) after the command is received (on the request channel). If the command cannot be completed in 32 MAC-symbol_times, the DCE shall report IDLE within 32 MAC-symbol_times followed by an ACK or NAK indication when command execution is completed. When processing a command other than IDLE, the DCE may report IDLE for a maximum of 3 seconds. If IDLE is reported for longer than this period, the DTE may terminate the command and assume a service failure. When responding to these commands, RXSYM0 is always a "1". The reset command aborts any previous commands and shall be responded to with an ACK or IDLE within 32 MAC-symbol_times.

NOTE: The 32 clock periods are measured at the DCE interface. The maximum delay as measured at the DTE interface is 32 clock periods plus 6.4 μs (see 10.5.4).

(b) For a serial SM data command, an ACK, NAK, IDLE, or PHYSICAL ERROR indication shall be reported within 41 MAC-symbol_times after the start bit has been received, and will then be followed with an appropriate SM data response (which is covered in more detail in the SM data protocol, 10.3.2.6).

(c) The DCE shall respond to the IDLE command by asserting IDLE on the indication channel within 32 MAC-symbol_times. If another command is received before IDLE has been received for a duration of 32 MAC-symbol_times, the DCE may respond to the new command without giving the IDLE indication.

(d) An ACK or NAK indication shall be held until receipt of another command or termination of management mode. After receipt of a new management command, the new (then current) status shall be indicated within 32 MAC-symbol_times.

(e) All unimplemented serial commands shall elicit a NAK response.

10.3.2.3 DCE-Initiated Operation. The DCE enters management mode, other than at the request of the DTE, to report an error event. The DCE indicates the error state when the DCE has a fault (for example, jabber-inhibit). In this case, the DCE shall indicate with the RXSYM3 circuit that normal (data) mode has been terminated and that the DCE is now in management mode; furthermore, the DCE shall indicate error through the encoding of the RXSYM circuits. The error indication shall be held until cleared by a reset command. For DCEs with serial data capability and several error possibilities, it is recommended that the ability to read the cause of the error after the reset command be provided.

10.3.2.4 Request Channel Encoding. The encoding of the TXSYM circuits in management mode is summarized in Table 10-3.

Table 10-3
Request Channel Encoding—Management Mode

	TXSYM3	TXSYM2	TXSYM1	TXSYM0
Reset	0	1	1	1
Disable loopback	0	1	0	1
Enable transmitter	0	0	1	1
Serial SM one (or stop bit or IDLE)	0	0	0	1
Serial SM zero (or start bit)	0	0	0	0

The electrical representations of a “1” and a “0” are specified in 10.4. All unused management mode states are not permitted.

10.3.2.5 Indication Channel Encoding. The encoding of the RXSYM circuits in management mode is summarized in Table 10-4.

Table 10-4
Indication Channel Encoding—Management Mode

	RXSYM3	RXSYM2	RXSYM1	RXSYM0
Physical layer error	0	1	1	1
NAK (nonacknowledgment)	0	1	0	*
ACK (acknowledgment)	0	0	1	*
Idle	0	0	0	0

*RXSYM0 contains SM data when responding to a serial data command. A “1” indicates a serial SM one or a stop bit. A “0” indicates a serial SM zero or a start bit. When responding to other commands, RXSYM0 is a “1”.

The electrical representations of a “1” and a “0” are specified in 10.4. All unused management mode states are not permitted.

NOTE: The DTE-DCE interface encoding is arranged to allow subscriber nodes to be implemented with a minimum hardware configuration. For simple nodes (such as most non-broadband nodes), the basic reset, loopback, and transmitter control functions can be implemented with a minimum number of circuits. In this case, the use of serial SM data transfer need not be supported. On the other hand, complex DCEs are supported as required.

10.3.2.6 Serial SM Data Protocol. In order to manage more sophisticated DCEs, additional information needs to be transferred between management entities in the DTE and the DCE. The simple commands provided with the basic interface are not sufficient to meet sophisticated management needs. The serial SM data interface allows data exchange between management and the DCE. For DCEs that do not support serial SM data transfer refer to 10.3.2.6.2.

10.3.2.6.1 DTE to DCE SM Data Protocol. When the DTE is in management mode (TXSYM3 = "0"), the serial SM data command (TXSYM2 = "0", TXSYM1 = "0") is used to send data to the DCE. Data is sent on circuit TXSYM0. TXSYM0 = "1" represents a one or a stop bit. TXSYM0 = "0" represents a zero or a start bit. The data transmission format (on the TXSYM0 circuit) consists of the following fields:

first bit transmitted

START BIT	DATA	STOP BIT

```
where
START BIT = a "0" indicating start of octet
     DATA = an 8-bit field that supplies either the address of
            the register to be accessed (read or written), or
            the data to be written (LSB first)
 STOP BIT = a "1" indicating end of octet
```

A DTE serial SM data command consists of two framed octets of data sent to the DCE and two framed octets of response data from the DCE (one framed response octet for each framed command octet). The sequence shall be as follows:

(1) The DTE sends a framed octet to the DCE containing the address of the register (in the DCE) to be accessed and a bit indicating whether that register is to be read or written.

(2) The response from the DCE (assuming an ACK) is an echo of the first framed octet received from the DTE.

(3) The DTE sends the second framed octet containing the data to be written. In the case of a register read, the second framed octet is a dummy framed octet (ignored by the DCE).

(4) The second response from the DCE is the value of the register. In the case of a register write command, the second response is the new register value after writing.

The format of the first framed octet is as follows:

first bit transmitted

START BIT	ADDRESS (6 BITS) LSB MSB	RESERVED (1 BIT)	READ/ WRITE* (1 BIT)	STOP BIT

```
where
  ADDRESS is the number of the register to be accessed.
  The RESERVED bit shall be set to "0".
  READ/WRITE* when a "1" indicates that the register in
    the DCE is to be read. READ/WRITE* when a "0" indicates
    that the register in the DCE is to be written.
```

10.3.2.6.2 DCE to DTE SM Data Protocol. A DCE that *does not* support serial commands is permitted to respond with a NAK and a start bit before receiving the stop bit from the DTE. DCEs that *do* support serial commands shall take the following actions in response to a serial command (start bit, eight data bits, and a stop bit):

(1) Receive the start bit, eight data bits, and the stop bit.
(2) Within 41 MAC-symbol_times of receiving the start bit, respond with an ACK, a NAK, IDLE, or a PHYSICAL ERROR.
(3) If an IDLE was used, change to ACK or NAK upon completion of command execution. When processing a command other than IDLE, the DCE may report IDLE for a maximum of 3 seconds. If IDLE is reported for longer than this period, the DTE may terminate the command and assume a service failure. The DTE may not initiate a new command until receiving an ACK, a NAK, or a PHYSICAL ERROR, or until timeout of the current command.
(4) While holding the ACK or NAK on RXSYM2, 1, output on RXSYM0 a start bit, a response octet, and a stop bit. The ACK or NAK shall be valid at or before the start bit and this ACK or NAK shall be held during transmission of the data and the stop bit.

Data Transmission Format. The data transmission format consists of the following fields:

first bit transmitted (START BIT)

START BIT	DATA	STOP BIT

```
where
START BIT = a "0" indicating start of octet
     DATA = an 8-bit field representing the data (LSB first)
 STOP BIT = a "1" indicating end of octet
```

The use of the data field is explained below.

First Octet Response. In response to the first framed octet from management —

- In an acknowledged response (RXSYM2,1= 0,1), the DATA field is equal to the first framed octet from management as received by the DCE.
- In a non-acknowledged response (RXSYM2,1= 1,0), the DATA field is don't care. The NAK signal plus the start bit represents confirmation that the command was bad or not implemented.

Second Octet Response. In response to the second framed octet from management—

- In an acknowledged (RXSYM2,1 = 0,1) register write command, the DATA field is the new contents of the register (LSB transmitted first). The ACK symbol plus the start bit represents confirmation.
- In an acknowledged (RXSYM2,1 = 0,1) register read command, the DATA field is the contents of the addressed register (LSB transmitted first). The ACK symbol plus the start bit represents confirmation.
- In an nonacknowledged (RXSYM2,1 = 1,0) register read or write command, the DATA field is don't care. The NAK symbol plus the start bit represents confirmation that the command was bad, non-implemented, or that the register does not support the requested action.

10.3.2.6.3 Error Recovery. Error detection and recovery on DTE-DCE SM transactions is the responsibility of the DTE. When a register is written, the new contents of that register are repeated to DTE management to facilitate the detection of incorrect completion of the write. If additional error checking is required and the register is readable, the register can be reread any number of times. If the first framed octet of a serial command is NAKed, the DCE shall also NAK the second (next) framed octet. The DCE shall always be prepared for the first framed octet of a serial command when first entering management mode, and after receiving a reset command, a disable loopback command, or an enable transmitter command. The following typical write sequence is provided as an illustrative example:

From DTE

first framed octet, write to register 6

```
  0      011000      0         0       1
START   ADDRESS   RESERVED   WRITE    STOP
```

To DTE

first framed octet echoed back

```
  0      011000      0         0       1
START   ADDRESS   RESERVED   WRITE    STOP
```

From DTE

second framed octet, hex data 62

```
  0     01000110     1
START     DATA      STOP
```

To DTE

contents of register 6 after writing = hex E2 (an error)

```
  0     01000111     1
START     DATA      STOP
```

DTE management is responsible for recovering from all errors including those due to noise. Therefore, after writing a register, it is recommended that the echo of the command be compared with the request. In addition, further error checking can be accomplished by reading the contents of readable registers. For this example, the contents as written were not the same as requested to be written.

10.3.2.6.4 Register Assignment. Only registers A–F are reserved, all other registers are user defined. Registers D–F are reserved for modem ID. If modem ID is not supported, the DCE shall NAK a read of locations D–F.

10.4 Electrical Characteristics. The electrical characteristics of the request and indication channel circuits are specified here. Two electrical signaling methods are specified. The first employs single-ended (unbalanced) transmission, the second employs differential (balanced) transmission. Associated with each electrical signaling method is a mechanical specification for the physical cabling. Single-ended signaling is intended for use over short distances within a physical enclosure. Differential signaling is intended for use over greater distances, especially between separate enclosures.

10.4.1 Single-ended Signaling. Single-ended transmission may be used when the DTE-DCE interface cable is 20 cm or less in length. It is intended for "board to board" communication within an enclosure. Care should be taken to ensure that the design of the DTE and DCE account for losses, reflections, and voltage drops in the DTE-DCE interface.

10.4.1.1 Signal Characteristics. A "0" impressed on any circuit shall be electrically represented as a low voltage as defined below. A "1" impressed on any circuit shall be electrically represented as a high voltage as defined below.

10.4.1.2 Driver Characteristics. The driver is a TTL compatible single-ended driver capable of driving the specified interface cable and load shown in Fig 10-2. The following shall be the line driver output characteristics:

Output Low < 0.5 V @ 8.0 mA
Output Hi Output Low < 2.5 V @ –1.0 mA

10.4.1.3 Receiver Characteristics. The suggested line input characteristics of the receiver circuit (including buffer and pullup resistor) shall be as follows:

Low level input current < –2.0 mA @ VIL = 0.5 V
High level input current < 50 μA @ VIH = 2.7 V
Input high voltage > 2.0 V
Input low voltage < 0.8 V

Failsafe operation is achieved through the use of a 15 kΩ pullup resistor connected as shown in Fig 10-2. If the cable is disconnected, the receiver output shall indicate a high state on the received line.

10.4.1.4 Power Characteristics. DTE supplied power is optional. Documentation shall be supplied with the DTE indicating the power sourcing capability. The following is the recommended minimum values for power supplied over the interface:

+5 V ±5 % @ 1.50 A
+12 V ±10% @ 0.50 A
–12 V ±10% @ 0.25 A

10.4.1.5 Signal Timing. Signal timing is shown in Figs 10-3 and 10-4. All signals on the indication channel are referenced to RXCLK where the frequency of RXCLK is nominally equal to the MAC-symbol rate within the requirements of the appropriate PLE defined by this standard. All signals on the request channel are referenced to TXCLK. As with RXCLK, the frequency of TXCLK is equal to the MAC-symbol rate. For TXCLK and RXCLK the high time and the low time shall each not be less than 40% of the nominal clock period. (The high times are measured at the 90% point and the low times are measured at the 10% point.) When the DCE is disconnected from the medium or loopback is enabled, the permitted frequency tolerance for RXCLK and TXCLK is relaxed to ±5% of that nominal frequency. TXCLK and RXCLK shall meet all other specifications of this standard for supported PLEs.

10.4.2 Differential Signaling. Differential signaling shall be used for cable distances greater than 20 cm. The differential signaling method defined here employs twisted shielded pairs making connection between separate enclosures practical.

10.4.2.1 Signal Characteristics. A "0" impressed on any circuit shall be electrically represented as a negative differential voltage. A "1" impressed on any circuit shall be electrically represented as a positive differential voltage. The differential voltage is equal to the voltage of the plus (+) signal minus the voltage of the minus (–) signal of a differential pair.

Fig 10-2
Single-ended Interface Connection

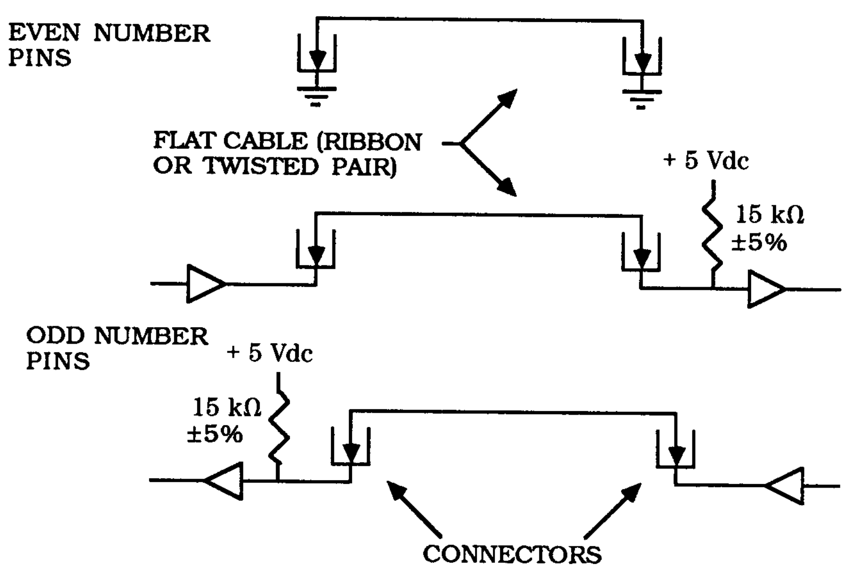

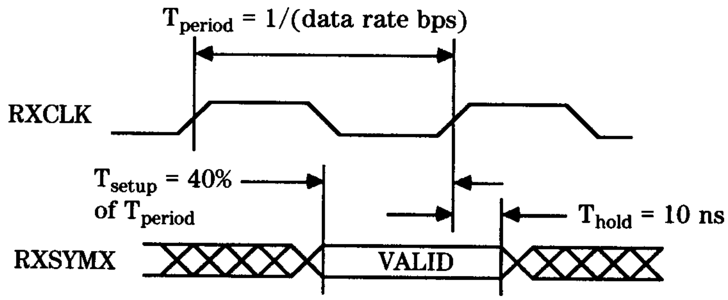

Fig 10-3
Indication Channel Timing
(Shown at DTE Connector)

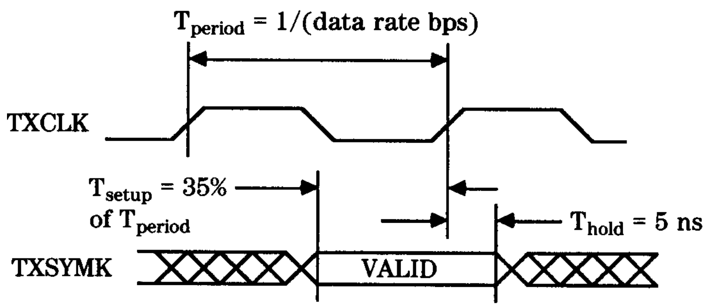

Fig 10-4
Request Channel Timing Using TXCLK
(Shown at DCE Connector)

10.4.2.2 Driver Characteristics. The driver shall provide a differential voltage of 580 mV minimum and 1220 mV maximum into a 100 Ω differential load. The dc common mode voltage shall be 3.66 ± 0.4 V.

NOTE: This requirement can be met by using emitter-coupled logic (ECL) gates of the 10K, 10KH, and 100K families operating at 5.0 V ±5% with 270 Ω pulldown resistors.

10.4.2.3 Receiver Characteristics. The receiver shall operate with a differential input of 350 mV minimum to 1220 mV maximum with a dc common-mode voltage of 3.66 ± 0.4 V. The termination impedance shall be 100 Ω ±5% for the single resistor termination, and ±1% matching for the dual resistor termination.

NOTE: This requirement can be met by using ECL gates of the 10K, 10KH, and 100K families operating at 5.0 V ± 5%. Fig 10-6 (b) is preferred for better common-mode rejection.

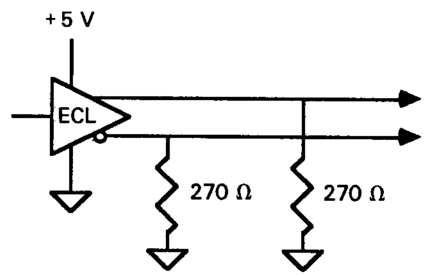

Fig 10-5
Typical Driver—Differential Interface

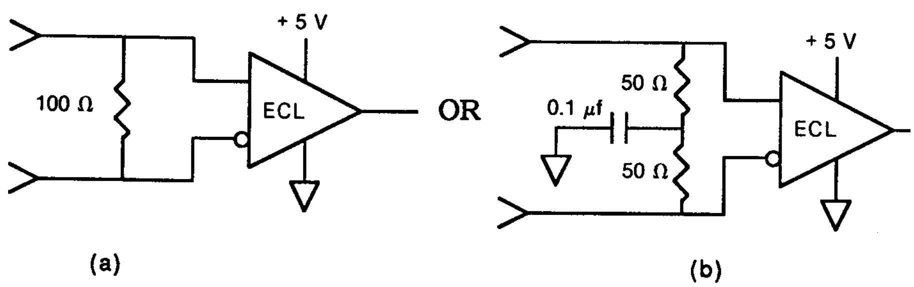

Fig 10-6
Typical Receivers—Differential Interface

10.4.2.4 DCE Power. The means by which power is supplied to the DCE is beyond the scope of this standard; however, the DCE dc power shall be referenced to the DTE signal return. This is to ensure that the common-mode level remains within a range suitable to the DTE line receivers.

10.4.2.5 Signal Timing. Signal timing is shown in Figs 10-3 and 10-7. All signals on the indication channel are referenced to RXCLK where the frequency of RXCLK is nominally equal to the MAC-symbol rate within the requirements of the appropriate PLE defined by this standard.

All signals on the request channel are referenced to RETURNCLK. Figure 10-8 shows a derivation of RETURNCLK from TXCLK. As with RXCLK, the frequency of TXCLK is equal to the MAC-symbol rate. For TXCLK and RXCLK the high time and the low time shall each be at least 40% of the nominal clock period. (The high times are measured at the 90% point and the low times are measured at the 10% point.) When the DCE is disconnected from the medium or loopback is enabled, the permitted frequency tolerance for RXCLK and TXCLK is relaxed to ±5% of that nominal frequency. TXCLK and RXCLK shall meet all other specifications of this standard for supported PLEs.

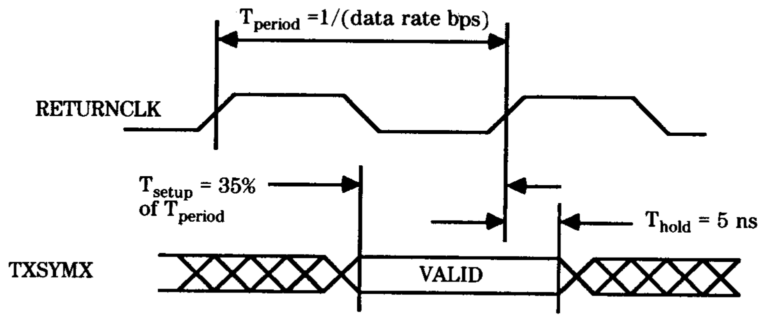

Fig 10-7
Request Channel Timing Using RETURNCLK
(Shown at DCE Connector)

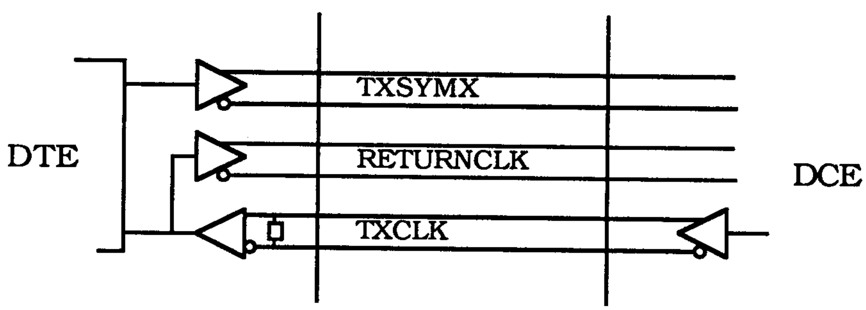

Fig 10-8
Typical Derivation of RETURNCLK

10.4.2.6 Failsafe Operation. Two signals are provided in the differential interface for the purpose of detecting a cable break or loss of power at either end of the interface. These signals are called DCE PRESENT and DTE PRESENT, and are static dc levels generated with a 470 Ω ±5% resistor pulled up to +5 V ±10% as shown in Fig 10-9. The DCE shall transmit silence when loss of DTE power or a cable break is detected.

10.5 Mechanical Characteristics. The mechanical characteristics of the DTE-DCE interface are specified here. Two separate mechanical interfaces are specified: One is for use with single-ended signaling and one is for use with differential signaling. Each interface consists of cable and connectors. In each interface several pins are reserved. Users shall not drive these pins. Cables need not implement reserved connections.

10.5.1 Single-ended Signaling Interface Cable. Flat cable (either ribbon or twisted pair) is recommended as being most suitable for use with the connector defined in 10.5.2. The interface connector described in 10.5.2 will accept a cable connector that is strain-relieved or not strain-relieved. Alternate conductors are grounded to minimize crosstalk. The minimum wire diameter shall be 0.32 mm.

10.5.2 Single-ended Signaling Interface Connectors. The connectors used shall be 40 pin dual-in-line type with 2.54 mm centers as specified in MIL-C-83503. The equipment-end connectors contain a shrouded set of male pins (similar to MIL-C-83503/20-19); the cable-end connectors shall have female receptacles (similar to MIL-C-83503/7-09). The connector shall use center polarization. A long latch may be used or no latch at all. This will ensure interoperability with cable connectors that are strain-relieved or not strain-relieved. (See Fig 10-10.)

NOTE: The creation of an international standard for this connector is under consideration.

Fig 10-9
Cable Detection Circuits

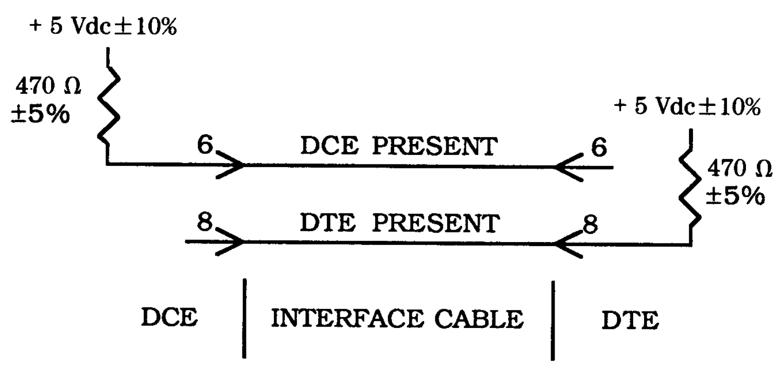

Fig 10-10
Single-ended Interface Connector Receptacle

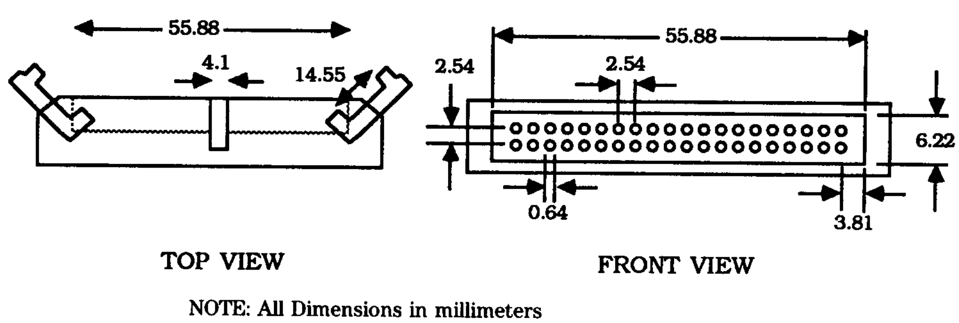

10.5.3 Connector Pinout. The signal pinout is shown in Table 10-5.

Table 10-5
Single-ended Connector Pinout

Pin Number	Signal Name	Source
1	RXSYM0	DCE
3	RXSYM1	DCE
5	RXSYM2	DCE
7	RXSYM3	DCE
9	RXCLK	DCE
11	+5 Vdc	DTE
13	+5 Vdc	DTE
15	+5 Vdc	DTE
17	TXCLK	DCE
19	RESERVED	—
21	TXSYM3	DTE
23	TXSYM2	DTE
25	TXSYM1	DTE
27	TXSYM0	DTE
29	RESERVED	—
31	RESERVED	—
33	RESERVED	—
35	RESERVED	—
37	+12 Vdc (Analog supply)	DTE
39	–12 Vdc (Analog supply)	DTE

NOTE: Even numbered pins 2–40 are connected to logic ground at both the DTE and the DCE.

10.5.4 Differential Signaling Interface Cable. Twisted pair cable of 100 Ω ±5% and minimum diameter of 0.32 mm or greater with an overall shield is recommended. The maximum attenuation shall be 4.4 dB or less at the nominal clock frequency. Specific shielding and crosstalk requirements are under study. The maximum one-way propagation delay of all signals shall be less than 3.2 µs.

10.5.5 Differential Signaling Interface Connectors. The connectors specified are 37 pin "D" subminiature type as specified in IEC Publication 807-2 (1985) [5]. Female contacts and a male shell shall be on the DTE equipment. Male contacts and a female shell shall be on DCE equipment. The cable shall use appropriate mating connectors. Screw-locks shall be used for cable retention. The female screw-lock shall be used on the DTE and DCE bulkheads and the male screw-lock shall be used on the cable ends. Both connectors shall use shielded housings. The connector pin assignments are given in Table 10-6.

Table 10-6
Differential Connector Pinout

Pin Number	Signal Name	Source
1	RXSYM0 +	DCE
20	RXSYM0 –	DCE
2	RXSYM1 +	DCE
21	RXSYM1 –	DCE
3	RXSYM2 +	DCE
22	RXSYM2 –	DCE
4	RXSYM3 +	DCE
23	RXSYM3 –	DCE
5	RXCLK +	DCE
24	RXCLK –	DCE
6	DCE PRESENT	DCE
25	SIGNAL GROUND	DTE
7	RESERVED	—
26	RESERVED	DTE
8	DTE PRESENT	DTE
27	SIGNAL GROUND	DTE
9	TXCLK +	DCE
28	TXCLK –	DCE
10	RETURNCLK +	DTE
29	RETURNCLK –	DTE
11	TXSYM3 +	DTE
30	TXSYM3 –	DTE
12	TXSYM2 +	DTE
31	TXSYM2 –	DTE
13	TXSYM1 +	DTE
32	TXSYM1 –	DTE
14	TXSYM0 +	DTE
33	TXSYM0 –	DTE
15	RESERVED	—
34	RESERVED	—
16	RESERVED	—
35	RESERVED	—
17	RESERVED	—
36	RESERVED	—
18	RESERVED	—
37	RESERVED	—
19	(Do Not Connect)	—
SHELL	FRAME GROUND	DTE & DCE

Appendix

(This Appendix is not a part of this International Standard or of IEEE Std 802.4-1990, but is included for information only.)

10.6 Appendix—Clarification of Management Mode. The following sequence diagrams have been added to help clarify the management mode functions.

Figure 10-11 shows a DTE entering management mode, issuing two commands to the DCE (which are acknowledged), and then returning to MAC mode. Figure 10-12 shows a DTE entering management mode, writing the data A5 (hex) to register 23 (hex), then returning to MAC mode. Figure 10-13 shows a DCE reporting an error event.

Figure 10-11
Typical Station Management

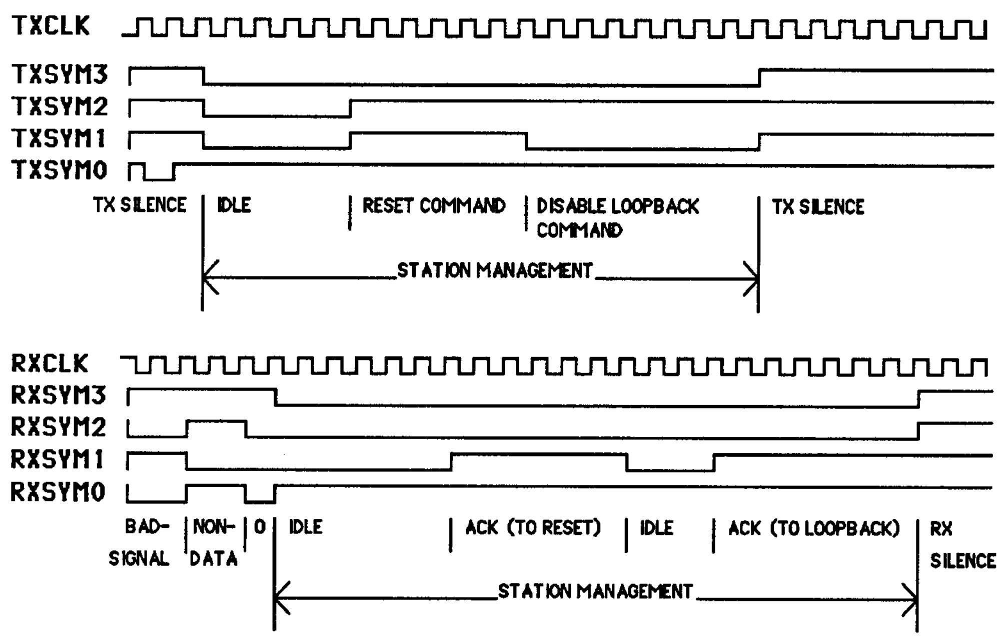

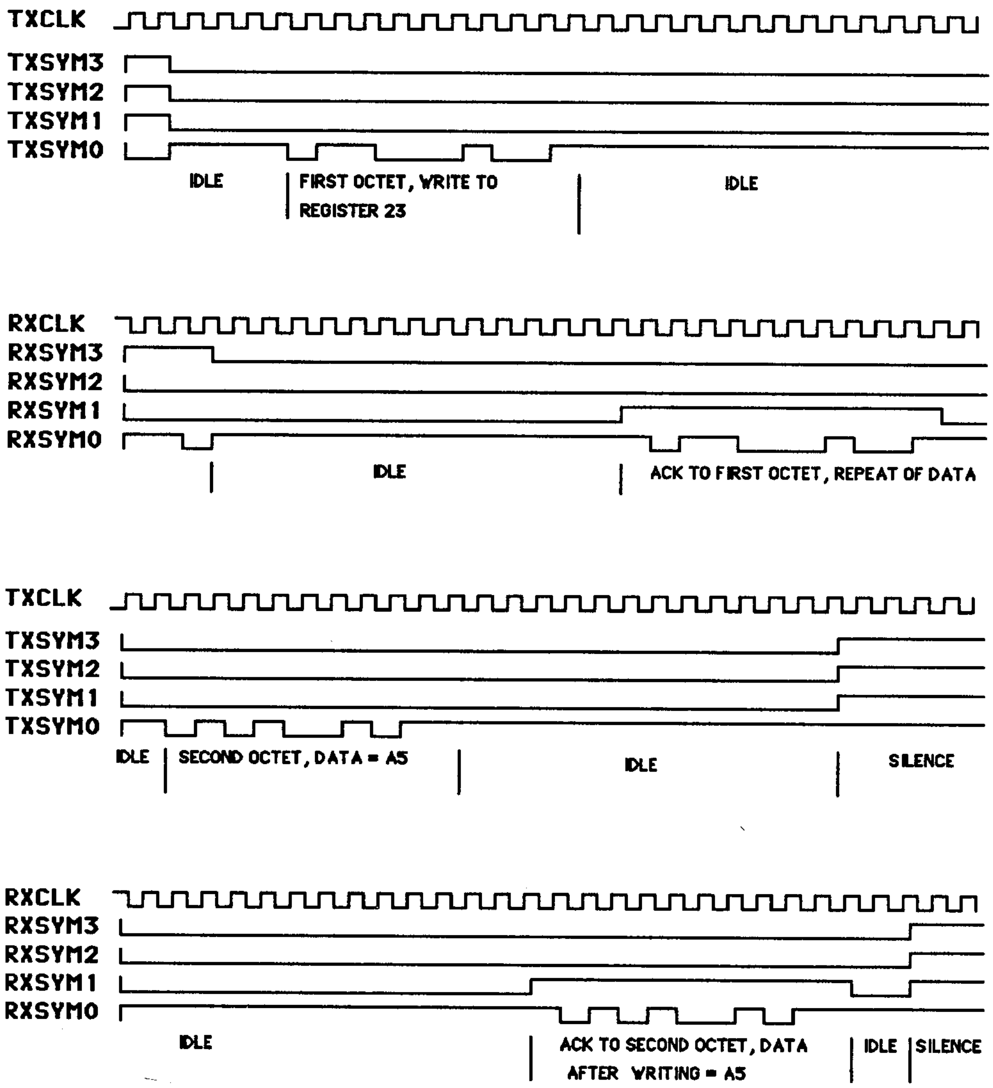

Fig 10-12
Writing to a Register

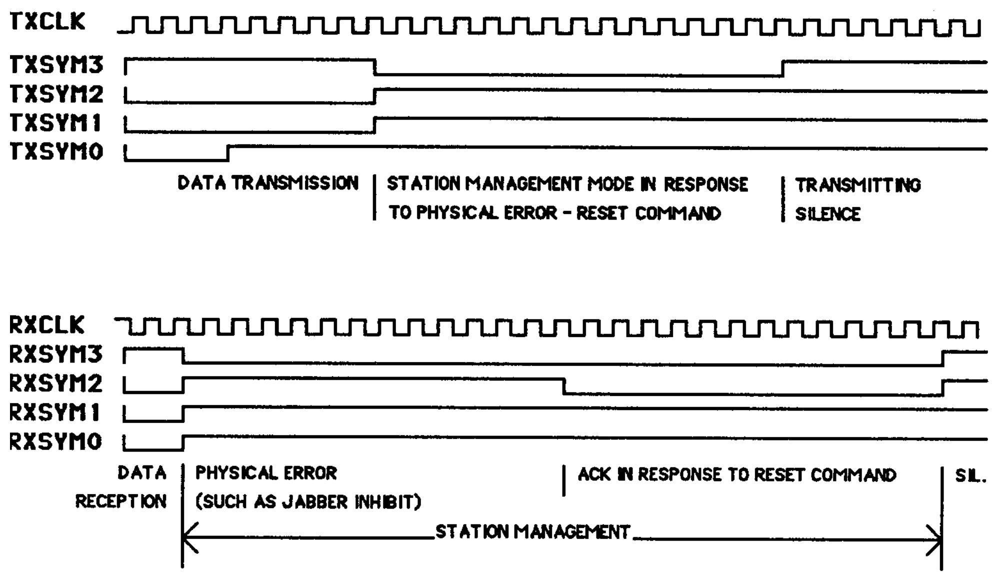

Fig 10-13
Typical Modem Failure Sequence

11. RESERVED

This section number is reserved in order to maintain the numbering of ensuing sections.

12. Single-Channel Phase-Coherent-FSK Bus Physical Layer Entity (PLE)

The functional, electrical, and mechanical characteristics of one specific form of PLE (single-channel phase-coherent-FSK bus) of this standard are specified in this section. This specification includes the PLEs found in stations that can attach to the single-channel phase-coherent-FSK bus Local Area Network (LAN). The relationship of this section to other sections of this standard and to LAN specifications is illustrated in Fig 12-1. The relationship of this section to the single-channel phase-coherent-FSK bus PLE and medium is illustrated in Fig 12-2.

This standard specifies these PLEs only as necessary to ensure

(1) The interoperability of implementations conforming to this specification, and

(2) The protection of the LAN and those using it.

Fig 12-1
Relation to LAN Model

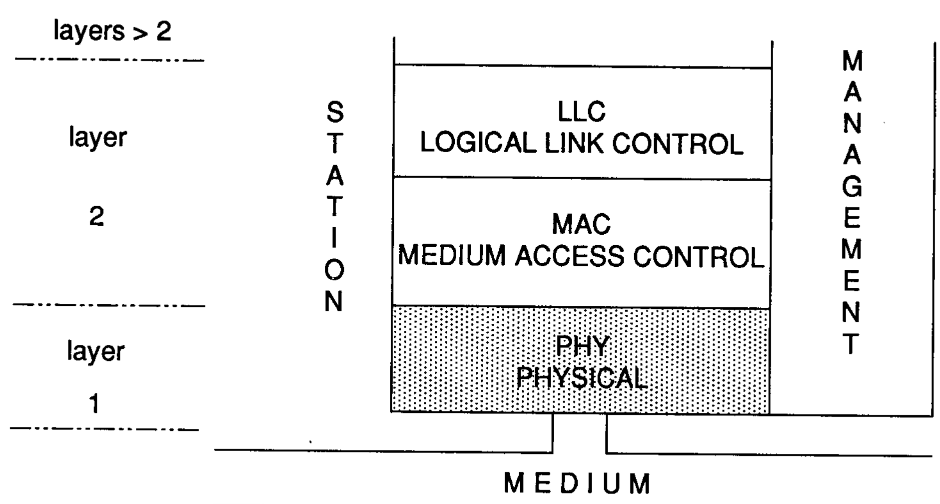

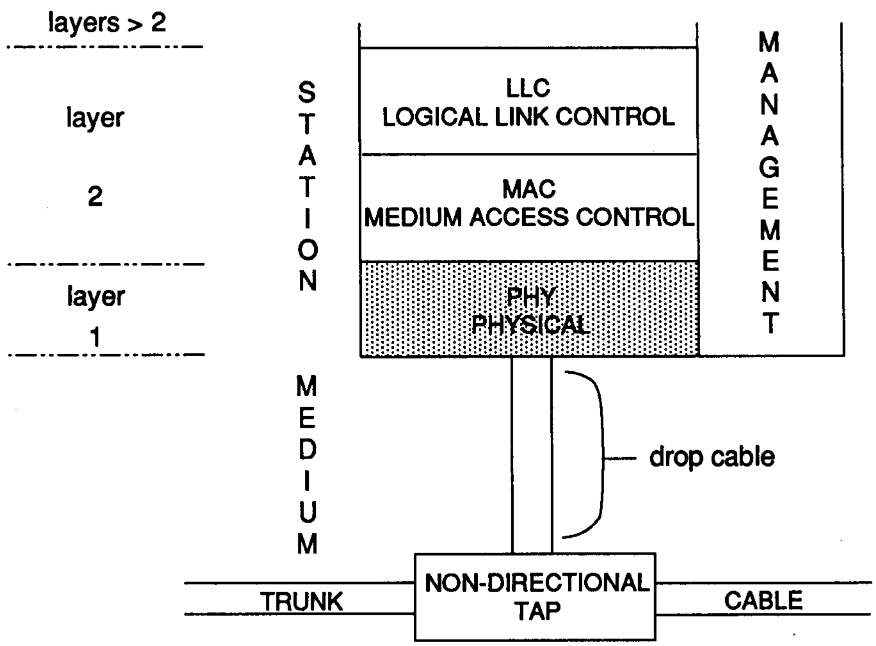

Fig 12-2
Physical Hardware Partitioning

12.1 Nomenclature. Terms used in this section, whose meanings within the section are more specific than indicated in the glossary of IEEE P802.1B [10], are as follows:

dBmV. A measure of rms signal level on a 75 Ω cable compared to 1 mV. In terms of SI units (International System of Units), dBmV is defined as dB (1 mV, 75 Ω) rms.

detected bit error. An error that is reported as *bad_signal*. *Bad_signal* reported during the preamble or during the four symbols following the last end delimiter of a transmission is not included.

drop cable. The coaxial cable of the medium that connects to a station.

FSK (frequency shift keying). A modulation technique whereby information is impressed upon a carrier by shifting the frequency of the transmitted signal to one of a small set of frequencies.

medium. The cable system of a LAN, including the trunk cable, taps, drop cables, and splitters. These components are defined in Section 13.

operating frequency range. The operating frequency range is defined as 2–15 MHz for a 5 Mb/s data rate and 4–30 MHz for a 10 Mb/s data rate.

phase-coherent FSK. A particular form of FSK where the two signaling frequencies are integrally related to the data rate and transitions between the two signaling frequencies are made at zero crossings of the carrier waveform.

regenerative repeater. A device used to extend the length, topology, or interconnectivity of a LAN beyond that imposed by the minimum transmit and receive level specifications of the station and the connectivity restrictions of the medium. Regenerative repeaters perform the basic actions of restoring signal amplitude, waveform, and timing.

single-channel system. A system where at any point on the medium only one information signal at a time can be present without disruption.

trunk cable. The main cable of the medium, which interconnects the taps.

undetected bit error. An error that is not reported as *bad_signal* by the PLE.

12.2 Object. The object of this specification is to

(1) Provide the physical means necessary for communication between local network stations employing the LAN token-passing bus access method and using the single-channel phase-coherent-FSK bus medium described in this standard.
(2) Define a physical interface that can be implemented independently among different manufacturers of hardware and achieve the intended level of compatibility when interconnected to a common single-channel phase-coherent-FSK bus LAN medium.
(3) Provide a communication channel capable of high bandwidth and low bit error rate.
(4) Provide for ease of installation and service in a wide range of environments.
(5) Provide for high network availability.
(6) Facilitate low-cost implementations.

12.3 Compatibility Considerations. This standard applies to PLEs that are designed to operate on a 75 Ω coaxial cable configured in a trunk and drop cable structure as defined in Section 13. Compatibility with this medium is determined at the medium interface. Compatible physical entities shall use the same signaling rate.

12.4 Medium Overview. The communications medium specified in Section 13 consists of a trunk and drop structure, with branching (*splitters*) possible. The medium's trunk cable is connected to the drop cables by non-directional passive impedance-matching networks (*taps*), and the drop cables in turn are connected to the stations. Extension of the topology or size of the medium is accomplished by active regenerative repeaters.

12.5 PLE Overview

12.5.1 General Description of Functions. The functions performed by the single-channel phase-coherent-FSK bus PLE are described informally here.

Jointly, these functions provide a means whereby Medium Access Control (MAC)-symbols presented at the MAC interface of one PLE can be conveyed to all of the PLEs on the bus for presentation to their respective MAC interfaces.

12.5.1.1 Symbol Transmission and Reception Functions. Successive MAC-symbols presented to the PLE at its MAC service interface are applied to an encoder that produces as output a three PHY-symbol code: {*H*}, {*L*}, {*off*}.

That output is then applied to a two-tone FSK modulator which represents each {*H*} as one full cycle of a tone whose period is exactly one-half of the MAC-symbol_time, each {*L*} as one half cycle of a tone whose full-cycle period is exactly the MAC-symbol_time, and each {*off*} as no tone for one-half MAC-symbol_time. This modulated signal is then ac-coupled to the single-channel phase-coherent-FSK bus medium and conveyed by the medium to one or more receivers.

Each receiver is also ac-coupled to the bus medium. It bandpass filters the received signal to reduce received noise, demodulates the filtered signal and then infers the transmitted PHY-symbol from the presence of carrier and the frequency of the received signal. It recovers the timing of the transmitted PHY-symbols directly from the representation of the PHY-symbols on the medium. It then decodes that inferred PHY-symbol by an inverse of the encoding process and presents the resultant decoded MAC-symbols at its MAC service interface.

For all MAC-symbols except *pad_idle*, this decoding process is an exact inverse of the encoding process in the absence of errors. The *pad_idle* symbols, which are referred to collectively as *preamble*, are transmitted at the start of each MAC frame, both to provide a training signal for receivers and to provide a nonzero minimum separation between consecutive frames. Since each transmission begins with *pad_idle* symbols, it is expected that some of these initial symbols may be "lost in transit" between the transmitter and receivers. Additionally, in phase-coherent-FSK systems, the MAC-symbol encodings for successive *pad_idle* symbols is identical with the encodings of an alternating series of *ones* and *zeros*. Receivers are permitted to decode the transmitted representation of successive *pad_idle* symbols as an alternating series of *ones* and *zeros* and report it as such to the MAC entity.

12.5.1.2 Regenerative Repeater Functions. Regenerative repeaters may be used to extend the network size beyond the maximum signal loss budget of a single medium segment. They do so by connecting two or more medium segments and repeating anything "heard" on one segment to the other segments. For the purposes of this standard, regenerative repeaters are considered stations, whether or not they have functionality beyond that of a repeater.

12.5.1.3 Jabber-Inhibit Function. To protect the LAN from most faults in a station, each station contains a jabber-inhibit function. This function serves as a "watchdog" on the transmitter; if the station does not turn off its transmitter after a prolonged time (approximately one-half second), then the transmitter output shall be automatically disabled for at least the remainder of the transmission.

12.5.1.4 Local Administrative Functions (OPTIONAL). These functions are activated either manually, or by way of the PLE's management interface, or both. They can include

(1) Enabling or disabling each transmitter output. (A redundant medium configuration has two or more transmitter outputs.)

(2) Selecting the received signal source: any medium (if redundant media are present) or any available loopback point.

NOTE: If a loopback point is selected, then all transmitter output to the medium shall be inhibited.

12.5.2 Basic Functions and Options. Symbol transmission and reception functions and jabber-inhibit functions are required in all implementations. All other functions are optional.

12.6 Application of Management. The following constraints are imposed on the parameters and actions specified in Section 9:

(1) In the capabilities group the dataRates parameter sequence shall specify one or more of the values 5 and 10.

When implemented, the receiveChannels and transmitChannels parameters shall specify identically either a null set or the set of frequencies with values equal to 1.5 times the values in the dataRates parameter sequence.

(2) In the local management information the value of minPostSilencePreamble Length shall be 6.

12.7 Functional, Electrical, and Mechanical Specifications. Unless otherwise stated, all voltage and power levels specified are in rms and dBmV, respectively (dBmV is defined as dB [1 mV, 75 Ω] rms in 12.1). Specifications shall be met by the fundamental frequency component during the continuous transmissions of all *one* symbols and also be met by the fundamental frequency component during the continuous transmission of all *zero* symbols.

12.7.1 Data Signaling Rates. The standard data signaling rates for phase-coherent-FSK systems are 5 Mb/s and 10 Mb/s. The permitted tolerance for each signaling rate is ±0.01% for an originating station, and is ±0.015% for a repeater station while repeating. When a composite PLE is embodied in a regenerative repeater, it shall originate signaling on all media at exactly the same data rate.

12.7.2 Symbol Encoding. The PLE transmits symbols presented to it at its MAC interface by the MAC sublayer entity. The possible input MAC-symbols are (see 8.2.1): *zero*, *one*, *non_data*, *pad_idle*, and *silence*. Each of these MAC-symbols is encoded into a pair of PHY-symbols from a different three-symbol {*H*}, {*L*}, {*off*} code and then transmitted. The encoding action to be taken for each of the input MAC-symbols is

(1) *Silence*—Each *silence* symbol shall be encoded as the sequence {*off off*}.

(2) *Pad_idle*—*Pad_idle* symbols are always originated in octets. Each pair of consecutive *pad_idle* symbols shall be encoded as the sequence {*L L*} {*H H*} with the {*L L*} as the first transmitted symbol.

With the exception of a regenerative repeater, when sending *pad_idle* MAC-symbols immediately after *silence* MAC-symbols, the MAC sublayer entity shall send at least 6 octets of *pad_idle* MAC-symbols. This minimum post silence preamble is composed of two parts:

(a) The first part, of 3 octets duration, is a delay to compensate for possible receiver blanking at remote stations as described in 12.7.6.3.

(b) The final 3 octets of *pad_idle* MAC-symbols are provided for receiver synchronization.

A regenerative repeater shall repeat silence during the receiver blanking period. This may cause the preamble of the succeeding transmission to be shortened by up to three octets.

(3) *Zero*—Each *zero* symbol shall be encoded as the sequence {*H H*}.

(4) *One*—Each *one* symbol shall be encoded as the sequence {*L L*}.

(5) *Non_data*—*Non_data* symbols are transmitted by the MAC sublayer entity in pairs. Each such pair of consecutive *non_data* symbols shall be encoded as the sequence {*H* L} {*L H*}. Thus the (start delimiter) subsequence

non_data non_data zero non_data non_data zero

shall be encoded as the sequence

{*H L*} {*L H*} {*H H*} {*H L*} {*L H*} {*H H*},

and the (end delimiter) subsequence

non_data non_data one non_data non_data one

shall be encoded as the sequence

{*H L*} {*L H*} {*L L*} {*H L*} {*L H*} {*L L}*.

Figure 12-3 shows representative encodings of the MAC-symbols.

12.7.3 Transmitted Signals. The PHY-symbols resulting from the encoding of 12.7.2 shall be converted directly to their line representations, as described in 12.7.3.1, and the resultant signaling shall be coupled to the single-channel phase-coherent-FSK bus medium, as specified in 12.7.5. The polarity of the carrier signal at the start of the frame does not imply specific information about the signal. The PLE shall be able to receive signals of either polarity.

Fig 12-3
MAC-Symbol Encodings

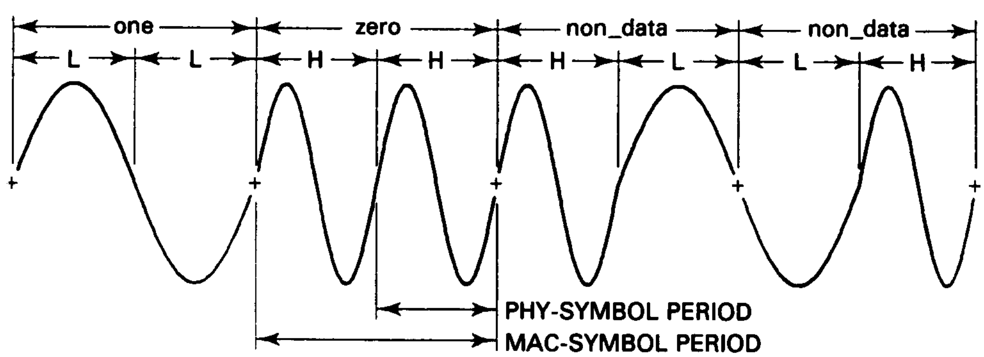

12.7.3.1 Signal Representation. The signal representations of the {*H*}, {*L*}, and {*off*} PHY-symbols shall be as follows:

(1) An {*H*} shall be represented as one full cycle of a signal, starting and ending with a nominal zero amplitude, whose period is equal to half the period of MAC-symbol delivered from the MAC entity at the MAC interface. Each {*H*} PHY-symbol shall be phased such that there is a zero-crossing between the {*H*} PHY-symbol and the preceding PHY-symbol.

(2) An {*L*} shall be represented as one half cycle of a signal, starting and ending with a nominal zero amplitude, whose full cycle period is equal to the period of MAC-symbol delivered from the MAC entity at the MAC interface. Each {*L*} PHY-symbol shall be phased such that there is a zero crossing between the {*L*} PHY-symbol and the preceding PHY-symbol.

(3) An {*off*} shall be represented by no signal for a period equal to one half of the period of MAC-symbol delivery from the MAC entity at the MAC interface. Table 12-1 summarizes this relationship.

Table 12-1
Data Rate Versus Signaling Frequencies

Data Rate (Mb/s)	Frequency of Lower Tone (MHz)	Frequency of Higher Tone (MHz)
5	5.0	20.0
10	10.0	

12.7.3.2 Jitter. The maximum jitter in the period of any {L} or any {H} shall be no more than ±1% of the MAC-symbol_time. Additionally, the jitter at the zero crossing in the middle of any {*H*} shall be no more than±1% of the MAC-symbol_time.

12.7.3.3 Output Level. The output level of the transmitted signal into a 75 Ω resistive load shall be between +63 dBmV and +66 dBmV, inclusive. The fundamental frequency power when transmitting continuous {*L*} PHY-symbols shall be within 1 dB of the fundamental frequency power when transmitting continuous {*H*} PHY-symbols.

The maximum variance of the transmitter amplitude between any two {*L*} or between any two {*H*} PHY-symbols measured within any one transmission shall be, at most, 0.5 dB.

12.7.3.4 Leakage. The residual or leakage transmitter-off output signal shall be no more than –20 dBmV in the operating frequency range.

12.7.3.5 Transition Times. The transmit signal of {*H*} and {*L*} PHY-symbols shall fit within the shaded portion of the eye pattern shown in Fig 12-4.

(1) The signal rise time, measured from the –80% to + 80% signal points (10% to 90% peak-to-peak signal amplitude), and the fall time, measured from +80% to –80% signal points (90% to 10% of peak-to-peak signal amplitude), shall be greater than 7.5% of the MAC-symbol_time. That is, the

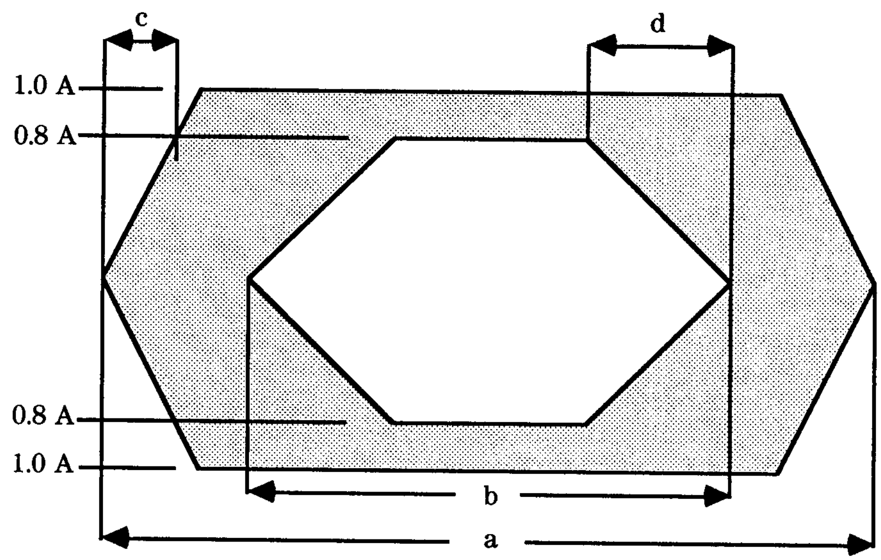

	5 Mbits/s				10 Mbits/s			
	a	b	c	d	a	b	c	d
{*L*}	102	98	7.5	35	51	49	3.75	17.5
{*H*}	52	98	7.5	17.5	26	24	3.75	8.75

NOTE: Values in nanoseconds.

Fig 12-4
Eye Patterns

rise and fall times of the transmitted signal shall be no less then 15 ns for 5 Mb/s data rate and no less then 7.5 ns for 10 Mb/s data rate.

In addition, the rise and fall times of the transmit signal, measured as above, shall be less than 35% of the period of one full cycle of the transmitted signal. That is, the rise and fall times shall be no greater than 70 ns for a 5 MHz signal, no greater than 35 ns for a 10 MHz signal, and no greater than 17.5 ns for a 20 MHz signal.

NOTE: The intent of this specification is to allow a transmitted waveshape anywhere between a slew rate limited square wave and a sinusoidal wave. If the slew rates for {*H*} and {*L*} PHY-symbols are equal, the slew rate should not exceed 12% of the nominal MAC-symbol_time in order to meet 12.7.3.3.

(2) The rise and fall times for all {*H*}s shall be within 10% of each other. The rise and fall times for all {*L*}s shall be within 10% of each other.

(3) The spectral envelope of both signaling frequencies shall be such that the harmonic content relative to the rms transmitted power is as follows:
   (a) 2nd and 3rd harmonic at least 10 dB below the fundamental frequency power.
   (b) 4th and 5th harmonic at least 15 dB below the fundamental frequency power.
   (c) 6th and 7th harmonic at least 20 dB below the fundamental frequency power.
   (d) All higher harmonics at least 25 dB below the fundamental frequency power.

(4) The transmitted signal ripple within each PHY-symbol shall be no greater than 20% of the peak amplitude of the transmitted signal.

(5) The physical signal of the transmitter shall turn on, that is, go between silence and full physical symbols, in two or less MAC-symbol_times. The symbols corresponding to the third and subsequent MAC-symbols shall be as specified by the eye patterns. With a 6-octet preamble, there shall be at least 92 full PHY-symbols before the start delimiter for a nonrepeater station and at least 44 full PHY-symbols for a repeater station.

The physical signal of the transmitter shall turn off, that is, go between full physical symbols and the receive silence level of +4 dBmV as specified in 12.7.6.2, in 2 or less MAC-symbol_times; and to the transmitter off level of –20 dBmV as specified in 12.7.3.4 in 8 or less MAC-symbol_times. The 9th and subsequent silence symbols shall meet the transmitter off specification.

12.7.4 Jabber Inhibit. Each PLE shall have a self-interrupt capability to inhibit modulation from reaching the LAN medium. Hardware within the PLE (with no external message other than the prolonged detection of an output-on condition within the transmitter) shall provide a nominal window of one-half second ±25% during which time a normal data link transmission may occur. If a transmission is in excess of this duration, the jabber-inhibit function shall operate to inhibit any further output from reaching the medium. Reset of this jabber-inhibit function is implementation-dependent.

12.7.5 Coupling to the Medium. The PLE functions are intended to operate satisfactorily over a medium consisting of a 75 Ω bidirectional coaxial trunk cable, non-directional impedance matching taps with a 20 dB trunk-to-drop loss, and 75 Ω drop cables. The mechanical coupling of the station to the medium shall be to a drop cable as specified in Section 13 by way of a 75 Ω female F-series connector on the station.

When transmitting, the station shall present an impedance resulting in a voltage standing wave ratio (VSWR) of 3:1 or less at that F-connector when terminated with a 75 Ω resistive load or when driven from a 75 Ω resistive source, over the operating frequency range. At all other times, the station shall present an impedance resulting in a VSWR of 1.5:1 or less at that F-connector, when driven from a 75 Ω resistive source, over the operating frequency range.

NOTE: The specifications of this paragraph shall also be met by inactive and unpowered PLEs.

Both the transmitter and the receiver shall be ac-coupled to the center conductor of the drop cable of the 75 Ω single-channel phase-coherent-FSK bus coaxial cable medium, and the breakdown voltage of that ac coupling means shall be at least 500 V ac rms at 50/60 Hz. In addition to this coupling, the shield of the coaxial cable medium shall be connected to chassis ground, and the dc impedance of that connection shall be less than 0.1 Ω.

12.7.6 Receiver Characteristics

12.7.6.1 Receiver Sensitivity and Selectivity. The PLE shall be capable of providing an undetected equivalent bit error rate of 10^{-9} or lower, and a detected bit error rate of 10^{-8} or lower, when

(1) The received signals are transmitted according to 12.7.3.

NOTE: This implies that a preamble following silence may be as short as 22 MAC-symbols.

(2) The received signals are conveyed by the medium defined in Section 13, with the signal characteristics defined in 13.5.3, the maximum noise specified in 13.5.5, and the maximum reflections specified in 13.5.4.

(3) The receiver shall report the presence of received signalling, perhaps as *bad_signal*, and shall not report silence to the MAC sublayer when the received signal level is greater than +10 dBmV.

In summary, the specified signal range is +10 dBmV to +66 dBmV and maximum noise is –10 dBmV in the operating frequency range for both 5 Mb/s and 10 Mb/s receivers.

NOTE: When measuring bit error rate, multiple errors in a frame may be counted as a single error.

12.7.6.2 Minimum Receiver Off Level. The receiver shall report silence within 2 MAC-symbol_times of the receiver going below + 4 dBmV. The intention of this specification is to prevent the PLE from reporting reflections on the medium as non-*silence* to the MAC sublayer entity.

12.7.6.3 Receiver Blanking. The PLE receiver function shall recognize the end of each transmission and report silence to the MAC sublayer entity for a period thereafter. This period of silence, or blanking, shall begin no later than 4 MAC-symbol_times after reporting the last MAC-symbol of the last end delimiter of the transmission. This period of silence shall continue to a point at least 24 but no more than 32 MAC-symbol_times after the reporting of the last MAC-symbol of the end delimiter. Subsequent to this blanking period, silence or non-silence, corresponding to the received signal, shall be reported to the MAC sublayer entity. Figure 12-5 shows the receiver blanking function pictorially.

Therefore the minimum slot time of a network is four octets.

12.7.7 Symbol Timing. During the recovery of *pad_idle* and *silence* symbols, it is permissible to vary the MAC-symbol reporting period by up to one nominal MAC-symbol_time. The period of each reported MAC-symbol shall nonetheless be within 90–210% of the nominal MAC-symbol_time at all times. Further, the period of each reported MAC-symbol after *pad_idle* symbol recovery and until either silence or *bad_signal* is reported shall be within 90–110% of the nominal MAC-symbol_time.

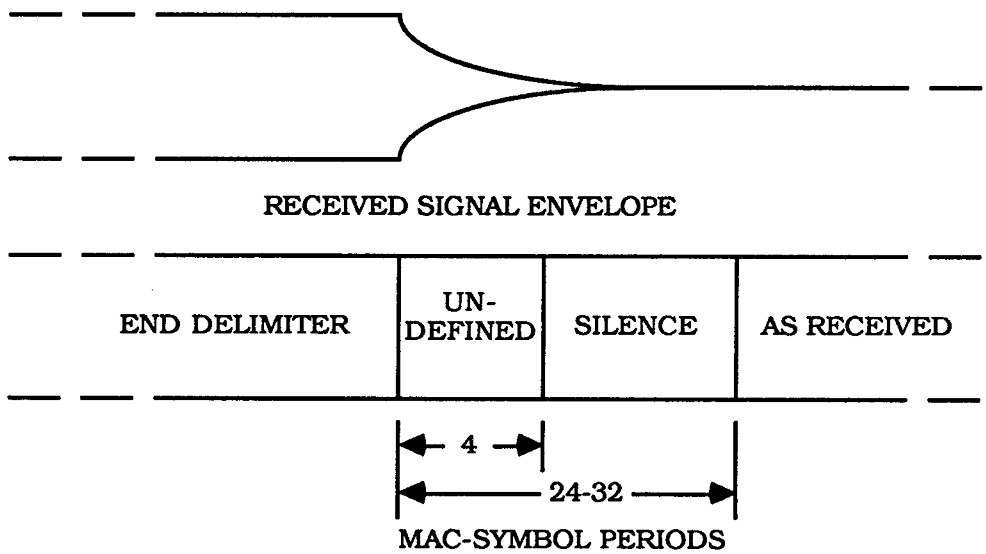

Fig 12-5
Receiver Blanking

12.7.8 Symbol Decoding. After demodulation and determination of each received PHY-symbol, MAC-symbols shall be decoded by the process inverse to that described in 12.7.2, and the decoded MAC-symbols shall be reported at the MAC interface. (As stated in 12.5.1.1, receivers are permitted to decode the transmitted representation of *pad_idle* symbol as an alternating sequence of *ones* and *zeros*.) Whenever a PHY-symbol sequence is received for which the encoding process has no inverse, those PHY-symbols shall be decoded as an appropriate number of *bad_signal* MAC-symbols and reported as such at the MAC interface. In such cases, the receiving PLE should resynchronize the decoding process as rapidly as possible.

The polarity of the carrier signal at the start of the frame does not imply specific information about the signal. The PLE shall be able to receive signals of either initial polarity.

12.7.9 Transmitter Enable/Disable and Received Signal Source Selection (OPTIONAL). The ability to enable and disable transmission onto the single-channel phase-coherent-FSK bus medium as directed by management is recommended but optional.

The ability to select the source of received signaling, either a loopback point within the PLE or one of the possibly redundant media, as directed by management, is recommended but optional. When such an option is invoked and the selected source is not one of the media, transmission to all connected single-channel phase-coherent-FSK bus media shall be disabled automatically while such selection is in force.

12.7.10 Redundant Media Considerations (OPTIONAL). Embodiments of this standard that can function with redundant media are not precluded, provided that the embodiment as *delivered* functions correctly in a nonredundant, single-cable environment. Where redundant media are employed, the provisions of 12.7.4 and 12.7.9 apply separately and independently to each single medium interface, and much of 12.7.9 shall be mandatory. Specifically, separate jabber-inhibit monitoring shall exist for each medium (although common inhibition is permissible), receiver signal source selection shall be provided and it shall be capable of selecting any one of the redundant media, and it shall be possible to enable or disable each single transmitter independently of all other redundant transmitters when the source of received signaling is one of the redundant media.

12.7.11 Reliability. The PLE shall be designed so that its probability of causing a communication failure among other stations connected to the medium is less than 10^{-6} per hour of continuous (or discontinuous) operation. For regenerative repeaters this requirement is relaxed to a probability of 10^{-5} per hour of operation. Connectors and other passive components comprising the means of connecting the station to the coaxial cable medium shall be designed to minimize the probability of total network failure.

12.7.12 Regenerative Repeater Considerations. The PLE of a regenerative repeater can be considered to be a composite entity, with separate electrical and mechanical transmit and receive functions for each connected trunk segment (that is, each port), that all work under a common encoding, decoding, timing-recovery and control function.

In performing its repeating functions, the regenerative repeater serves as a relay station. When a PHY-symbol other than *off* is received, the composite PLE determines which trunk carried the PHY-symbol, and it then selects that trunk as the source of reported PHY-symbols. Concurrently, the MAC entity begins to retransmit onto the other trunks all MAC-symbols reported. When a collision or noise is detected (for example *bad_signal* is reported), the repeater's MAC entity sends an abort sequence (see 4.1.8) in lieu of the received MAC-symbols.

The basic mode of operation, originating or repeating, shall be determined by the superior MAC entity and conveyed by the PHY-MODE invoke primitive (see 8.2.3).

When originating, the repeater PLE shall originate the symbol timing provided to the MAC entity and transmit the encoded MAC-symbols onto all connected trunk segments. The repeater PLE shall use either loopback or any one of the attached trunks as the source of PHY-symbols that are decoded and reported by way of the PHY-DATA indicate primitive.

When repeating and when switching to repeating, the repeater PLE shall delay for the equivalent of the receiver blanking period (see 12.7.6.3) from the end of the most recent transmission to prevent repeating the reflection of the just-prior transmission, and shall then scan the connected ports for one on which signaling is being received. During the delay period, and while this scan for signal is unsuccessful, the repeater PLE shall indicate *silence* symbols to its MAC-entity using its locally originated symbol timing. Upon detecting signaling at one or more ports, the repeater entity shall select one of those active ports as the source of its received signaling. It shall then temporarily disable that selected port's transmitter function, decode the received signaling, retime that signaling before or after decoding, and indicate the decoded MAC-symbols to its associated MAC entity. It should then vary the frequency of the MAC-symbol timing, within the bounds of this subsection (12.7), as necessary to maintain the proper relationship with the frequency of the received PHY-symbol timing.

Regenerative repeaters also prefix enough *pad_idle* symbols to a transmission to provide a minimum of 3 octets of preamble following its receiver blanking period. In summary, when the MAC entity is originating

(1) The PLE alone shall determine the MAC-symbol timing
(2) Transmission occurs on all attached trunks (unless disabled by the provisions of 12.7.9)
(3) Loopback or any one of the attached trunks can be used as the source of PHY-symbols that are decoded and reported by way of the PHY-DATA indicate primitive.

When the MAC entity is repeating

(a) The PLE first delays for the receiver blanking period (i.e., long enough to ensure that the prior transmission is not repeated), then scans all attached trunks for signaling and selects one of those trunks with signaling as the source of received signaling.
(b) Transmission to the selected trunk is temporarily inhibited.
(c) The received signaling from the selected trunk is decoded and indicated to the MAC entity.
(d) The frequency of MAC-symbol timing is varied as necessary (at most, ±0.015%) to track the frequency of the peer transmitter's MAC-symbol timing until the MAC entity requests transmission of *silence*, and then repeats the whole procedure.

12.8 Environmental Specifications

12.8.1 Electromagnetic Emanation. Equipment shall comply with local and national requirements for limitation of electromagnetic interference. Where no local or national requirements exist, equipment shall comply with CISPR Publication 22 (1985) [3].

12.8.2 Safety Requirements. All stations meeting this standard shall comply with relevant local, national, and international safety codes and standards such as IEC Publication 950 (1986) [7].

12.8.3 Electromagnetic Environment. Sources of interference from the environment include electromagnetic fields, electrostatic discharge, transient voltages between earth connections, etc. Several sources of interference contribute to voltage buildup between the coaxial cable and the earth connection, if any, of the station.

The PLE embodiment shall meet its specifications when operating in an ambient plane wave field of

(1) 2 V/m from 10 kHz through 30 MHz.
(2) 5 V/m from 30 MHz through 1 GHz.

12.9 Labeling. It is recommended that each embodiment (and supporting documentation) of a PLE conformant to this standard be labeled in a manner visible to the user with at least these parameters:

(1) Data rate capabilities in Mb/s.

(2) Worst-case, round-trip delay (for nonrepeaters) or one-way delay (for repeaters) that this equipment induces on a two-way transmission exchange between stations, as specified in 6.1.9.

(3) Operating modes and selection capabilities as defined in 12.7.9 and 12.7.10.

Additionally, when the station has multiple connectors (for example, redundant media), the role of each such connector shall be designated clearly by markings on the station in the vicinity of that connector.

13. Single-Channel Phase-Coherent-FSK Bus Medium

The functional, electrical, and mechanical characteristics of one specific form of medium are specified in this section. This specification includes the medium embodiments of a single-channel phase-coherent-FSK bus medium for a Local Area Network (LAN). The relationship of this section to other sections of this standard and to LAN specifications is illustrated in Fig 13-1. The relationship of this section to the single-channel phase-coherent-FSK bus PLE and medium is illustrated in Fig 13-2.

This standard specifies the medium only insofar as necessary to ensure

(1) The interoperability of PLEs conforming to Section 12 when connected to a medium conformant to this section, and

(2) The protection of the LAN medium itself and those using it.

Fig 13-1
Relation to LAN Model

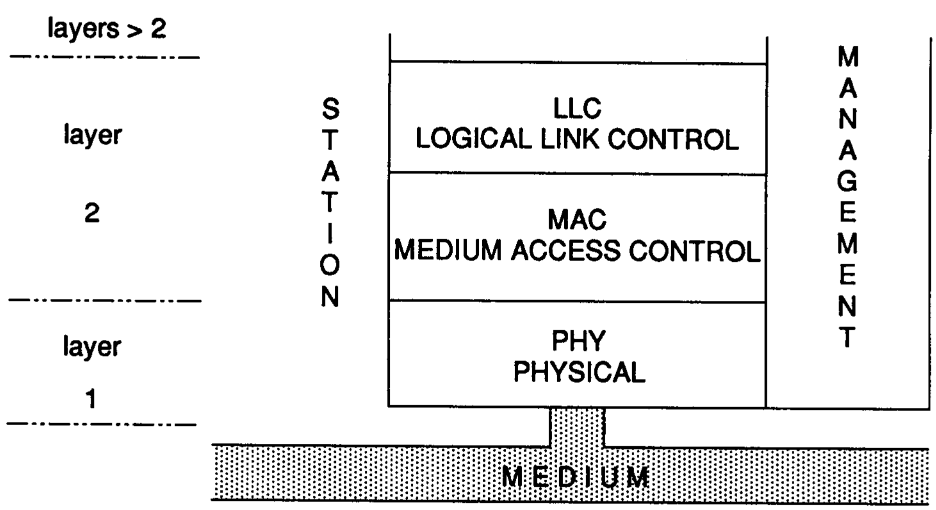

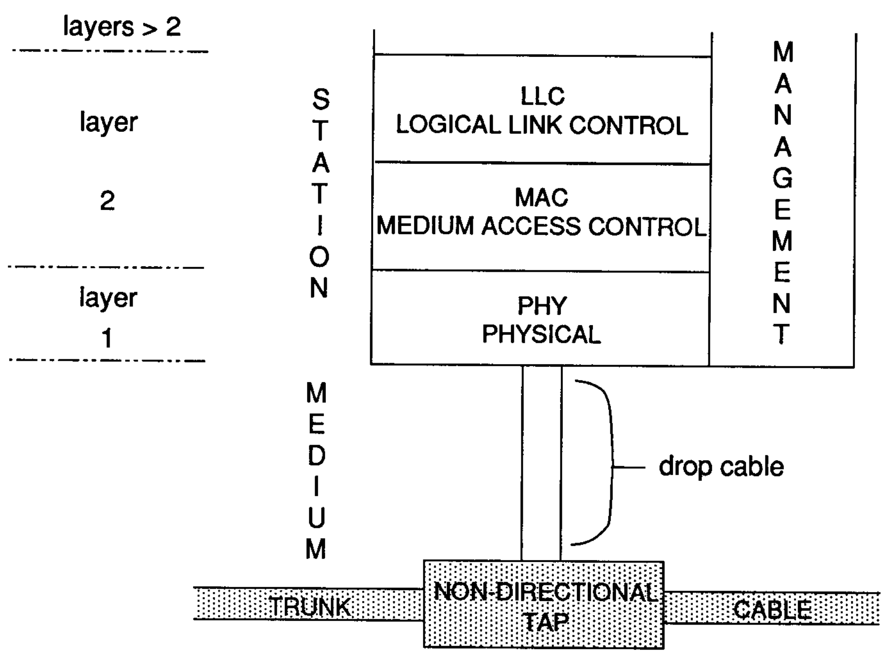

Fig 13-2
Physical Hardware Partitioning

13.1 Nomenclature. Terms used in this section whose meanings within the section are more specific than indicated in the glossary of IEEE P802.1B [10] are as follows:

dBmV. A measure of rms signal level on a 75 Ω cable compared to 1 mV. In terms of SI units, dBmV is defined as dB (1 mV, 75 Ω) rms.

drop cable. The cable of the medium which connects a tap to a station.

F-connector. A 75 Ω F-series coaxial cable connector such as those used for television and video equipment.

frequency shift keying (FSK). A modulation technique whereby information is impressed upon a carrier by shifting the frequency of the transmitted signal to one of a small set of frequencies.

medium. The cable system of a LAN including the trunk cable, taps, drop cables, and splitters.

operating frequency range. The operating frequency range is defined as 2–15 MHz for a 5 Mb/s data rate and 4–30 MHz for a 10 Mb/s data rate.

phase-coherent FSK. A particular form of FSK where the two signaling frequencies are integrally related to the data rate and transitions between the two signaling frequencies are made at zero crossings of the carrier waveform.

regenerative repeater. A device used to extend the length or topology of the medium beyond the limits imposed by signal degradation on a single segment of the medium.

single-channel system. A system where at any point on the medium, only one information signal at a time can be present without disruption.

splitter. A module that electrically and mechanically couples one trunk cable to other trunk cables, providing a branching topology for the single-channel phase-coherent-FSK trunk. A splitter combines signals received at its ports and splits any signal energy received from a trunk symmetrically among the other trunks. It contains only passive electrical components (R, L, C).

tap, nondirectional. A module that electrically and mechanically couples the trunk cable to one or more drop cables. It splits the signal received from each trunk cable very asymmetrically, with the bulk of the signal passed to the other trunk cable and only a small percentage passed to the drop cable(s). A small portion of the signal received by the drop cables from the station is split equally among the trunk cables. The tap contains only passive electrical components (R, L, C).

trunk cable. The main cable of the medium, interconnecting the taps.

13.2 Object. The object of this specification is to

(1) Provide the physical medium necessary for communication between local network stations employing the token-passing bus access method and the single-channel phase-coherent-FSK bus PLE defined in this standard.
(2) Provide for high network availability.
(3) Provide for ease of installation and service in a wide range of environments.
(4) Permit use of components and installation practices similar or identical with those used in the cable TV industry.

13.3 Compatibility Considerations. This standard applies to media that are designed to operate as CATV-like, but nondirectional, single-channel phase-coherent-FSK bus coaxial cable systems. Such systems use nondirectional taps in standard CATV tap housings, and standard CATV connectors and coaxial cable. This specification applies to a single trunk system in which two-way alternate communication is accomplished through the use of nondirectional taps and splitters and, in large systems, multidirectional regenerative repeaters.

All implementations of media conformant to this standard shall be compatible at their drop cable station interfaces. Specific implementations based on this standard may be conceived in different ways provided compatibility at the station interfaces is maintained.

13.4 Overview

13.4.1 General Description of Functions. The functions performed by the single-channel phase-coherent-FSK bus medium entity are described informally here. Jointly, these functions provide a means for signals generated by a station at the medium interface to be conveyed to all of the other stations on the medium. Thus, stations connected to these drop cables can communicate with each other.

13.4.1.1 Operational Overview. Stations are connected to the *trunk* coaxial cable(s) of a single-channel phase-coherent-FSK bus system by *drop* cables and impedance-matching taps. These taps are passive devices that are nondirectional (that is, omnidirectional) with regard to signal propagation. The nondirectional characteristics of the taps permit the signals from a station to propagate in both directions along the trunk cable.

The topology of the single-channel phase-coherent-FSK bus system is that of a highly-branched tree without a root, with the stations connected as leaves to the tree's branches. Branching is accomplished in the trunk itself by way of *regenerative repeaters* (described below) and *splitters*, which provide nondirectional coupling of the signaling carried on the trunk cables similar to that of the just described taps. Like the taps, the splitters also employ only passive electrical components (R, L, C, only).

13.4.1.2 Regenerative Repeater Functions. In a single-channel phase-coherent-FSK bus system, regenerative repeaters may be used to connect trunk segments into a highly branched topology, or to extend the length of the medium, or to increase the number of taps on a trunk. The regenerative repeater is connected to the trunks by way of taps and drop cables. Regenerative repeaters are discussed in detail in 12.7.12.

13.4.2 Basic Characteristics and Options. All signal-conveyance characteristics and station-interface characteristics are mandatory. All other characteristics are optional.

13.5 Functional, Electrical, and Mechanical Specifications. The single-channel phase-coherent-FSK bus medium is an entity whose sole function is signal transport between the stations of a single-channel phase-coherent-FSK bus network. Consequently, only those characteristics of the medium that impinge on station-to-station, signal transport, or on human and equipment safety, are specified in this standard.

An implementation of the medium shall be deemed conformant to this standard if it provides the specified signal transport services and characteristics for the stations of a single-channel phase-coherent-FSK bus medium, and if it meets the relevant safety and environmental codes.

Unless otherwise specified, all measurements in 13.5 are to be made at the point of station or regenerative repeater connection to the medium. Unless otherwise stated, all voltage and power levels specified are in rms and dBmV, respectively, based on measuring the fundamental signal content of continuous transmissions of either all *one* or all *zero* symbols (dBmV is defined as dB [1 mV, 75 Ω] rms in 13.1).

13.5.1 Connection to the Station. The connection of the single-channel phase-coherent-FSK bus medium to the station shall be by way of a 75 Ω drop cable not to exceed 50 m in length. At the station, the drop cable shall have a male F-series 75 Ω connector. This connector shall mate with a female F-series 75 Ω connector mounted on the station.

In addition to this coupling, the shield(s) of the coaxial drop cable medium shall be connected to the shell of the terminating male F-series connector and the dc impedance of that connection shall be less than 0.001 Ω. Also, the dc impedance of a connection between the shell of that male F-series connector and the outer barrel of a mated female F-series connector shall be less than 0.001 Ω.

NOTES: (1) At the time this standard was written, there were no recognized standards for F-connectors. A standard for F-connector mating interface dimensions was being drafted by the IEC as Publication 169: (to be determined), Radio-frequency connectors (Type F). A comprehensive Type FD connector standard (intermateable with Type F) was being drafted by the Electronic Industries Association to be published as EIA 550. F-connectors vary greatly in quality. It is suggested that high quality F-connectors be used. Further, since the center conductor of the cable in some F-connectors becomes the male part of the connector, caution should be used in mixing different types of cables with F-connectors. Where possible, the use of connectors with captive center conductors is recommended.

(2) Under certain environmental conditions F-connectors may be unsuitable. Although not part of this standard, the use of 50 Ω N-connectors will meet the electrical and performance requirements of Sections 12 and 13.

13.5.2 Cable Characteristic

13.5.2.1 Characteristic Impedance. The characteristic impedance of the single-channel phase-coherent-FSK bus medium shall be 75 ± 3 Ω.

13.5.2.2 Trunk Return Loss Requirements. Unidirectional reflected signal energy, that reflected from one direction on the trunk cable, shall be less than –22 dB below the incident energy level at any point on the trunk cable over the operating frequency range. This reflection is the sum of the reflections caused by the structural return loss of the cable itself and all the taps and other components on the trunk cable.

13.5.2.3 Medium Interface Return Loss. The medium as seen at any drop cable to station connection point shall present an impedance resulting in less than –14 dB return loss (VSWR of 1.5:1) when driven from a 75 Ω resistive source as measured over the operating frequency range.

13.5.3 Signal Characteristics

13.5.3.1 Amplitude. When conveying the signal of a single station or regenerative repeater whose transmit level is as specified in 12.7.3.3, the single-channel phase-coherent-FSK bus medium shall present that signal to any connected station or regenerative repeater at an amplitude of between +10 and +66 dBmV for either a 5 Mb/s data rate or a 10 Mb/s data rate system.

13.5.3.2 Tilt. The signal amplitude of the two fundamental signaling frequencies due to media attenuation (corresponding to the data rate and to twice the data rate) at any receiving station shall differ by no more the 3.5 dB.

13.5.3.3 Distortion. The maximum percent phase delay distortion (PDD) shall fall within the mask specified in Fig 13-3, which consists of a straight line mask defined by the break points.

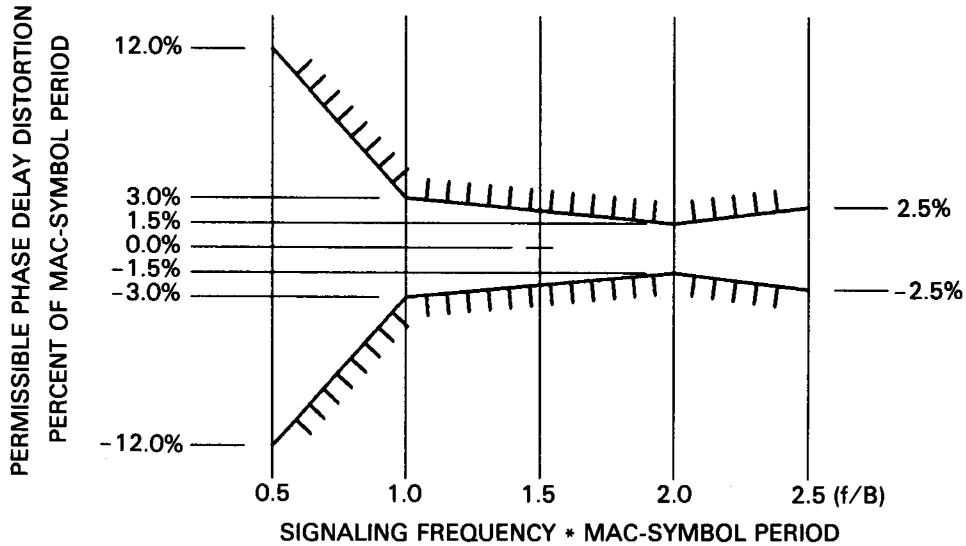

f/B	Maximum PDD
0.5	± 12%
1.0	± 3.0%
2.0	± 1.5%
2.5	± 2.5%

NOTE: The 0% reference is defined at $1.5 f/B$.

where

f = signaling frequency (Hz)
B = MAC-symbol data rate (symbols/second)
PDD = the maximum permissible phase delay distortion as a percent of $1/B$ seconds, which is the MAC-symbol_time

Fig 13-3
Limits for Percent PDD

13.5.4 Received Signal Reflection Component. Excluding noise, the signal received by any station is composed of two parts: the nonreflected component of the signal, and the reflected component. The nonreflected component is that part of the signal that would be received from a perfect medium. The reflected component is due to impedance mismatches in the medium, both on the trunk cable and the drop cable. The transmitting station and the receiving stations are expected to operate as required in Section 12 in the presence of these reflections. As a consequence of the station and medium impedance mismatch characteristics defined in this section and in Section 12, the maximum reflected signal components are as follows:

At a transmitting station, the reflected component of the signal from the medium, during transmission and 1 μs subsequent to the end of a transmission, will be at least 14 dB below the transmit level. From 1–2 μs following a transmission, the reflected component of a signal will be at least 42 dB below the end of transmission transmit level. After 2 μs, the reflection will be below +4 dBmV. These reflection levels are a consequence of the 14 dB return loss permitted for the interface to the medium (13.5.2.3) and the 50 meter maximum length of the drop cable (13.5.1).

If a receiving station shares the same tap as a transmitting station and the isolation between the drop cable ports to the stations is less than 40 dB, then the reflected signal level at the receiving station approaches that of the transmitting station.

13.5.5 Noise Power. The in-band noise power shall be –10 dBmV or less in the operating frequency range as measured at any point a station or regenerative repeater is connected to the medium.

13.5.6 Power Handling Capability. The total power over the entire cable spectrum, as presented to the station or regenerative repeater, shall be less than 0.25 W.

13.5.7 Tap Requirements. The loss between any trunk port and any drop cable port shall be 20 ± 0.5 dB.

For taps with multiple drop cable ports, the isolation between drop ports on the same tap may be between 0 and 41 dB.

13.5.8 Compatibility with the Stations and Regenerative Repeaters. An embodiment of a single-channel phase-coherent-FSK bus medium entity is deemed to support a specific single-channel phase-coherent-FSK bus LAN if the requirements of 13.5.1 through 13.5.7 (inclusive) are met when measured from each point of station connection to the medium, independent of which one of the points of station connection is chosen for test signal origination.

13.5.9 Redundancy Considerations. As stated in 12.7.10, redundant single-channel phase-coherent-FSK bus media are not precluded by this standard. Where redundant media are employed, the provisions of 13.5.1 through 13.5.7 shall apply separately and independently to each medium interface.

13.5.10 Reliability. Connectors and other passive components connecting the station to the coaxial cable medium shall be designed to minimize the probability of total network failure.

13.6 Environmental Specifications

13.6.1 Electromagnetic Emanation. LAN cable systems shall comply with local and national requirements for limitation of electromagnetic interference. Where no local or national requirements exist, equipment shall comply with CISPR Publication 22 (1985) [3].

13.6.2 Safety Requirements. All media meeting this standard shall comply with relevant local, national, and international safety codes and standards.

13.6.3 Electromagnetic Environment. Sources of interference from the environment include electromagnetic fields, electrostatic discharge, transient voltages between earth connections, etc. Several sources of interference contribute to voltage buildup between the coaxial cable and the earth connection, if any, of the station.

The medium embodiment shall meet its specifications when operating in an ambient plane wave field of

(1) 2 V/m from 10 kHz through 30 MHz.
(2) 5 V/m from 30 MHz through 1 GHz.

NOTE: Operating in this environment implies that quad-shield cable or comparable shielding measures may be required.

13.7 Transmission_path_delay Considerations. When specifying an embodiment of a medium that conforms to the specifications of this section, a vendor shall state the transmission_path_delay as the maximum *one-way* delay that the single-channel phase-coherent-FSK bus medium could be expected to induce on a transmission from any connected station through any intervening regenerative repeaters to any other station. The delays induced by the transmitting and receiving stations should not be included in the transmission_path_delay.

For each potentially worst-case path through the medium, a path delay is computed as the sum of the medium-induced delay and repeater-induced delay, if any, in propagating a signal from one station to another. The transmission_path_delay used for determining the network's slot_time (see 6.1.9) shall be the largest of these path delays for the cable system.

These path delay computations shall take into account all circuitry delays in all relevant regenerative repeaters and medium splitters, etc., and all signal propagation delays within the cable segments. The regenerative repeater delays shall be determined as specified in 12.9 and 6.1.9.

The transmission_path_delay shall be expressed in terms of the network's symbol signaling rate on the medium. When not an integral number of signaled symbols, it shall be rounded up to such an integral number. When uncertain of the exact value of the delay, vendors shall state an upper bound for the value.

13.8 Documentation. It is recommended that each vendor of an embodiment of a medium entity conformant to this standard provide to the user supporting documentation with at least the transmission_path_delay, as specified in 13.7 and 6.1.9.

13.9 Network Sizing

13.9.1 Topology Considerations. Very small systems can be constructed using only flexible coaxial cable and passive impedance-matching networks. Larger systems require semirigid trunk cable and flexible drop cables, or regenerative repeaters, or both. Highly-branched topologies are achieved in this medium using CATV-like components.

13.9.2 Signal Loss Budget Considerations. The placement of phase-coherent-FSK regenerative repeaters and taps should take into account

(1) Each station's minimum transmit level and receive level specifications,

(2) The desired current and anticipated future placement of stations and regenerative repeaters, and
(3) The presented signal level and signal-to-noise specifications.

Appendix

(This Appendix is not a part of this International Standard or of IEEE Std 802.4-1990, but is included for information only.)

13.10 Appendix—Guidelines for Configuring the Medium

13.10.1 Network Size. Both overall length of cable and number of taps should be explicitly combined in the system signal loss budget. Each tap is matched at all ports, and the attenuation between all ports is specified. Thus each tap makes a known contribution to the network loss budget, as does each length of the trunk cable. In addition, each tap includes a definite tap loss between the trunk and the drop, which should also be included in the loss budget.

The tap trunk to drop loss value specified in this standard is 20 ± 0.5 dB. Typical values for commercially available two-drop taps conforming to this specification are 0.3 dB insertion loss over a frequency range of 1–30 MHz. An example calculation of the signal loss budget for a simple, unbranched network, including 20 such two-drop taps, might look like this:

Drop loss at first tap:	20.5 dB
Insertion loss of 18 intervening taps:	5.4 dB
Drop loss at 20th tap:	20.5 dB
Total lumped losses:	46.4 dB

The minimum transmit level and receive sensitivity specified in Section 12 are +63 dBmV and +10 dBmV, respectively. Thus, the cable loss in this example network should not exceed 63 – 10 – 46.4 = 6.6 dB. Given the actual intended length of the installed trunk cable, one should select a suitable cable based on attenuation at the highest frequency component for the particular data rate planned or anticipated. For a 5 Mb/s data rate, an RG-11 type cable with an attenuation of 1.5 dB/100 meters at 10 MHz limits the cable length to 440 meters. A longer cable length would be possible if semirigid cable were to be used.

If fewer taps were used, the cable length could be longer. In this case, phase delay distortion rather than cable attenuation might be the limiting factor. A typical phase delay distortion limitation for RG-11 type cable is 700 meters. The cable manufacturer should specify the least cable length at which

(1) The cable attenuation equals 12 dB at a frequency of 10 MHz (for 5 Mb/s systems) or 20 MHz (for 10 Mb/s systems), or
(2) The cable tilt equals the maximum permitted by 13.5.3.2, or

(3) The cable phase delay distortion equals the maximum permitted by 13.5.3.3.

13.10.2 Network Topology. The simplest network topology is a long unbranched trunk, requiring the trunk cable to be routed to each station site in turn. Branched topologies may be implemented by impedance-matched nondirectional splitters, which are three-port passive networks that divide the signal incident at one port into two equal parts that are transmitted to the other two ports. The insertion loss between any two ports of a typical, commercially available, nondirectional splitter is 6.5 dB. When branches are implemented by way of such splitters, a separate loss budget should be calculated for each possible end-to-end path, so that the highest loss path can be used to select the trunk cable.

13.10.3 Repeaters. Regenerative repeaters are specified in Section 12 for the purposes of branching and the extension of the network beyond the basic signal loss budget.

NOTE: Simply cascading 45 taps in an unbranched trunk configuration (at 0.3 dB insertion loss for each of 43 taps, and 20 dB drop loss for each of the other two) consumes the entire dynamic range of 53 dB. Repeaters may be necessary on larger networks even when end-to-end cable lengths are relatively short.

13.10.4 Drop Cables. The drop cables specified in this section are 75 Ω cables that do not exceed 50 m in length, so that loss in the drop cable is less than 1 dB. This length of cable permits relative freedom in routing the trunk cable and locating the station. Since the drop cable length is not negligible, the drop cable should be terminated in its characteristic impedance of 75 Ω to preserve the impedance-matched conditions at the tap. This 75 Ω termination is provided by the stations connected to the drop cable. When no station is connected to a drop cable, or when no drop cable is attached to a drop cable port of a tap, the 75 Ω termination should be provided by other means. See NOTES in 13.5.1.

13.10.5 Trunk Connection. The drop cable is coupled to the trunk cable through the tap, a passive, nondirectional, coupling network that is impedance matched to the 75 Ω cables at all ports. A fixed small fraction of the signal travelling in either direction on the trunk cable is transferred to the drop cable. A signal from a station on the drop cable is attenuated at the tap and then propagates out in both directions on the trunk.

Multi-port taps return a portion of the signal to other drops as well.

Tap insertion losses essentially represent the way the incident power is divided between the ports, they are not due to dissipation in the tap. Typical values, mentioned earlier in connection with network sizing, are 0.3 dB insertion loss between the trunk ports for a two-port tap, and 20 dB drop loss between the trunk and drop ports. In every case, all ports should be matched for proper operation. This means that a 75 Ω termination should be connected to every unused port, and that the cable attached to any port should be properly terminated.

The coupling networks consist only of passive R, L and C elements. Power connections are not permitted.

The coupling networks are enclosed by sealed housings to provide both environmental protection and electrical shielding. The housings, typically metal castings, include integral connectors for the trunk cable and drop cable(s). Housings designed for flexible trunk cables typically use Type F connectors for all ports, while those designed for semirigid trunk cables use Type F connectors for the drop cable, and special connectors that mate directly with properly prepared ends of the trunk cable. If F-connectors are planned, see NOTES in 13.5.1.

13.10.5.1 Tap Placement. It is recommended that the taps and cable splices should not be placed at regular 1/2 wavelength intervals of the higher signaling frequency.

NOTE: It is the ultimate responsibility of the system installer to assure the specifications of this section are met.

13.10.6 Earthing. The trunk cable shield may be floating, single-point, or multiple-point grounded as far as signal transmission is concerned. Earths thus may be installed to eliminate electromagnetic interference (EMI) and comply with safety codes and other regulations applicable to the particular installation. This usually means earthing where the cable enters or leaves a building, and at intervals not exceeding approximately 100 m within the building. Earths should be applied carefully by a clamp that does not crush or damage the cable, because such cable damage causes serious reflections. Suitable clamps are available from suppliers of CATV system hardware.

13.10.7 Surge Protection. It is good practice to protect the cable against ground surges due to lightning. Suitable surge protectors that meet the requirements of IEEE C37.90.1-1989 [8] should be used at each end of the cable. The capacitive loading of surge protectors should be small to avoid affecting PLE (modem) performance, and should not exceed values permitted for a standard tap. For maximum surge protection, a low-impedance, heavy-duty earth connection is required.

13.10.8 Termination. The trunk cable should be terminated in 75 Ω (resistive) at all ends. All drop cables should be properly terminated at the station end. All unused tap ports should be terminated in 75 Ω. Shielded 75 Ω coaxial terminations with good broadband characteristics are commercially available for most coaxial cables. Since maximum transmit levels are +66 dBmV, power ratings of 0.25 W are sufficient.

13.10.9 Joining Cable Sections. In general, the trunk cable will consist of a number of separate sections of coaxial cable. Some sections may be joined by connections to tap housings, while others may be joined by splicing connectors (for semirigid cable) or straight-through connectors (for flexible cable). Flexible cables will be fitted with matched connectors at each end, while semirigid cables will simply have their ends properly prepared to mate with corresponding connectors.

A good engineering practice is to maintain constant impedance between cable sections by using one cable type from one manufacturer for the entire trunk. This practice avoids significant reflections where cables are joined. When dissimilar types are to be joined, it is suggested that a lossy (attenuating) impedance-matched connector be employed to reduce repeated reflections. In order to minimize discontinuities, cable should be obtained from the same lot (same extruder at about the same time of manufacture). It is the extruder-to-extruder differences that contribute to the mismatch. A minimal-loss splice should only be used between two pieces of cable from the same roll, or at worst from the same lot.

13.10.10 Pretested Cable. It is a good practice to pretest all trunk cable before installation. The objective is to ensure that the attenuation does not exceed the expected values at frequencies of interest, and to ensure that concealed (that is, internal) discontinuities that can cause reflections do not exist. For a nominal charge, most cable suppliers will pretest or certify all cable before shipment.

On-site testing after installation is also recommended, since any damage may degrade operating margins or cause outright failure. A recommended method for testing the installed cable for damage, improper termination, shorts, or discontinuities is to use a time domain reflectometer or a return loss meter.

14. Broadband Bus Physical Layer Entity (PLE)

The functional, electrical, and mechanical characteristics of one specific form of PLE (broadband bus) of this standard are specified in this section. This specification includes

(1) The PLEs found in stations that could attach to the broadband bus local area network (LAN),

(2) The PLEs of remodulators that are part of the broadband bus LAN.

The relationship of this section to other sections of this standard and to LAN specifications is illustrated in Fig 14-1. The relationship of this section to the broadband bus PLE and medium is illustrated in Fig 14-2. This standard specifies these PLEs only in so far as necessary to ensure

(1) The interoperability of implementations conforming to this specification,

(2) The protection of the LAN and those using it.

Fig 14-1
Relation to LAN Model

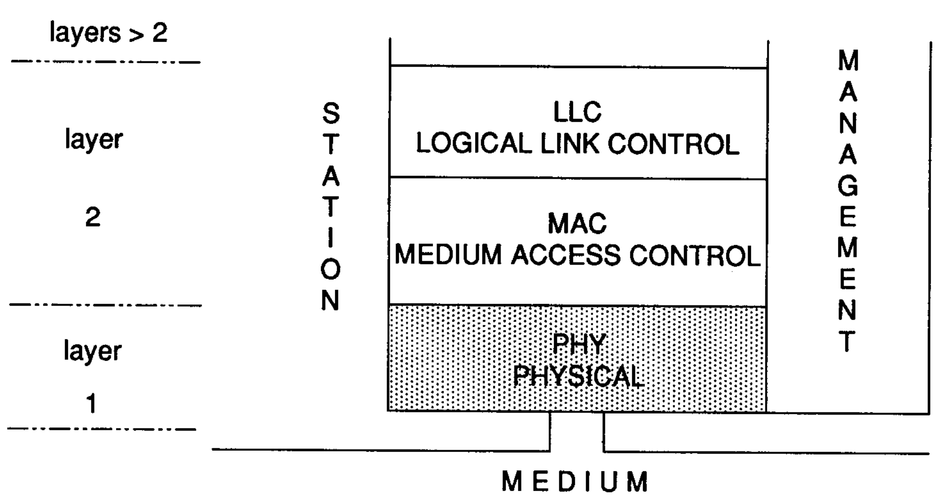

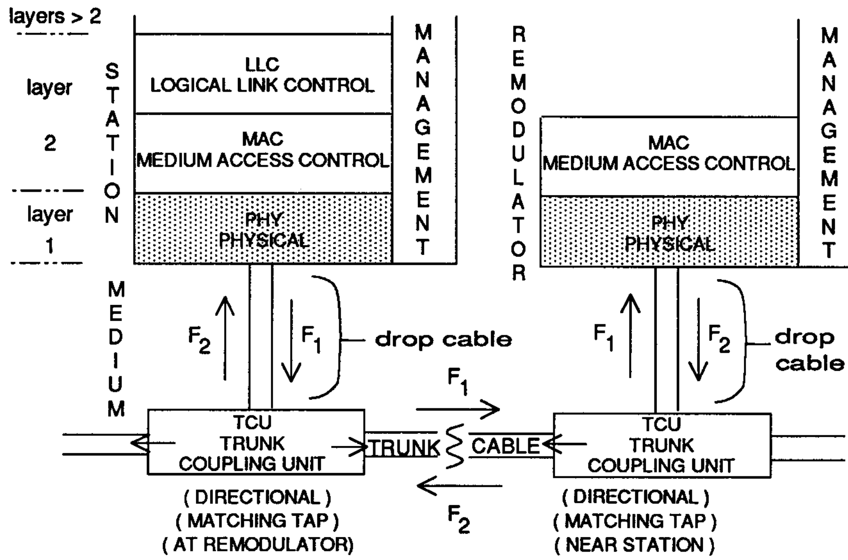

Fig 14-2
Physical Hardware Partitioning

14.1 Nomenclature. Terms used in this section whose meanings, within the section, are more specific than indicated by glossaries referred to in 1.2 are as follows:

AM/PSK modulation. A form of modulation in which an RF carrier is both amplitude modulated (AM) and phase-shifted (PSK, phase shift keyed)

NOTE: On single (bidirectional) cable systems, $F_1 < F_2$ and both F_1 and F_2 are carried on common trunk and drop cables. On dual (unidirectional) cable systems, F_1 and F_2 are carried on separate trunk and drop cables.

broadband coaxial system. A system whereby information is encoded, modulated onto a carrier, and band-pass filtered or otherwise constrained to occupy only a limited frequency spectrum on the coaxial transmission medium. Many information signals can be present on the medium at the same time without disruption provided that they all occupy nonoverlapping frequency regions within the cable system's range of frequency transport.

dBmV. A measure of rms signal on a 75 Ω cable compared to 1 mV. In terms of SI units (International System of Units), dBmV is defined as dB (1 mV, 75 Ω) rms.

detected bit error. An error which is reported as *bad_signal*. *Bad_signal* reported during the preamble or during the four symbols following the last end delimiter of a transmission is not included.

drop cable. The smaller diameter flexible coaxial cable of the broadband medium which connects to a station.

duobinary signaling. A method of representing information to be transmitted by pulses shaped so as to reduce the frequency spectrum required to signal the information. Also known as Class 1 Partial Response coding.

duobinary AM/PSK modulation. As used in this section, a form of modulation in which data is precoded and signaled as duobinary pulses AM/PSK-modulated onto an RF carrier. The particular precoding used is such that receivers can demodulate the modulated signal without having to recover the phase of that signal. In essence, the PSK component of the modulation is used to reduce the RF signal bandwidth, not to carry additional data.

remodulator. As depicted in Fig 14-2, the unit located at the logical root of a broadband bus LAN which demodulates the F_1 frequency signals transmitted by other network stations and rebroadcasts them as F_2 frequency signals back to those stations. In this standard, it transmits continuously and in doing so specifies the precise data rate at which all stations shall transmit.

multilevel duobinary AM/PSK. A form of duobinary AM/PSK modulation which uses more than two distinct amplitude levels (in this specification, independent of phase) to represent information. This section specifies a three-level duobinary AM/PSK system capable of signaling at one Medium Access Control (MAC)-symbol/PHY-symbol; Appendix 14.11 discusses an extension capable of signaling at two MAC-symbols/PHY-symbol.

trunk cable. The main (larger diameter) cable of a broadband coaxial cable system.

undetected bit error. An error that is not reported as *bad_signal* by the PLE.

14.2 Object. The object of this specification is to

(1) Provide the physical means necessary for communication between local network stations employing the token-passing bus access method defined in this standard and a broadband bus medium,
(2) Define a physical interface that can be implemented independently among different manufacturers of hardware, and achieve the intended level of compatibility when interconnected to a common broadband bus local area network medium,
(3) Provide a communication channel capable of high bandwidth and low bit error rate,
(4) Provide for ease of installation and service in a wide range of environments,
(5) Provide for high network availability,
(6) Use a medium that can be shared by totally unrelated applications (such as voice, data, and analog video), and conserve the capacity of that medium where reasonable.

14.3 Compatibility Considerations. This standard applies to PLEs that are designed to operate on conventional bidirectional (split frequency) community antenna television (CATV)-like broadband coaxial cable systems, or on similar unidirectional (split cable) dual-cable systems, or both. Such networks use standard CATV taps, connectors, amplifiers, power supplies, and coaxial cable.

The use of a coaxial broadband system permits the assignment of different frequency bands to multiple simultaneous applications. For example, a portion of the cable spectrum can be used by LANs while other portions are used for point-to-point or multipoint data links, and others are used to convey television and audio signals. With the proper selection of signal levels and proper equipment design, all of the applications can be supported simultaneously.

Compatibility with this medium is determined at the medium interface. Compatible PLEs shall use the same signaling rate, the same frequency allocations, and the same number of cables.

NOTE: CATV-like media typically provide separate coupling points for the remodulator transmit and receive functions due to the location of this station in the topology. A remodulator configured in this way is compatible with other stations having single connections to the medium.

14.4 Operational Overview of Single-Cable Medium. The use of a coaxial cable operating as a broadband medium introduces special system considerations that are not encountered in single-channel transmission systems. A conventional CATV-like system uses bidirectional amplifiers to achieve two-way transmission on a single coaxial cable. The amplifier actually consists of two independent unidirectional amplifiers, together with the appropriate crossover filters, so that one amplifier transmits the upper portion of the cable spectrum in the *forward* direction, outbound from the remodulator, while the other amplifier transmits the lower portion in the *reverse* direction, toward the remodulator. Therefore the PLE requires an RF transmitter and receiver each operating in a different portion of the cable spectrum.

At the remodulator, where forward direction signals originate and reverse direction signals converge, a remodulator receives the lower-frequency reverse channel signal and retransmits it as the higher-frequency forward channel signal to the “downstream” stations.

Stations are connected to the *trunk* coaxial cable in such systems by smaller-diameter *drop* cables and impedance-matching *taps*. These taps are passive devices that are highly directional with regard to signal propagation, as shown in Fig 14-2. The directional characteristics of the taps improve the impedance matching at the *ports* of the tap and minimize the effects of reflections due to any impedance mismatches along the trunk or on other drop cables.

The directionality of the tap is such that the transmission loss between the station and the reverse direction trunk. Hence, communication between devices connected to different taps along the coax occurs through the transmission of a frequency F_1 in the low-band toward the remodulator, reception at the remodulator and retransmission in the forward direction at a high-band frequency F_2, and reception by the “listening” stations.

14.5 Operational Overview of a Dual-Cable Medium. Dual-cable systems are similar to single-cable systems except that bidirectionality is provided by separate cables rather than by separate portions of the frequency spectrum on the same cable. Consequently, amplifiers need be only unidirectional and, with respect to Fig 14-2, F_1 need not be a lower frequency than F_2.

14.6 General Overview

14.6.1 General Description of Functions. The functions performed by the broadband bus PLE are described informally here. Some of these functions are common to all stations of a broadband bus LAN; others are found only in the remodulator station. Jointly they provide a means whereby symbols presented at the MAC interface of one PLE can be conveyed to all of the PLEs on the bus for presentation to their respective MAC interfaces.

14.6.1.1 Symbol Transmission and Reception Functions. Successive MAC-symbols presented to the PLEs at its MAC service interface are applied to an encoder that produces as output PHY-symbols from the set: {*0*} {*2*} {*4*}. The encoder includes a scrambler that is applied to consecutive *zero* and *one* input MAC-symbols to reduce the autocorrelation of the encoder's output. The output is then applied to a modulator that uses a duobinary signal shaping process to produce a multilevel AM/PSK signal within the designated RF channel, where the amplitude of each output pulse corresponds directly with the relative numeric value of the associated PHY-symbol. This RF signal is then ac-coupled to the broadband bus medium and conveyed by the medium to one or more receivers.

NOTE: Except for the remodulator, sequences of the MAC-symbol *silence* are conveyed by turning off the transmitter.

Each receiver is also ac-coupled to the broadband bus. It bandpass filters the received signal to eliminate noise from other channels and then infers the transmitted PHY-symbol from the amplitude of the received signal. It decodes the inferred PHY-symbols by an approximate inverse of the encoding process and presents the resultant decoded MAC-symbols at its MAC service interface.

For all MAC-symbols except *pad_idle*, this decoding process is an exact inverse of the encoding process in the absence of errors. The *pad_idle* symbols, which are referred to collectively as *preamble*, are transmitted at the start of each MAC frame, both to provide a training signal for receivers and to provide a nonzero minimum separation between consecutive frames. Since each transmission begins with *pad_idle* symbols, it is expected that some of these initial symbols may be "lost in transit" between the transmitter and the remodulator. Additionally, in broadband bus systems, the encoding for the sequence of *pad_idle* MAC-symbols is identical to the encoding for a particular sequence of *one* and *zero* MAC-symbols, and receivers are permitted to decode the transmitted representation of *pad_idle* as a *one* or *zero* MAC-symbol, descramble it if desired, and report it as such to the MAC entity.

As stated above, the remodulator actively transmits sequences of the MAC-symbol *silence* as a specific repeating signaling sequence known as pseudo-silence. This sequence is used to provide a continuous signal for the automatic gain control (AGC) and carrier and data recovery circuitry of the LAN's other stations, and to indicate the signaling type (one MAC-symbol/PHY-symbol, or the extended type defined in 14.11) currently in use on the LAN. When received by the other stations in the network, pseudo-silence is decoded and reported to their local MAC entities as the MAC-symbol *silence*.

14.6.1.2 Remodulator Functions. In an actual single-cable broadband bus system, stations do not transmit and receive on the same frequency channel. One central station, known as a remodulator, receives on a low-frequency channel and transmits on a high-frequency channel. All of the other stations transmit on the low-frequency channel and receive on the high-frequency channel. Thus, the system really consists of a pair of directional channels, with the low-frequency (reverse) channel having many transmitters and one receiver, and the high-frequency (forward) channel having one transmitter and many receivers.

NOTE: The remodulator may also function like any other station and source transmissions of its own.

To allow stations to communicate, the remodulator (see 14.8.18) serves as a relay station; when not transmitting for itself, its MAC sublayer entity interprets the symbols received in the reverse channel and retransmits them (in some form) in the forward channel. When a collision or noise is detected (for example, *bad_signal* reported), the remodulator's medium access entity sends an abort sequence (see 4.1.8) in lieu of the received symbols. When *silence* is detected, the remodulator sends a pseudo-silence pattern which the other stations report to their medium access entities as *silence*. Thus, during normal operation the remodulator transmits continuously. The other stations in the LAN determine the exact transmit and receive data rate from that continuously on, forward-channel signal and the type of signaling (one MAC-symbol/PHY-symbol, or the extended type as defined in 14.11) from the transmitted pseudo-silence pattern.

In a dual-cable system, the forward channel is carried by the forward-direction cable, the reverse channel by the reverse-direction cable, and the reverse channel frequency does not need to be less than the forward channel frequency. In other respects, dual-cable operation is similar to single-cable operation.

14.6.1.3 Jabber-Inhibit Function. To protect the LAN from most faults in a station, each station other than the remodulator contains a jabber-inhibit function. This function serves as a "watchdog" on the transmitter; if the station does not turn off its transmitter after a prolonged time (roughly one-half second), then the transmitter output shall be disabled automatically for at least the remainder of the transmission.

14.6.1.4 Local Administrative Functions (OPTIONAL). These functions are activated either manually, or by way of the PLE's management interface, or both. They can include

(1) Selecting the transmit and receive channels (frequencies and bandwidth).
(2) Enabling or disabling each transmitter output. (A redundant media configuration would have two or more transmitter outputs.)
(3) Adjusting the transmitted power level of each transmitter output.
(4) Selecting the received signal source: any medium (if redundant media are present), or any available loopback point.

NOTE: If a loopback point is selected, then all transmitter output to the medium shall be inhibited.

(5) Reporting the level of the received RF signal.
(6) Selecting the signaling type (one MAC-symbol/PHY-symbol or the extended type defined in 14.11) at the remodulator.

14.6.2 Basic Functions and Options. Symbol transmission and reception functions and jabber-inhibit functions are required in all implementations. Within the symbol transmission and reception functions, three-symbol signaling (one MAC-symbol/PHY-symbol) capability is mandatory. All other functions are optional.

14.7 Application of Management. The following constraints are imposed on the parameters and actions specified in Section 9:

(1) In the capabilities group the dataRates parameter sequence shall specify one or more of the values 1, 5, and 10.
(2) In the relativePowerLevel type, the values *tooLow* and *tooHigh* shall be used to indicate signal levels outside the appropriate range as specified in Table 14-1, whereas the values *low* and *high* shall be used to indicate signal levels within that range but deviating by more than 1 dB from a value 7 dB less than the maximum level specified in Table 14-6. The accuracy of this assessment need be no greater than ±2 dB.
(3) In the local management information, the value of minPostSilencePreambleLength shall be 4.

14.8 Functional, Electrical, and Mechanical Specifications. Unless otherwise stated, all voltage and power level specifications are in rms and dBmV, respectively, based on transmissions of random data patterns (dBmV is defined as dB (1 mV, 75 Ω) rms in 14.1). Measurements for voltage levels and power ratios may be made using a data pattern of unscrambled ones, with a correction factor of +3 dB over equivalent measurements using random data.

14.8.1 Data Signaling Rates. The standard data signaling rates are 1 Mb/s, 5 Mb/s, and 10 Mb/s. The permitted tolerance for each signaling rate is ±0.005%. The exact rate in any given LAN shall be determined by that network's remodulator.

14.8.2 Symbol Encoding. The PLE transmits symbols presented to it at its MAC interface by the MAC sublayer entity. The possible MAC-symbols are *zero*, *one*, *non_data*, *pad_idle*, and *silence*. Each of these MAC-symbols is encoded into a different code of PHY-symbols {*0*} {*2*} {*4*} and then transmitted.

Table 14-1
Relative Power Level Reporting

Channel Bandwidth (MHz)	tooLow	low	OK	high	tooHigh
1.5	<–13	–13 to –4	–4 to –2	–2 to 4	>4
6.0	<–7	–7 to 2	2 to 4	4 to 10	>10
12.0	<–4	–4 to 5	5 to 7	7 to 13	>13

NOTE: All values in dBmV.

14.8.2.1 Three-Symbol Encoder. The encoding action to be taken for each of the input MAC-symbols when three-symbol (one MAC-symbol/PHY-symbol) encoding is employed is

(1) ***Silence—***

(a) The remodulator shall encode successive *silence* symbols as successive PHY-symbols of the repeating sequence {*2*} {*2*} {*0*} {*4*}, always restarting with the first {*2*} of the sequence for each new period of transmitted pseudo-silence. (It is permissible to terminate the sequence after any PHY-symbol.)

(b) All stations other than the remodulator shall encode the first MAC-symbol of *silence* after non-*silence*, and the last MAC-symbol of *silence* before non-*silence*, as the PHY-symbol {*2*}. Such stations shall encode all other MAC-symbols of *silence* as the PHY-symbol {*0*} and shall cease transmission while transmitting only such {*0*} symbols representing *silence*.

(2) ***Pad_idle—***Consecutive *pad_idle* symbols shall be encoded as successive PHY-symbols of the repeating sequence {*4*}{*0*}, always restarting the sequence with the {*4*} for each new period of *pad_idle*. When sending *pad_idle* MAC-symbols immediately after sending *silence* MAC-symbols, the MAC sublayer entity shall send at least as many octets of *pad_idle* MAC-symbols as are required to result in a minimum of 32 PHY-symbols.

NOTE: For this basic signaling, this minimum is four octets of *pad_idle* MAC-symbols.

(3) ***Non_data—***The scrambler described in 14.8.2.3 shall be reinitialized (preset to all ones) at each *non_data* symbol. Each *non_data* symbol shall be encoded as a {*2*}, and the eight-symbol (frame delimiter) sequence of 8.2.1 shall be encoded as follows: each

$non\_data\ non\_data\ data_1\ non\_data\ non\_data\ data_2\ data_3\ data_4$

sequence shall be encoded as the sequence

$\{2\}\ \{2\}\ \{d_1\}\ \{2\}\ \{2\}\ \{d_2\}\ \{d_3\}\ \{d_4\}$,

where each $data_i$ is either a *zero* or a *one*, and each d_i is the corresponding encoding of the $data_i$, either a {*0*} or a {*4*} as described below.

After each such eight-symbol (frame delimiter) sequence, as defined in 8.2.1, the encoder shall start to maintain a sense of PHY-octet alignment, with the next output PHY-symbol considered to be the first of a new PHY-octet.

NOTE: This PLE sense of PHY-octet alignment is identical to the MAC entity's sense of MAC-octet alignment.

(4) ***Zero and one—***Each *zero* and *one* symbol shall be encoded as follows:

(a) The MAC-symbol shall be applied as input to the scrambler described in 14.8.2.3, resulting in a "scrambled" *zero* or *one* MAC-symbol;

(b) The scrambled MAC-symbol shall be converted to one of the PHY-symbol output codes as follows:

zero → {*0*}
one → {*4*}

This is the $data_i$ to $\{d_i\}$ mapping referred to in 14.8.2.1(3).

(c) The following procedure shall be performed on each resultant PHY-octet of {*0*} and {*4*} PHY-symbol codes. Each such octet of these PHY-symbol codes shall be compared with the immediately prior octet of PHY-symbols as *transmitted* (that is, after execution of this procedure on that prior octet, if applicable). If both PHY-octets consist solely of the identical PHY-symbol code (that is, either sixteen {*0*}s or sixteen {*4*}s), then the last three codes in the PHY-octet under examination shall be altered before transmission as follows:

{*0*} {*0*} {*0*} shall be replaced by {*4*} {*2*} {*2*}
{*4*} {*4*} {*4*} shall be replaced by {*0*} {*2*} {*2*}

NOTES: (1) This procedure can alter at most every other octet of PHY-symbols, and the second octet after the encoded (frame delimiter) sequence of 8.2.1 is the first such octet that can be altered. This procedure has no effect on the scrambler described in 14.8.2.3 and in the previous subsection.

(2) The constraints on the occurrence of *non_data* and *pad_idle* MAC-symbols imposed by 8.2.1 and 8.2.2, and the resulting constraints on the occurrence of {*2*} PHY-symbols imposed by 14.8.2.1, may be used in post-detection processing to decrease the error rate of the receiving PLE's PHY-symbol to MAC-symbol decoding process.

14.8.2.2 Pseudo-Silence Provision for the Enhanced Encoder. One additional pseudo-silence sequence (other than that of 14.8.2.1) has been reserved for a two MAC-symbol/PHY-symbol enhancement to the basic signaling method. The enhancement itself is described in 14.11, and is for further study.

The pseudo-silence sequence of 14.8.2.1 is

{*2*} {*2*} {*0*} {*4*}—one MAC-symbol/PHY-symbol three-symbol signaling.

The reserved pseudo-silence sequence is

{*2*} {*2*} {*4*} {*0*}—two MAC-symbol/PHY-symbol three-symbol quadrature signaling.

14.8.2.3 Data Scrambler. In order to reduce the probability of long transmitted sequences of identical PHY-symbol codes and to randomize spectral components of the transmitted modulation, the symbolic binary data transmitted as *zero* and *one* MAC-symbols for the MAC entity shall be pseudo-randomized by dividing the equivalent message polynomial by the generating polynomial $1+X^{-6}+X^{-7}$. The coefficients of the quotient of this division are taken in descending order as the MAC-symbol sequence to be coded in 14.8.2.1 (4). The logical arrangement of such a scrambler and corresponding descrambler is described in 14.12.

As specified in 14.8.2.1, this scrambler (or the corresponding descrambler) is reinitialized at each *non_data* symbol requested by (or indicated to, respectively) the associated MAC entity. After reinitialization, which is presetting to all *ones*, the scrambler polynomial guarantees that the next six MAC-symbols output will be unaffected by the scrambling process. Thus, the *zeros* and *ones* within the frame delimiter octets of 8.2.1 are unaffected by the scrambling process, even though, in a technical sense, they have been scrambled.

The alteration in the transmitted PHY-symbol codes described in the last paragraphs of 14.8.2.1 has no effect on this scrambler or the corresponding descrambler, as the alteration is made subsequent to scrambling and is reversed prior to descrambling.

NOTE: The scrambler function shall only be applied to data (*one* and *zero*) symbols. *Pad_idle*, *silence*, and *non_data* symbols shall not be scrambled.

14.8.3 Baseband Modulation. The PHY-symbol coded signal shall be converted to an appropriately shaped pulse whose phase and amplitude are chosen from the set of pulses of phase and relative amplitude *–2*, *0*, and *+2* as determined by the output PHY-symbol code and by the phase and amplitude of the immediately prior one or two pulses, as specified by Table 14-2.

The equivalent baseband signal shaping function shall be based upon the partial response pulse often referred to as duobinary or Class I partial response, of which time and spectral functions are defined by Eq 14-1, Eq 14-2, and Eq 14-3.

This shaping process shall be effected in such a way that the decoding can be achieved by full wave rectification of the demodulated line signal (that is, without recovery or reconstruction of the phase of the transmitted equivalent baseband signal).

NOTE: The pulse-coding rules of Table 14-2, when combined with the above duobinary signal-shaping function, give rise to a combined signal whose amplitude, when sampled at the correct intervals, corresponds directly to the coded PHY-symbol as follows (where MAX is the maximum such amplitude):

{*0*} → 0/4 of MAX (= 0)
{*2*} → 2/4 of MAX (= MAX/2)
{*4*} → 4/4 of MAX (= MAX)

Table 14-2
Baseband Pulse Coding Rules

2nd Prior Pulse	Prior Pulse	Desired PHY-Symbol	Required Pulse	MAC-Symbol Usage (post-scrambler)
—	*+2*	{*0*}	*–2*	*zero* or
—	*–2*	{*0*}	*+2*	second *pad_idle* (of pair)
—	*0*	{*0*}	*0*	other *silence* or (transmitter off)
—	*+2*	{*2*}	*0*	first *non_data* (of pair)
—	*–2*	{*2*}	*0*	or first *silence*
+2	*0*	{*2*}	*+2*	second *non_data*
–2	*0*	{*2*}	*–2*	(of pair)
—	*0*	{*2*}	*+2*	last *silence*
—	*0*	{*2*}	*–2*	
—	*+2*	{*4*}	*+2*	*one* or first
—	*–2*	{*4*}	*–2*	*pad_idle* (of pair)

$$g(t) = \frac{4}{\pi} \cdot \frac{\cos\left(\pi \frac{t}{T}\right)}{1 - \left(2\frac{t}{T}\right)^2} \qquad \text{(Eq 14-1)}$$

$$\left| G(f) \right| = \begin{cases} 2T\cos(\pi\ f\ T), & |f| \le \frac{1}{2T} \\ 0, & |f| > \frac{1}{2T} \end{cases} \qquad \text{(Eq 14-2)}$$

$$\angle G(f) = \begin{cases} -\pi\ f\ T, & |f| \le \frac{1}{2T} \\ 0, & |f| > \frac{1}{2T} \end{cases} \qquad \text{(Eq 14-3)}$$

respectively, where
f = denotes the frequency (Hz)

$\frac{1}{T}$ = denotes the PHY-symbol signaling rate (symbols/second)

The reference to equivalent baseband signals recognizes that the Physical Layer implementation may be such that the MAC-symbols encoded as in 14.8.2 are converted to and from the line signals without appearing as actual baseband signals.

14.8.4 Recommended Frequency Allocation/Channel Spacing

(1) This standard recommends a communication medium configuration that is known in the CATV industry as a single-cable, mid-split configuration. In this configuration, the unidirectional forward and reverse channels are paired to provide bidirectional channels. This standard recommends a frequency offset of 192.25 MHz between paired forward and reverse channels in both mid-split and high-split configurations. Other configurations, including dual-cable where forward and reverse channels are on separate unidirectional cables, also are permitted.

(2) The relationship between data signaling rate and minimum channel band width is

A data rate of 1 Mb/s requires 1.5 MHz of channel bandwidth. A data rate of 5 Mb/s requires 6 MHz of channel bandwidth. A data rate of 10 Mb/s requires 12 MHz of channel bandwidth. Higher data rates for a given channel bandwidth are for future study.

NOTE: The remainder of 14.8.4 is not part of this International Standard. Frequency allocations are a subject for national standardization.

(3) Table 14-3 shows the preferred pairing for the usual North American 6 MHz channels. The conventional North American CATV nomenclature has been used for the channel names.

The 1.5 MHz channels are formed by equally subdividing any given 6 MHz channel. The 12 MHz channels are composed of adjacent 6 MHz channels, paired as indicated by the lines in the center column of Table 14-3.

The preferred channel assignments for dual-frequency, mid-split and high-split North American system configurations are

(a) For a 10 Mb/s data rate, a 12 MHz channel pair should be used, with reverse channels 3' and 4' paired with forward channels P and Q.

(b) For a 5 Mb/s data rate, a 6 MHz channel pair should be used, with reverse channel 3' paired with forward channel P, or reverse channel 4' paired with forward channel Q.

(c) For a 1 Mb/s data rate, a 1.5 MHz channel pair should be used, chosen from one of the eight equally spaced 1.5 MHz subchannels of reverse channels 3' and 4', paired with corresponding forward channels P and Q.

For dual-cable, single-frequency systems, either the forward channel frequencies or the reverse channel frequencies specified can be used on both forward-direction and reverse-direction cables.

Where multiple channel assignments are needed for multiple LANs on a single shared broadband medium, it is suggested that they be assigned different channel pairs in accordance with these recommendations.

Table 14-3
Usual North American Mid-split Channels—Nomenclature and Pairing

Reverse Channel	Frequency* (MHz)		Forward Channel	Frequency* (MHz)
T10	23.75		J	216
T11	29.75		K	222
T12	35.75		L	228
T13	41.75		M	234
T14	47.75		N	240
2'	53.75		O	246
3'	59.75	†	P	252
4'	65.75	†	Q	258
4A'	71.75	†	R	264
5'	77.75	†	S	270
6'	83.75	†	T	276
FM1'	89.75	†	U	282
FM2'	95.75		U	282
FM3'	101.75		W	294

* The frequency specified is that of the lower band-edge.
' The primed reverse-direction channels are offset from the conventional forward-direction channels of the same (unprimed) name.
† These channel pairings are recommended for use with this standard.

14.8.5 Line Signal in RF Band (at the line output of the station). The equivalent baseband modulation of 14.8.3 shall be applied to an RF carrier centered in the assigned channel (for example, 3 MHz from each channel edge of a 6 MHz channel), and the resultant modulated carrier shall be coupled to the broadband bus medium as specified in 14.8.7.

14.8.5.1 In the appropriate 1.5 MHz wide, 6 MHz wide, or 12 MHz wide RF channel, the line signal shall correspond to a double-sideband signal with its carrier frequency at the channel mid-point frequency ±0.01%.

14.8.5.2 The relationship between the multilevel signals at the real or hypothetical output of the multisymbol signal coder and the transmitted line signal shall be such that the amplitude of the modulated signal at the proper sampling intervals is directly proportional to the numeric symbol of the coded signal.

In a practical case this means that the voltage that results from the full wave rectification of the demodulated line signal corresponds at the proper sampling intervals with the numeric value of the coded signal symbol, with the maximum amplitude corresponding to a coded {*4*}, a zero amplitude corresponding to a coded {*0*}, and intermediate amplitudes corresponding proportionately.

14.8.5.3 The amplitude of the theoretical line signal pulse spectrum, corresponding to an isolated maximum-amplitude pulse appearing at the output of the multisymbol signal coder, is to be cosine-shaped in the frequency domain, as indicated in Eq 14-2, with maxima at the carrier frequency and zeros at ±0.5 MHz, ±2.5 MHz, or ±5.0 MHz from that carrier frequency for 1.5 MHz, 6 MHz, and 12 MHz channel bandwidths, respectively.

14.8.5.4 The output level of the transmitted signal into a 75 Ω resistive load shall be set at the value specified in Table 14-4. Once set, the permissible power variance from this level due to all factors shall be at most ±2 dB. The maximum variance of the transmitter amplitude for any coded symbol, measured over one transmission, shall be at most 0.5 dB. Embodiments with adjustable transmitter output power shall be able to provide these power and stability figures within the range of such adjustment.

14.8.5.5 In the 1 MHz, 5 MHz, or 10 MHz signaling band centered at the actual RF carrier frequency, the amplitude distortion of the real spectrum relative to the theoretic spectrum as defined under 14.8.3 shall fall within the limits specified by Table 14-5 and Fig 14-3, and the group-delay distortion relative to

Table 14-4
Required Transmit Level

Channel Bandwidth (MHz)	Required Transmit Level (dBmV)
1.5	+41
6.0	+47
12.0	+50

Table 14-5
Breakpoints of Amplitude Distortion Limits

Distance from Band Center (% of HBW)	Permissible Distortion from Theoretic Spectrum (in frequency domain (dB)
± 20%	±0.5
± 50%	±1.0
± 80%	±2.0
± 90%	±3.0
±100%	(at least 20 dB attenuation)

where
HBW = 1/2 signaling bandwidth, that is, 0.5, 2.5, or 5 MHz

$$\mathrm{PGD} = \pm 90 \log_{10} \quad \cos(\pi f T),\ f \le \frac{1}{2T} \qquad \text{(Eq 14-4)}$$

where
PGD denotes the permissible group delay distortion relative to the center of the signaling band, measured as a percent of T
f denotes the frequency (Hz)
$1/T$ denotes the PHY-symbol signaling rate (symbols/second)

$$\mathrm{RA} = \min\left(55, 30 + 10\log_{10} B + 80\frac{\text{MF-NCEF}}{B}\right) \qquad \text{(Eq 14-5)}$$

where
RA is the attenuation relative to total transmitted power in dB measured in any 30 kHz band
B is the channel bandwidth (MHz)
MF is the measurement frequency (MHz)
NCEF is the frequency (MHz) of the nearest edge of the channel

the center of the band shall fall within the limits specified by Eq 14-4. At the signaling band edges, where the theoretic spectrum has zeros (infinite attenuation), the actual signal amplitude shall be attenuated at least 20 dB with respect to the signal amplitude at the center of the band.

14.8.5.6 The transmitted power outside the channel bandwidth shall meet or exceed the attenuation specified in Fig 14-4 and Eq 14-5 relative to total transmitted power independent of the data transmitted.

14.8.5.7 When a transmitter ceases generating and sending the equivalent baseband pulses described in 14.8.3, either through receiving the MAC-symbol *silence*, or through a management command to cease transmitting, it shall take the following actions:

(1) It shall send the PHY-symbol {*2*} as the first symbol of the "silent" period (see 14.8.3),
(2) It shall continue to transmit the decaying energy of all previously generated equivalent baseband pulses until the power represented by the sum

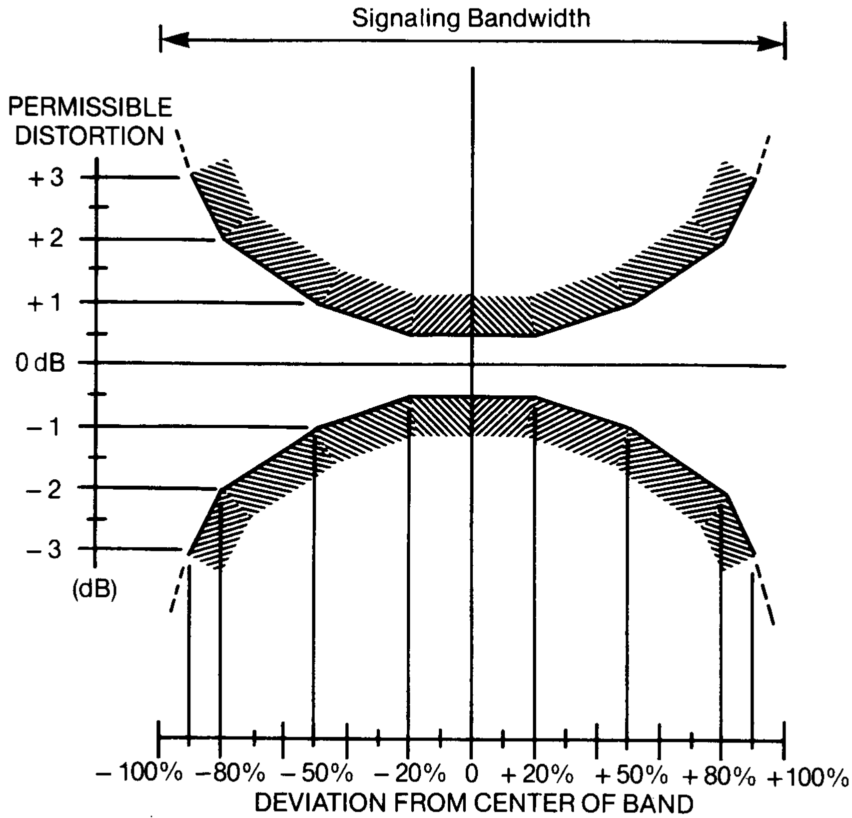

Fig 14-3
Limits for Amplitude Distortion

of these pulses, averaged over one symbol time, is less than a level 26 dB below the power level set by 14.8.5.4,

(3) After this time, but before 10 symbol times have elapsed from the last transmitted pulse, the total transmitted power, in the transmit channel bandwidth, shall be less than a level 70 dB below the set power level.

When a transmitter begins sending PHY-symbols requiring the generation of pulses, it shall take the following actions:

(a) It shall remain at a total power output less than a level 70 dB below the set transmit power until not more than 10 symbol times before the first transmitted pulse,
(b) It shall begin transmitting the energy contained in the pulses before the power represented by the sum of these pulses, averaged over one symbol time, has reached a level 26 dB below the set power level,
(c) It shall send the PHY-symbol {*2*} as the last symbol of the "silent" period.

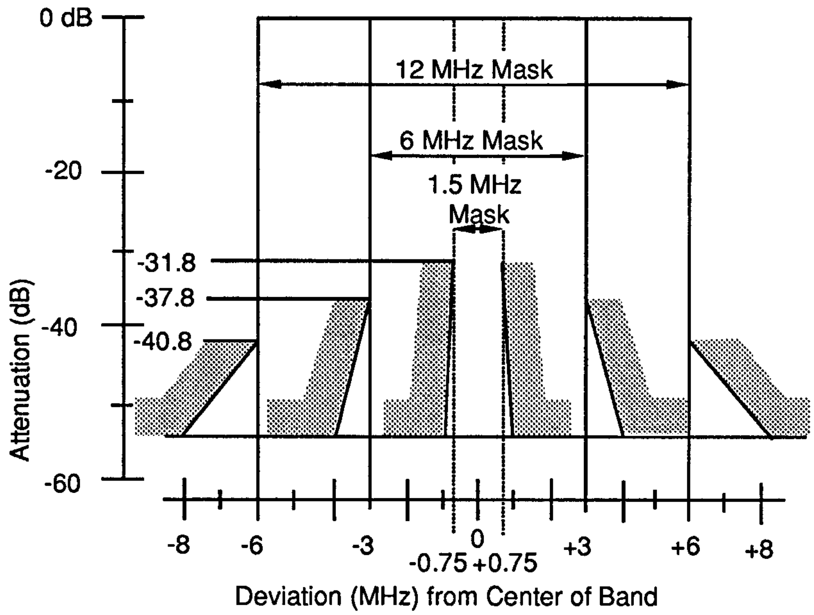

Fig 14-4
Transmit Spectrum Mask

This does not apply to transmissions terminated through jabber-inhibit.

14.8.5.8 When the transmitter is disabled or is transmitting only the MAC-symbol *silence*, the transmitted power in the transmit portion of the medium spectrum that is outside the channel bandwidth shall not exceed –25 dBmV measured in any 30 kHz band. When the transmitter is disabled or is transmitting only the MAC-symbol *silence*, the transmitted power in the receive portion of the medium spectrum that is outside the channel bandwidth shall not exceed –10 dBmV measured in any 30 kHz band.

14.8.5.9 Where separate input and output physical couplings to the medium exist (for example, in dual-cable systems), the leakage output into the input coupling

(1) During transmitter-off periods shall be the same as that specified for the output coupling in 14.8.5.7.

(2) During transmitter-on periods shall be the same as the maximum out-of-band attenuation specified for the output coupling in 14.8.5.6.

14.8.5.10 The residual power level of the suppressed carrier at the transmitter output shall be at least 32 dB below the total transmitted power.

14.8.6 Jabber-Inhibit. Each PLE other than a remodulator (which transmits continuously) shall have a self-interrupt capability to inhibit modulation from reaching the LAN medium. Hardware within the PLE (with no external message other than the prolonged detection of an output-on condition within the transmitter) shall provide a nominal window of one-half second ±25% during which time a normal data link transmission may occur. If a transmission is in excess of this duration, the jabber-inhibit function shall operate to inhibit any further output from reaching the medium. Reset of this jabber-inhibit function is implementation-dependent.

14.8.7 Coupling to the Medium. The PLE functions are intended to operate satisfactorily over 75 Ω CATV-like bidirectional broadband cable installations employing mid-split, sub-split, or high-split line amplifiers and filters, or dual unidirectional cables employing (no-split) line amplifiers.

Both the transmitter and the receiver shall be ac-coupled to the center conductor of the 75 Ω broadband medium, and the breakdown voltage of that ac coupling means shall be at least 500 V ac rms at 50/60 Hz. In addition to this coupling, the shield of the coaxial cable medium shall be connected to chassis ground, and the impedance of that connection shall be less than 0.1 Ω.

The mechanical coupling of the station to the medium shall be to a drop cable by way of a female F-series connector on the station, as specified in Section 15. The maximum voltage standing wave ratio (VSWR) at that F-connector shall be 1.5:1 or less when that F-connector is terminated with a 75 Ω resistive load, as measured over the transmit and receive channel bandwidth when power is applied to the station. This is only applicable to the transmit channel when the station is transmitting.

14.8.8 Alternate Broadband Media (OPTIONAL). Sections 14 and 15 specify the standard (preferred) broadband medium configuration as a single trunk cable and single drop cable, bidirectional by frequency, mid-split CATV-like configuration with channel pairings as indicated in 14.8.4. Other configurations are possible that meet the intent of this standard, including alternate channel pairings, sub-split, high-split, or dual-cable media, and redundant media.

Any channel pairing or range of channel pairings compatible with any of mid-split, sub-split, or high-split single-cable operation, or with dual-cable operation, or both, is permissible, provided that the pairing(s) is (are) clearly labeled both on the embodiment and in accompanying literature.

Embodiments of this standard that can be configured to function with either single-cable or dual-cable systems, or with redundant media, or both, are not precluded, provided that the embodiment *as delivered* functions correctly in a minimal (that is, nonredundant, single-cable, or both) environment.

14.8.9 Receiver Sensitivity and Selectivity. The PLE shall be capable of providing an undetected bit error rate of 10^{-9} or lower, and a detected bit error rate of 10^{-8} or lower, while

(1) Receiving in-band signals in the range specified by Table 14-6, transmitted by a PLE conforming to the requirements of Section 14, and
(2) The in-band signal-to-noise ratio of the medium is 30 dB (V rms/(V rms)) or greater, as measured at the F-connector, and
(3) The adjacent channels on both sides of the in-band signal contain signals conforming to Section 14 of this standard and whose levels are 14 dB greater than the in-band signal, yet not exceeding +4 dB above the maximum received signal power specified in Table 14-6.

14.8.10 Minimum Length of Preamble. The *pad_idle* sequence is used by PLEs to accurately access the received signal levels. In the remodulator, this initial assessment shall be completed at least four PHY-symbols before the end of this *pad_idle* sequence (see 8.2.2.5). The minimum length of that *pad_idle* sequence is a function of both the data rate and whether that *pad_idle* sequence follows *silence* or other non-*silence* symbols. Table 14-7 expresses that relationship in both microseconds and signaled MAC-symbols for three-symbol (one MAC-symbol/PHY-symbol) signaling.

Table 14-6
Required Noise Floor and In-band Signal Power

Channel Bandwidth (MHz)	Max In-band Noise Floor (for remodulators only)	Received Signal Power (dBmV)	
	(dBmV)	Min	Max
1.5	–40	–13	+ 4
6.0	–34	– 7	+10
12.0	–31	– 4	+13

NOTES: (1) Specifications of signal power, both out-of-band and in-band due to adjacent channels, are for future study. It is anticipated that the specification format will be to include an unweighted signal power over the total bandwidth (5–450 MHz), and a weighted signal power over the same bandwidth. The specification of unweighted power is to constrain the maximum power presented to a receiver. The specification of weighted power is to constrain the noise power contribution by all other signal sources as presented to a receiver's demodulator.

(2) When measuring bit error rate, multiple errors in a frame may be counted as a single error.

Table 14-7
Minimum Length of *Pad_idle* Preamble

Data Rate (Mb/s)	At Start of Transmission		Between Frames	
	Minimum # of MAC-Symbols	Minimum Time (μs)	Minimum # of MAC-Symbols	Minimum Time (μs)
1	32	32.0	8	8.0
5	32	6.4	16	3.2
10	32	3.2	24	2.4

14.8.11 Symbol Timing. Each PLE whose receiver is connected to the forward channel shall recover the PHY-symbol timing information contained within the received signaling, and shall use this recovered timing information to determine the precise rate of symbol transmission. The instantaneous transmit and receive data rates of each such entity shall be identical, with exactly one MAC-symbol (possibly *silence*) transmitted for each MAC-symbol received. (This relationship usually is termed *loopback timing*.) There shall be an implementation-dependent constant-phase relationship between the transmitted PHY-symbol timing and the received PHY-symbol timing. The jitter of the transmitted PHY-symbol timing relative to the PHY-symbol timing within the received signaling shall not exceed a total excursion of 5% of the period of one PHY-symbol.

The remodulator shall originate its own transmit symbol timing. The jitter of the transmitted symbol timing shall not exceed 0.1% of the period of one PHY-symbol. The remodulator may use the known loopback timing arrangement of the other stations in the LAN (as specified in the preceding paragraph) in determining the PHY-symbol timing of any received signaling. In a practical case this means that the remodulator may determine the phase difference between its transmitted and received PHY-symbol timings upon first receipt of a transmission, and then use its self-generated transmit timing plus that measured phase delay to infer the PHY-symbol timing of the received signaling.

14.8.12 Signaling Type Determination. The remodulator shall determine the type of signaling (one MAC-symbol/PHY-symbol, or the extended type as defined in 14.11) to be employed in the LAN before commencing transmission in the forward channel and shall convey that determination to all connected stations by means of the pseudo-silence signaling sequence used in lieu of "transmitted" silence as specified in 14.8.2.1 and 14.8.2.2. The remodulator may change this determination at any time, either due to manually invoked command, or through a management interface, or both; when so doing it shall cease all transmissions for a minimum period of 5 ms before making such a redetermination. (As a practical case, this may be used to fallback from extended signaling to one MAC-symbol/PHY-symbol signaling in an extended-capability system that is experiencing unusually high noise, as inferred from error rates.)

14.8.13 Symbol Decoding. After demodulation and determination of each received signaled PHY-symbol, that PHY-symbol shall be decoded to a corresponding MAC-symbol by the process inverse to that described in 14.8.2, and the decoded MAC-symbols shall be reported at the MAC interface at a data rate determined jointly by the provisions of 14.8.11 and 14.8.12.

All PLEs except the remodulator shall decode the pseudo-silence sequence as an appropriate number of *silence* MAC-symbols and shall report them at the MAC-PHY interface. True silence on the forward channel (an abnormal condition) shall be decoded as an appropriate number of *bad_signal* MAC-symbols and shall be reported as such at the MAC-PHY interface.

Whenever a signaled PHY-symbol sequence is received for which the encoding process has no inverse, those PHY-symbols shall be decoded as an appropriate number of *bad_signal* MAC-symbols and reported as such at the MAC interface.

In such cases, the receiving entity shall resynchronize the decoding process as rapidly as possible.

NOTE: Receivers are permitted to decode each transmitted *pad_idle* MAC-symbol, which was encoded as a {*0*} or a {*4*} PHY-symbol, as a *zero* or a *one* MAC-symbol.

14.8.14 Transmitter Enable/Disable and Received Signal Source Selection. The ability to enable and disable the transmission of modulation onto the broadband medium as directed by the management entity is mandatory in remodulators, where signaled *silence* does not turn off the transmitter, and recommended but optional in all other equipment.

The ability to select the source of received signaling, either a loopback point within the PLE or one of the possibly redundant media, as directed by the management entity, is recommended but optional. When such an option is invoked and the selected source is not one of the media, transmission to all connected media shall be disabled automatically while such selection is in force.

14.8.15 Transmit and Receive Channel Selection (OPTIONAL). The determination of a PLE's transmit and receive channels, either channel center frequencies or channel bandwidth or both, may be alterable. In such cases the alteration may be accomplished either manually, or upon command through the PLE's management interface, or both. All transmissions to the medium shall be inhibited while such changes are being effected, and any signaling received during such a period of change shall be discarded and not reported to the MAC interface.

14.8.16 Redundant Media Considerations. As implied by 14.8.8, redundant broadband media are not precluded from this standard. Where redundant media are employed, the provisions of 14.8.5.4, 14.8.6, and 14.8.14 shall apply separately and independently to each single medium interface, and much of 14.8.14 shall be mandatory. Specifically, the transmitted signal level adjustment of 14.8.5.4 shall be made separately and independently for each connected medium, separate jabber-inhibit monitoring shall exist for each medium (although common inhibition is permissible).

14.8.17 Reliability. The PLE shall be designed so that its probability of causing a communication failure among other stations connected to the medium is less than 10^{-6} per hour of continuous (or discontinuous) operation. For remodulators this requirement is relaxed to a probability of 2×10^{-5} per hour of operation. Connectors and other passive components comprising the means of connecting the station to the coaxial cable medium shall be designed to minimize the probability of total network failure.

14.8.18 Remodulator Considerations. A remodulator can be considered to be a specialized form of regenerative repeater, which either repeats signaling from the reverse channel on the forward channel or originates signaling of its own on the forward channel. In either case, the remodulator originates the symbol timing provided to its local MAC entity and to all the other stations on the LAN.

The basic mode of operation, originating or repeating, shall be determined by the remodulator's associated MAC entity and conveyed by the PHY-MODE invoke primitive (see 8.2.3). When originating, the remodulator's PLE shall use internal loopback as the source of the PHY-symbols which it reports to its associated MAC entity by way of the PHY-UNITDATA indicate primitive.

When repeating, the remodulator shall decode the signaling received on the reverse channel and indicate the decoded MAC-symbols to its associated MAC entity.

In summary, when the MAC entity is originating, internal loopback is used as the source of PHY-symbols which are reported by way of the PHY-UNITDATA indicate primitive. When the MAC entity is repeating, the received signaling from the reverse channel is decoded and reported by way of the PHY-UNITDATA indicate primitive.

In either case, the following apply:

(1) The remodulator determines its own MAC-symbol timing and the exact PHY-symbol signaling rate for itself and for all other stations on the local network (see 14.8.1 and 14.8.11.).

(2) The remodulator determines the signaling type (one MAC-symbol/PHY-symbol, or the extended type as defined in 14.11) for the entire network, and may optionally change that signaling type after disabling its transmitter(s) for a period of at least 5 ms (see 14.8.12).

(3) The remodulator transmits a specific repeating sequence (pseudo-silence) during periods when its MAC entity requests the repeated transmission of the MAC-symbol *silence* (see 14.8.2 and 14.8.12). This pseudo-silence sequence is used

 (a) To provide a continuous signal for the AGC and carrier and data recovery circuitry of the LAN's other stations.

 (b) To indicate the signaling type (one MAC-symbol/PHY-symbol, or the extended type as defined in 14.11) currently in use on the LAN (see 14.8.2.2).

(4) The remodulator can use the known exact signaling rate of the other stations (that is, its own exact signaling rate as retransmitted by those other stations, see 14.8.11) to minimize the time required to train on a new transmission and be able to report the received signaling at the specified bit error rates (see 14.8.9).

NOTE: Designers are cautioned that the remodulator shall be designed to achieve these bit error rates at least four MAC-symbol_times (see 8.2.2) before the end of reception of a minimum-length preamble (see 14.8.10) when the transmission originates from a station whose transmitter output, as measurable at the station's F-connector, falls within the bounds of Table 14-6.

(5) The receive levels may be exceeded during periods when the distributed token bus protocol has more than one station transmitting. A remodulator should be designed to tolerate at least a 30 dB increase in the received signal level over the maximum specified in Table 14-6 minus the path loss (see 15.5.3). Furthermore, the remodulator shall report the presence of received signaling, perhaps as *bad_signal*, and shall not report silence to the MAC sublayer during such periods of overload.

The same phenomenon may occur in any other channel that employs a distributed protocol with potential collisions. Thus, a remodulator should be designed to sharply reject signaling in other channels.

14.9 Environmental Specifications

14.9.1 Electromagnetic Emanation. Equipment shall comply with local and national requirements for limitation of electromagnetic interference. Where no local or national requirements exist, equipment shall comply with CISPR Publication 22 (1985) [3].

14.9.2 Safety Requirements. All stations meeting this standard shall comply with relevant local, national, and international safety codes and standards, such as IEC Publication 950 (1986) [7].

14.9.3 Electromagnetic Environment. Sources of interference from the environment include electromagnetic fields, electrostatic discharge, transient voltages between earth connections, etc. Several sources of interference contribute to voltage buildup between the coaxial cable and the earth connection, if any, of the station.

The PLE embodiment shall meet its specifications when operating in an ambient continuous wave plane wave field of

(1) 80 mV/m within the operational channel bandwidth.

(2) 2 V/m for other frequencies from 10 kHz through 30 MHz.

(3) 5 V/m for other frequencies from 30 MHz through 1 GHz outside the operational channel bandwidth.

14.10 Labeling. It is recommended that each embodiment (and supporting documentation) of a PLE conformant to this standard be labeled in a manner visible to the user with at least these parameters:

(1) Data rate capabilities in Mb/s.

(2) Types of signaling—one MAC-symbol/PHY-symbol only, or one MAC-symbol/ PHY-symbol plus the extended type, as described in 14.11.

(3) Transmit and receive channel assignments, both frequency and bandwidth.

(4) Worst-case, one-way repeater delay (for remodulators) or round-trip delay (for other stations) which this equipment induces on a two-way transmission exchange between stations, as specified in 6.1.9.

(5) Operating modes and selection capabilities as defined in 14.8.14 through 14.8.16.

Additionally, when the station has multiple F-connectors (for example, for dual cables, redundant media, or both), the role of each F-connector shall be designated clearly by markings on the station in the vicinity of that F-connector.

Appendix

(This Appendix is not a part of this International Standard or of IEEE Std 802.4-1990, but is included for information only.)

14.11 Appendix—Provisions for Two MAC-Symbol/PHY-Symbol Signaling. This Appendix gives preliminary considerations for a method of enhancing the basic three-symbol multilevel duobinary AM/PSK signaling to double the data rate within the same signaling bandwidth. The method transmits two three-symbol channels in quadrature (that is, on carriers that are in phase-quadrature to each other) at a theoretic signal-to-noise ratio (S/N) penalty of about 5 dB, but with a requirement for coherent or quasi-coherent detection within receivers.

Standardization of this method is for further study; other methods may be considered as well.

In systems where the remodulators and other stations support this extended signaling method (in addition to the required one MAC-symbol/PHY-symbol method of 14.8.2.1 and 14.8.3), additional requirements for interoperability may be needed. The following requirements, each for further study, are proposed:

(1) Each PLE (including a remodulator) conforming to the provisions of 14.11 should be able to fall back to half-speed under the provisions of 14.8.2.1 and 14.8.3. Such fallback should be under management control.

(2) Each PLE (other than a remodulator) capable of signaling by the extended method of 14.11 should be able to recognize and differentiate between the pseudo-silence sequences (see 14.8.2) of its supported signaling methods.

(3) When the signaling method of the network, as determined by the active remodulator, is of a type supported by the PLE, then that entity should infer the type of signaling then being employed in the LAN from the pseudo-silence signaling sequence used in lieu of "transmitted" silence in the forward channel, as specified in 14.8.1 and 14.8.2. This determination should be remade after any period in which continuous silence (lack of modulation) has been detected in the forward channel for at least 2 ms, or after any loopback mode has been deselected and the broadband medium reselected as the source of received symbols (as described in 14.8.14).

If the inferred network signaling type differs from the current signaling type of the PLE, then that entity should report the new signaling type to management (see 9.2.7) and should inhibit transmission while those types differ.

14.11.1 Symbol Encoding. For the enhanced encoding, two separate baseband signals are formed and independently modulated onto in-phase and quadrature carriers for transmission. These are referred to in the following discussion as the in-phase and quadrature subchannels, respectively. Successively encoded MAC-symbols are transmitted alternately in the in-phase and quadrature subchannels.

The encoding action to be taken for each of the input MAC-symbols when three-symbol quadrature (two MAC-symbol/PHY-symbol) encoding is employed is as follows:

(1) ***Silence—***

(a) The remodulator shall encode successive *silence* symbols as successive PHY-symbols of the repeating sequence {*2*} {*0*} {*2*} {*0*} {*4*} {*0*} {*0*} {0}, always restarting with the first {*2*} of the sequence for each new period of transmitted pseudo-silence. (Thus, the in-phase subchannel transmission sequence is {*2*} {*2*} {*4*} {*0*}, and the quadrature subchannel transmission sequence is {*0*} {*0*} {*0*} {*0*}. Also, it is permissible to terminate the sequence after any even number of PHY-symbols.)

(b) All stations other than the remodulator shall encode the first two MAC-symbols of *silence* after non-*silence*, and the last two MAC-symbols of *silence* before non-*silence* as the PHY-symbol {*2*}. Such stations shall encode all other MAC-symbols of *silence* as the PHY-symbol {*0*} and shall cease transmission while transmitting only such {*0*} symbols representing silence.

(2) ***Pad_idle—***Consecutive *pad_idle* symbols shall be encoded as successive PHY-symbols of the repeating sequence {*4*} {*0*} {*0*} {*0*}, always restarting the sequence with the {*4*} for each new period of *pad_idle*, and always with the {*4*} transmitted in the in-phase subchannel. (Thus, the in-phase subchannel transmission sequence is {*4*} {*0*}, and the quadrature subchannel transmission sequence is {*0*} {*0*}.)

When sending *pad_idle* MAC-symbols immediately after sending *silence* MAC-symbols, the MAC-sublayer entity shall send at least as many octets of *pad_idle* MAC-symbols as are required to result in a minimum of 64 PHY-symbols.

NOTE: For this enhanced signaling, this minimum is eight octets of *pad_idle* MAC-symbols (that is, 32 baud).

(3) ***Non_data—***The scrambler described in 14.8.2.3 shall be reinitialized (preset to all *ones*) at each *non_data* symbol. Each *non_data* symbol shall be encoded as a {*2*}, and the eight-symbol (frame delimiter) sequence of 8.2.1 shall be encoded as follows: each

$non\_data\ \ non\_data\ \ data_1\ \ non\_data\ \ non\_data\ \ data_2\ \ data_3\ \ data_4$

sequence shall be encoded as the sequence

{*2*} {*2*} {*2*} {*2*} $\{d_1\}$ $\{d_2\}$ $\{d_3\}$ $\{d_4\}$,

where each $data_i$ is either a *zero* or a *one* and each d_i is the corresponding encoding of the MAC-symbol $data_i$, either a {*0*} or a {*4*} as described below, and always with the first {*2*} transmitted in the primary channel. (Thus, the in-phase subchannel transmission sequence is

{*2*} {*2*} $\{d_1\}$ $\{d_3\}$,

and the quadrature subchannel transmission sequence is

{*2*} {*2*} $\{d_2\}$ $\{d_4\}$.)

After each such eight-symbol (frame delimiter) sequence, as defined in 8.2.1, the encoder shall start to maintain a sense of PHY-octet alignment, with the next output PHY-symbol considered to be the first of a new PHY-octet.

NOTES: (1) This PLE sense of PHY-octet alignment is identical to the MAC entity's sense of MAC-octet alignment.

(2) Since *non_data* symbols occur only in frame delimiters (see 8.2.1), the PLE may infer that the MAC sublayer is outputting a frame delimiter sequence upon receipt of its first *non_data* symbol.

(4) ***Zero and one***—Each *zero* and *one* symbol shall be encoded as follows:

(a) The MAC-symbol shall be applied as input to the scrambler described in 14.8.2.3, resulting in a "scrambled" *zero* or *one* MAC-symbol;

(b) The scrambled MAC-symbol shall be converted to one of the PHY-symbol output codes as follows:

zero → {0}
one → {4}

This is the $data_i$ to $\{d_i\}$ mapping referred to in 14.11.1(3).

(c) The following procedure shall be performed on each resultant PHY-octet of {*0*} and {*4*} PHY-symbol codes. Each such octet of these PHY-symbol codes shall be compared with the immediately prior octet of PHY-symbols *as transmitted* (that is, after execution of this procedure on that prior octet, if applicable). If the pair of PHY-octets is identical to a single ordered pair of PHY-symbol codes repeated eight times (that is, $(\{0\}\ \{0\})^8$, $(\{0\}\ \{4\})^8$, $(\{4\}\ \{0\})^8$, or $(\{4\}\ \{4\})^8$), so that the in-phase and quadrature subchannel transmission sequences each consist of a single PHY-symbol code, possibly different in each subchannel, repeated eight times, then the last three codes in the PHY-octet under examination shall be altered as follows:

{*0*} {*0*} {*0*} shall be replaced by {*2*} {*4*} {*2*}
{*0*} {*4*} {*0*} shall be replaced by {*2*} {*0*} {*2*}
{*0*} {*0*} {*4*} shall be replaced by {*2*} {*4*} {*2*}
{*4*} {*4*} {*4*} shall be replaced by {*2*} {*0*} {*2*}

In other words, the last code of the in-phase subchannel transmission sequence is altered before transmission as follows:

{*0*} is replaced by {*4*}
{*4*} is replaced by {*0*}

and the last two codes of the quadrature subchannel transmission sequence are changed to {*2*} {*2*} before transmissions.

NOTE: This procedure can alter, at most, every other octet of PHY-symbols; and the second octet after the encoded (frame delimiter) sequence of 8.2.1 is the first such octet that can be altered. This procedure has no effect on the scrambler described in 14.8.2.3 and in the previous subsection.

14.11.2 Baseband Modulation. The PHY-symbol coded signals shall be converted to appropriately shaped pulses whose amplitude at the correct sampling intervals is determined by the output PHY-symbol codes as follows (where MAX is the amplitude of the largest pulse):

{*0*} → 0/4 of MAX (= 0)
{*2*} → 2/4 of MAX
{*4*} → 4/4 of MAX (= MAX)

The sequence of pulses corresponding to the first, third, and so forth (odd numbered) PHY-symbols of each transmission sequence shall be shaped for transmission in the in-phase subchannel, as described in 14.8.3, so that their decoding can be achieved by full-wave rectification of the demodulated in-phase component of the line signal. This sequence is known as the baseband in-phase modulation.

The sequence of pulses corresponding to the second, fourth, and so forth (even numbered) PHY-symbols of each transmission sequence shall be shaped for transmission in the quadrature subchannel, as described in 14.8.3, so that their decoding can be achieved by full-wave rectification of the demodulated quadrature component of the line signal. This sequence is known as the baseband quadrature modulation.

14.11.3 Line Signal in RF Band (at the line output of the station). The equivalent baseband in-phase modulation of 14.11.2 shall be applied to an (in-phase) RF carrier centered in the assigned channel, the equivalent baseband quadrature modulation of 14.11.2 shall be applied to an RF carrier in phase-quadrature with the preceding (in-phase) RF carrier, and the resultant modulated carriers shall be combined and coupled to the broadband bus medium as specified in 14.8.7.

NOTE: Any equivalent method of impressing in-phase and quadrature components on the RF carrier is acceptable.

Appendix

(This Appendix is not a part of this International Standard or of IEEE Std 802.4-1990, but is included for information only.)

14.12 Appendix—Detailed Scrambling and Descrambling Process

14.12.1 Scrambling. The message polynomial is divided by the generating polynomial $1 + X^{-6} + X^{-7}$. (See Fig 14-5.) The coefficients of the quotient of this division are taken in descending order from the data sequence to be transmitted.

14.12.2 Descrambling. At the receiver the incoming bit sequence is multiplied by the generating polynomial $1 + X^{-6} + X^{-7}$ to form the recovered message polynomial. The coefficients of the recovered polynomial, taken in descending order, form the output data sequence.

14.12.3 Elements of the Scrambling Process. The factor $1 + X^{-6} + X^{-7}$ randomizes the transmitted data over a sequence length of 127 bits. Figure 14-6 is given as an indication only, since with another technique the logical arrangement might take another form.

Fig 14-5
$1 + X^{-6} + X^{-7}$ Scrambler

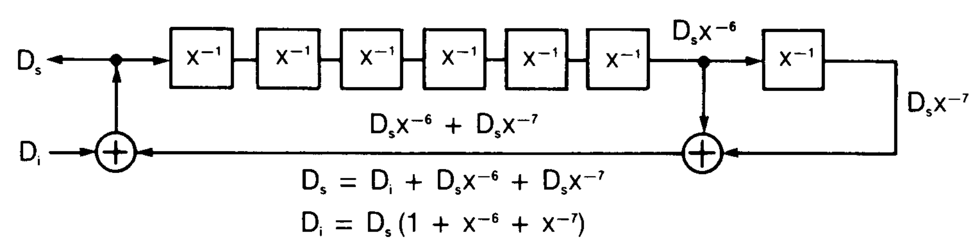

Fig 14-6
Exemplary Scrambler/Descrambler Circuitry

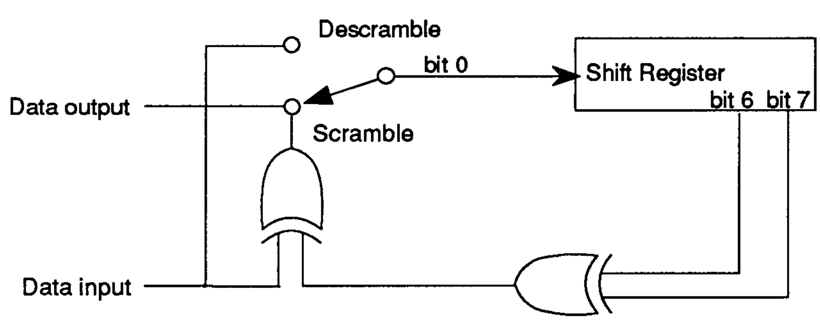

15. Broadband Bus Medium

The functional, electrical, and mechanical characteristics of one specific form of broadband bus medium of this standard are specified in this section. This specification includes the medium embodiments of a broadband bus Local Area Network (LAN). The relationship of this section to other sections of this standard and to LAN specifications is illustrated in Fig 15-1. The relationship between this specification and the broadband bus Physical Layer and the medium is illustrated in Fig 15-2.

This standard specifies the medium only insofar as necessary to ensure

(1) The interoperability of PLEs conforming to Section 14 when connected to a medium conformant to this section, and

(2) The protection of the LAN and those using it.

IEEE Std 802.7-1989 [11] gives additional guidance on the use of broadband cable systems for LAN data communications.

Fig 15-1
Relation to LAN Model

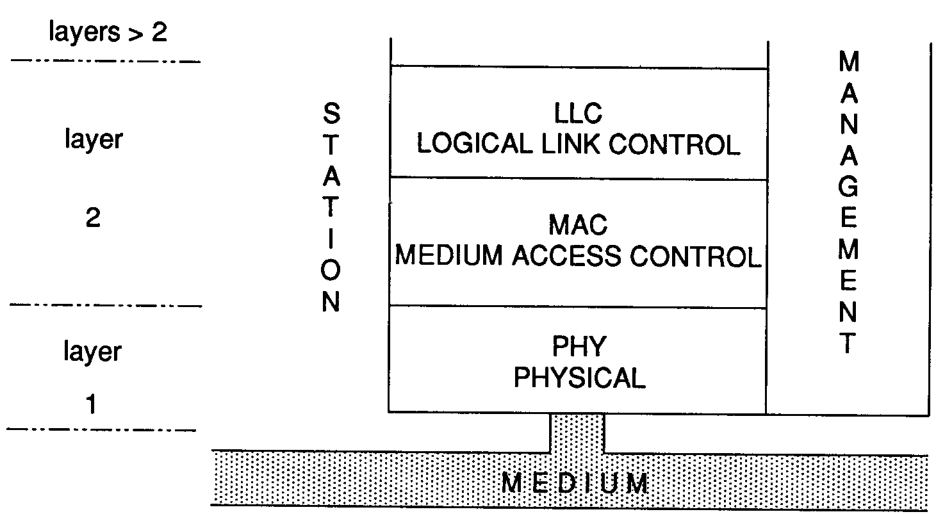

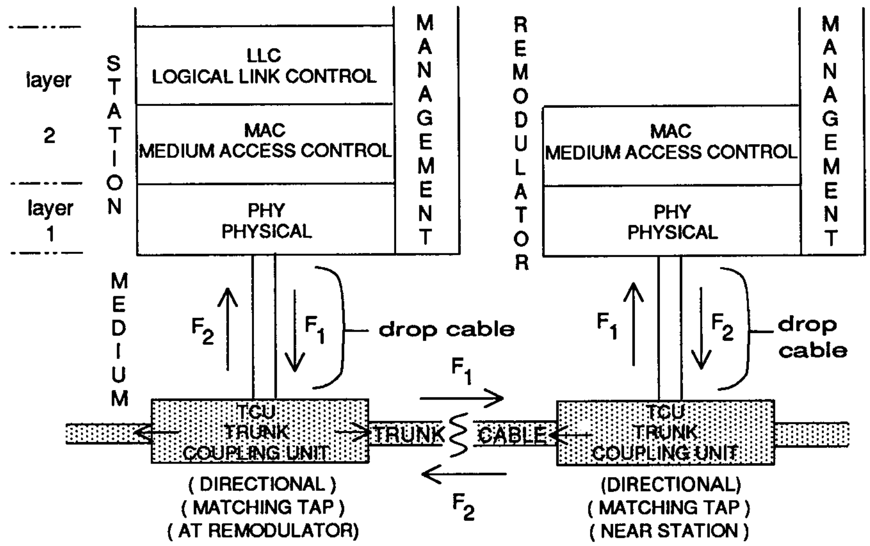

Fig 15-2
Physical Hardware Partitioning

15.1 Nomenclature. Terms used in this section whose meanings within the section are more specific than indicated by glossaries referred to in 1.2 are as follows:

bidirectional broadband amplifier. An assemblage of amplifiers and filters that amplifies and re-equalizes in the *forward* direction all signals received in the higher-frequency portion of the broadband spectrum, and simultaneously amplifies and re-equalizes in the *reverse* direction all signals received in the lower-frequency portion of the broadband spectrum.

broadband coaxial system. A system whereby information is encoded, modulated onto a carrier, and band-pass filtered or otherwise constrained to occupy only a limited frequency spectrum on the coaxial transmission medium. Many information signals can be present on the medium at the same time without disruption provided that they all occupy nonoverlapping frequency regions within the cable system's range of frequency transport.

dBmV. A measure of rms signal level on a 75 Ω cable compared to 1 mV. In terms of SI units, dBmV is defined as dB (1 mV, 75 Ω) rms.

drop cable. The smaller diameter flexible coaxial cable of the broadband bus medium that connects to a station.

dual-cable. A broadband coaxial cable system in which separate coaxial cables are used for the *forward* and *reverse* directions of signal transmission. Connection of a dual-cable system to a station requires dual F-connectors at the station—one for transmission and one for reception.

F-connector. A 75 Ω F-series coaxial cable connector.

forward. The direction of transmission originating at the remodulator of a broadband cable system and relayed "outbound" by the system's bidirectional broadband amplifiers to the system's "subscribers." Transmission on the higher-frequency channels is supported in this direction.

remodulator. The unit located at the logical root of a broadband bus LAN which demodulates the F_1 frequency signals transmitted by other network stations and rebroadcasts them as F_2 frequency signals back to those stations. In this standard, it transmits continuously and in doing so specifies the precise data rate at which all stations shall transmit.

high-split. The broadband system configuration, as determined by the system's bidirectional amplifiers, in which signaling in the spectrum from 5–174 MHz is relayed in the *reverse* direction and signaling in the spectrum from 234 MHz up is relayed in the *forward* direction. (The frequency range used is determined by convention. The values given here are commonly found in North American CATV systems.)

mid-split. The broadband system configuration, as determined by the system's bidirectional amplifiers, in which signaling in the spectrum from 5–108 MHz is relayed in the *reverse* direction and signaling in the spectrum from 162 MHz up is relayed in the *forward* direction. This is the configuration preferred in the broadband portion of this standard. (The frequency range used is determined by convention. The values given here are commonly found in North American CATV systems.)

(impedance-matching) power splitter. A small module that electrically and mechanically couples one large diameter trunk cable to other large diameter trunk cables, providing a branching topology for the broadband trunk. A power splitter splits the signal energy received in the forward direction among the outgoing trunks, and combines any signal energy received in the reverse direction. It passes low-frequency (<1 kHz) ac power between the trunk cables. It contains only passive electrical components (R, L, C).

reverse. The direction of transmission originating at the "subscribers" of a broadband cable system, terminating at the remodulator and relayed "inbound" by the system's bidirectional broadband amplifiers. Transmission on the lower-frequency channels is supported in this direction.

(impedance-matching) splitter. A smaller version of the power splitter used to couple drop cables together. It does not pass low-frequency ac power between the drop cables.

sub-split. The broadband system configuration, as determined by the system's bidirectional amplifiers, in which signaling in the spectrum from 5–30 MHz is relayed in the *reverse* direction and signaling in the spectrum from 54 MHz up is relayed in the *forward* direction. (The frequency range is determined by convention. The values given here are commonly found in North American community residential CATV systems.)

(impedance-matching) tap. A small module that electrically and mechanically couples the large diameter trunk cable to smaller diameter drop cables, passes low-frequency (<1 kHz) ac power between input and output trunk cable sections, and isolates that ac power from drop cable sections. It splits the signal energy received in the forward direction very asymmetrically, with the bulk of that signal energy passed to the outgoing trunk cable and only a small percentage going to the drop cables, and combines with similar asymmetry any signal energy received in the reverse direction. It contains only passive electrical components (R, L, C).

trunk cable. The main (larger-diameter) semirigid coaxial cable of a broadband coaxial cable system. Both ac power and RF signaling are carried on the cable.

15.2 Object. The object of this section is to

(1) Provide a description of the physical medium necessary for communication between local network stations employing the token-passing bus access method defined in this standard and a broadband bus PLE.
(2) Provide a medium that can be shared by multiple LANs and by totally unrelated applications (such as voice, data, and analog video).
(3) Provide for high network availability.
(4) Provide for ease of installation and service in a wide range of environments.
(5) Permit use of existing components mass produced for the CATV industry, and of the corresponding set of installation practices, where feasible.

15.3 Compatibility Considerations. This standard applies to media that are designed to operate as conventional bidirectional (by frequency) CATV-like broadband coaxial cable systems. Such systems generally use standard CATV taps, connectors, amplifiers, power supplies, and coaxial cable. This specification primarily applies to a single trunk system in which two-way alternate communication usually is accomplished through the use of bidirectional amplifiers whose filters permit differing parts of the available cable spectrum to be transmitted in each direction. Dual-trunk systems in which each trunk is used unidirectionally are also permitted, but require both stations and the remodulator to have the necessary dual-cable connector fittings.

The use of a coaxial broadband system permits the assignment of different frequency bands to multiple simultaneous applications. For example, a portion of the cable spectrum can be used by LANs while other portions are used for point-to-point or multipoint data links, and others are used to convey television and audio signals. With the proper selection of signal levels and proper equipment design all of the applications can be propagated without deleterious interference.

Compatibility of medium conformant to this standard is determined at their station interfaces. Specific implementations based on this standard may be conceived in different ways provided compatibility at the actual station interfaces is maintained.

15.4 Overview

15.4.1 General Description of Functions. The functions performed by the broadband bus medium are described informally here. Jointly, these functions provide a means whereby signaling in the reverse-direction spectrum presented at the station interfaces of drop cables can be combined and conveyed as corresponding forward-channel spectrum signaling (with the assistance of remodulators functioning as relay entities) to all of the stations on all of the medium's drop cables. Thus, stations connected to these drop cables can communicate.

15.4.1.1 Operational Overview. The use of a coaxial cable operating as a broadband medium introduces special system considerations that are not encountered in single-channel transmission systems. A conventional CATV-like system uses bidirectional amplifiers to achieve two-way transmission on a single coaxial cable. The amplifier actually consists of two independent unidirectional amplifiers, together with the appropriate crossover filters, so that one amplifier transmits the upper portion of the cable spectrum in the *forward* direction, outbound from the remodulator, while the other amplifier transmits the lower portion in the *reverse* direction, toward the remodulator. (Therefore, each station requires an RF transmitter and receiver, each operating in a different portion of the cable spectrum.)

The remodulator is where forward direction signals originate and reverse direction signals converge. The remodulator receives the lower-frequency reverse channel signal and retransmits it as the higher-frequency forward channel signal to the "downstream" stations. One remodulator is needed for each separate channel pair used as a broadband bus LAN.

Stations are connected to the *trunk* coaxial cable in such systems by smaller-diameter *drop* cables and impedance-matching *taps*. These taps are passive devices that are highly directional with regard to signal propagation as shown in Fig 15-2. The directional characteristics of the taps improve the impedance matching at the *ports* of the tap and minimize the effects of reflections due to any impedance mismatches along the trunk or on other drop cables.

The directionality of the tap is such that the transmission loss between the station and the reverse direction trunk is considerably less than that between the station and the forward direction trunk. Hence, communication between devices connected to different taps along the coax occurs through the transmission of a frequency F_1 in the low-band toward the remodulator, reception at the remodulator and retransmission in the forward direction at a remodulator frequency F_2, and reception by the "listening" stations.

NOTE: Remodulator taps are connected with forward and reverse directions interchanged to support the remodulator's role as a relay entity.

The topology of the broadband system is that of a highly branched tree, with the remodulator equipment located at the tree's roots, and the other equipment connected as leaves to the tree's branches. Branching is accomplished in the trunk itself by way of *power splitters*, which provide nondirectional coupling of the ac power carried on the trunk cables, and directional coupling of the RF power similar to that of the just-described *taps*. Like the taps, the power splitters also employ only passive electrical components (R, L, C only).

Branching in the drop cables is provided by (drop cable) splitters that block ac power flow among their ports. They also employ only passive electrical components.

The amplifiers of the broadband system are powered from the trunk cable itself. Low-voltage (30 V or 60 V) ac power usually is carried on the center conductor and grounded shield of the cable at a nominal frequency of 50 Hz or 60 Hz; the amplifiers use some of that power and pass the rest to the connected "downstream" trunk cable. Standard CATV-like power supplies and power-combining and power-blocking components permit the power to be supplied to the trunk cabling at one or more points in the system.

15.4.1.2 Remodulator Functions. In an actual broadband bus system, stations do not transmit and receive on the same frequency channel. For each broadband bus LAN, one central station, designated the *remodulator*, receives on a low-frequency channel and transmits on a high-frequency channel. All of the other stations of that particular network receive on the high-frequency channel and transmit on the low-frequency channel. Thus, any one network really uses a pair of directional channels, with the low-frequency (*reverse*) channel having many transmitters and one receiver, and the high-frequency (*forward*) channel having one transmitter and many receivers.

To support its functioning as a relay, the remodulator is located at the "root" of the tree-like structure of the medium, and is connected to the trunk by way of a tap so installed that the remodulator receives signaling originating "downstream" from it, and also transmits its signal "downstream."

15.4.2 Basic Characteristics and Options. All signal-conveyance characteristics and station-interface characteristics are mandatory. All other characteristics are optional.

15.5 Functional, Electrical, and Mechanical Specifications. The broadband bus medium is an entity whose sole function (relative to this standard) is signal transport between the stations and remodulator of a broadband bus LAN. Consequently, only those characteristics of the medium that impinge on station to remodulator and remodulator to station signal transport, or on human and equipment safety, are specified in this standard.

An implementation of the medium shall be deemed conformant to this standard if it provides the specified signal transport services and characteristics for the stations and remodulator of a broadband bus LAN, and if it meets the relevant safety and environmental codes.

NOTE: IEEE Std 802.7-1989 [11] provides some guidelines in this area.

All measurements specified in 15.5 are to be made at the point of station or remodulator connection to the medium. Unless otherwise indicated, these measurements shall be made across that central portion of the 1.5 MHz, 6 MHz, or 12 MHz channel actually used for signaling, that is, the central 1 MHz, 5 MHz, or 10 MHz, respectively. Unless otherwise stated, all voltage and power levels specified are in rms and dBmV, respectively, based on transmissions of random data patterns (dBmV is defined as dB [1 mV, 75 Ω] rms in 15.1). (See 14.8 for more details.)

15.5.1 Coupling to the Station/Remodulator. The connection of the broadband bus medium to the station or remodulator shall be by way of a flexible 75 Ω drop cable terminated in a male F-series 75 Ω connector; this combination shall mate with a female F-series 75 Ω connector mounted on the station or remodulator.

In addition to this coupling, the shield(s) of the coaxial drop cable medium shall be connected to the shell of the terminating male F-connector and the impedance of that connection over the range from dc to 450 MHz shall be less than 0.1 Ω. Also, the impedance of a connection between the shell of that male F-series connector and the outer barrel of a mated female F-series connector over the range from dc to 450 MHz shall be less than 0.1 Ω. In addition to this coupling, the shield(s) of the coaxial drop cable medium shall be connected to the shell of the terminating male F-series connector and the dc impedance of that connection shall be less than 0.001 Ω. Also, the dc impedance of a connection between the shell of that male F-series connector and the outer barrel of a mated female F-series connector shall be less than 0.001 Ω.

NOTE: At the time this standard was written, there were no recognized standards for F-connectors. A standard for F-connector mating interface dimensions was being drafted by the IEC as Publication 169 : (to be determined), Radio-frequency connectors (Type F). A comprehensive Type FD connector standard (intermateable with Type F) was being drafted by the Electronic Industries Association to be published as EIA 550. F-connectors vary greatly in quality. It is suggested that high quality F-connectors be used. Further, since the center conductor of the cable in some F-connectors becomes the male part of the connector, caution should be used in mixing different types of cables with F-connectors. Where possible, the use of connectors with captive center conductors is recommended.

15.5.1.1 Station-to-Station Isolation. The isolation from any point of station connection to any other point of station connection within the transmit and receive channel bandwidth shall be a minimum of 25 dB.

15.5.2 Characteristic Impedance. The nominal characteristic impedance of the broadband medium shall be 75 Ω. The maximum VSWR at each of the medium's F-connector shall be 1.5:1 or less when the F-connector is terminated with a 75 Ω resistive load, as measured over the entire broadband cable frequency spectrum supported by the medium.

NOTE: This spectrum shall include both the forward and reverse channels in use in the network.

15.5.3 Path Loss. The broadband bus medium shall exhibit a total path loss from any single station to the remodulator or from the remodulator to any single station of 44 dB with a permissible uncertainty due to all factors of ±5 dB. The maximum short-term variance component of this path loss shall be no more than ±1 dB/minute.

15.5.4 Signal-to-Noise (S/N) Level. It is recommended that the signal-to-noise level be at least 40 dB (V rms/[V rms]). In no case shall it be less than 30 dB (as measured at any station or remodulator). The noise floor, averaged across the channel bandwidth, shall not exceed –57 dBmV as measured in any 30 kHz band at the remodulator.

15.5.5 Power Handling Capability. The 50 Hz or 60 Hz ac power commonly carried on broadband trunk cables shall not be carried on the (drop-cable) medium that connects directly to the station or remodulator. The total power over the entire cable spectrum, as presented to the station or remodulator, shall be less than +55 dB dBmV.

15.5.6 Compatibility with the Stations and Remodulator. An embodiment of a broadband bus medium is deemed to support a specific broadband bus LAN if

(1) It provides a *forward* channel of suitable bandwidth for the LAN and a *reverse* channel of equal bandwidth.

(2) The stations of that LAN can receive on that forward channel and transmit on that reverse channel.

(3) The network has a remodulator connected closer to the cable system's head-end than all (other) stations, and that remodulator can transmit on that forward channel and receive on that reverse channel.

(4) The requirements of 15.5.1 through 15.5.5 (inclusive) are met on the reverse channel when measured from the point of remodulator connection to the medium, independent of which point of station connection to the medium is chosen for test signal origination.

(5) The requirements of 15.5.1 through 15.5.5 (inclusive) are met on the forward channel when measured from each point of station connection to the medium, with test signals originated at the point of remodulator connection to the medium.

15.5.7 Alternate Broadband Media (OPTIONAL). Section 14 specifies the standard (preferred) broadband bus medium configuration as a single trunk cable and single drop cable, bidirectional by frequency, mid-split CATV-like configuration. Other configurations are possible which meet the intent of this standard, including sub-split, high-split, or dual-cable media, and redundant media.

15.5.8 Redundancy Considerations. As implied by 14.8.8 and 14.8.16, redundant broadband media are not precluded from this standard. Where redundant media are employed, the provisions of 15.5.1 to 15.5.6 shall apply separately and independently to each single nonredundant medium interface. Additionally, the forward channel frequency and bandwidth, and reverse channel frequency and bandwidth, for the LAN shall be the same on all of the redundant media.

15.5.9 Reliability. All active (powered) medium equipment should be designed so that the aggregate probability of that equipment causing a communication failure at more than one station connected to the medium is less than 10^{-4} per hour of continuous (or discontinuous) operation. Connectors and other passive components comprising the means of connecting the station to the coaxial cable medium shall be designed to minimize the probability of total network failure.

15.6 Environmental Specifications

15.6.1 Electromagnetic Emanation. LAN cable systems shall comply with local and national requirements for limitation of electromagnetic interference. Where no local or national requirements exist, equipment shall comply with CISPR Publication 22 (1985) [3].

15.6.2 Safety Requirements. All media meeting this standard shall comply with relevant local, national, and international safety codes and standards such as IEC Publication 950 (1986) [7].

15.6.3 Electromagnetic Environment. Sources of interference from the environment include electromagnetic fields, electrostatic discharge, transient voltages between earth connections, etc. Several sources of interference contribute to voltage buildup between the coaxial cable and the earth connection, if any, of the station.

The medium embodiment shall meet its specifications when operating in an ambient continuous wave plane wave field of

(1) 80 mV/m within the operational channel bandwidth,

(2) 2 V/m for other frequencies from 10 kHz through 30 MHz,

(3) 5 V/m for other frequencies from 30 MHz through 1 GHz outside the operational channel bandwidth.

NOTE: IEEE Std 802.7-1989 [11] provides some guidelines in this area.

15.7 Transmission_path_delay Considerations. When specifying an embodiment of a medium that conforms to the specifications of this Section, a vendor shall state the transmission_path_delay as the maximum delay that the broadband bus medium could be expected to induce on a transmission from any connected station through the remodulator and back to that station. The delays induced by the transmitting and receiving stations themselves should not be included in the transmission_path_delay.

For each potentially worst-case path through the medium, a path delay is computed as the sum of the medium-induced delay and remodulator-induced delay in propagating a signal from one station to another. The transmission_path_delay used for determining the network's slot_time (see 6.1.9) shall be the largest of these path delays for the cable system.

These path delay computations shall take into account all circuitry delays in the remodulator and the medium amplifiers, taps, splitters, etc., and all signal propagation delays within the cable segments themselves.

The transmission_path_delay shall be expressed in terms of the network's symbol signaling rate on the medium. When not an integral number of signaled symbols, it shall be rounded up to such an integral number. When uncertain of the exact value of the delay, vendors shall state an upper bound for the value.

15.8 Documentation. It is recommended that each vendor of an embodiment of a medium conformant to this standard provide to the user supporting documentation with at least these parameters:

(1) Specific sections of the standards to which the embodiment conforms.
(2) Configuration type—sub-split, mid-split, high-split, or dual-cable.
(3) Upper limit of forward channel frequency spectrum (for example, 300 MHz, 400 MHz, 450 MHz).
(4) The transmission_path_delay, as specified in 15.7 and 6.1.9.
(5) For remodulators, the proper range of transmitter output levels so that the signal-level requirements of 14.8.5.4, 14.8.9, and 14.8.19 are all met.

15.9 Network Sizing

15.9.1 Topology Considerations. Very small systems can be constructed using only flexible coaxial cable and passive impedance-matching networks. Larger systems require amplifiers and power supplies, together with semirigid trunk cable and flexible drop cables. Highly-branched topologies are achieved easily in this medium using standard CATV-like components.

Each remodulator shall be "upstream" from all of the (other) stations on its LAN, but need not actually be at the head-end of the entire cable system. It is only necessary that the remodulator be closer to the "root" of the tree-shaped cable system than any of the other stations of its network. Thus, in large systems (for example, campus-wide or city-wide), a large cable tree could have many subtrees, with a remodulator at the root of each subtree controlling a LAN within that subtree, and with all the networks using the same pair of forward and reverse channels. Band-stop (notch) filters would be needed in the cable segments that connect those subtrees to isolate the forward and reverse channels of each subtree's network to that subtree.

15.9.2 Signal Loss Budget Considerations. The placement of broadband amplifiers, and the placement of taps, choice of number of drops in each tap, and determination of the tap's insertion and drop losses shall take into account the following:

(1) Each station's minimum transmit level and receive level specifications (see 14.8.5.4 and 14.8.9).
(2) The desired current and anticipated future placement of station's and remodulator.
(3) Each station's transmit level and receive level specifications (see 14.8.5.4 and 14.8.9).

16. Fiber Optic Physical Layer Entity (PLE)

The functional, optical, electrical, and mechanical characteristics of one specific form of PLE (fiber optic) of this standard are specified in this section. This standard specifies these PLEs only in so far as necessary to ensure

(1) The interoperability of implementations conforming to this specification, and

(2) The protection of the Local Area Network (LAN) itself and those who use it.

The relationship of this section to other sections of this standard and to LAN specifications is illustrated in Fig 16-1. The relationship of this section to the fiber optic PLE and medium is illustrated in Fig 16-2.

Fig 16-1
Relation to LAN Model

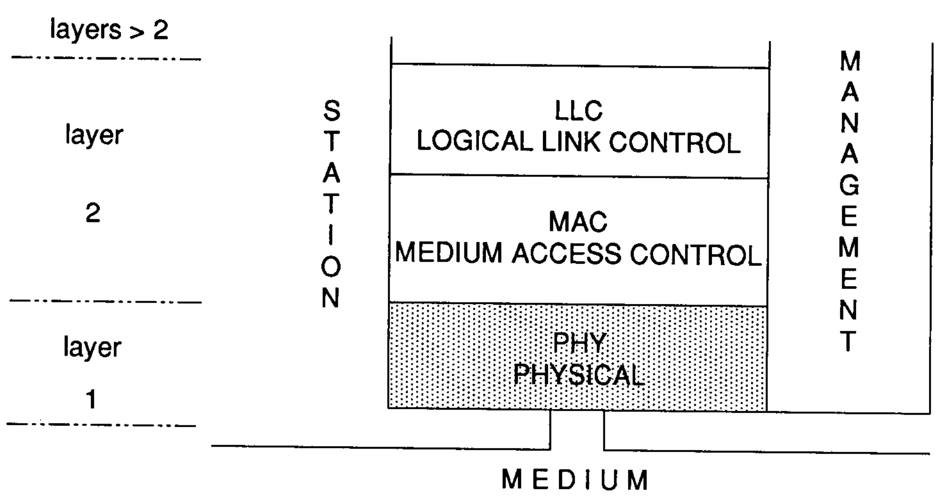

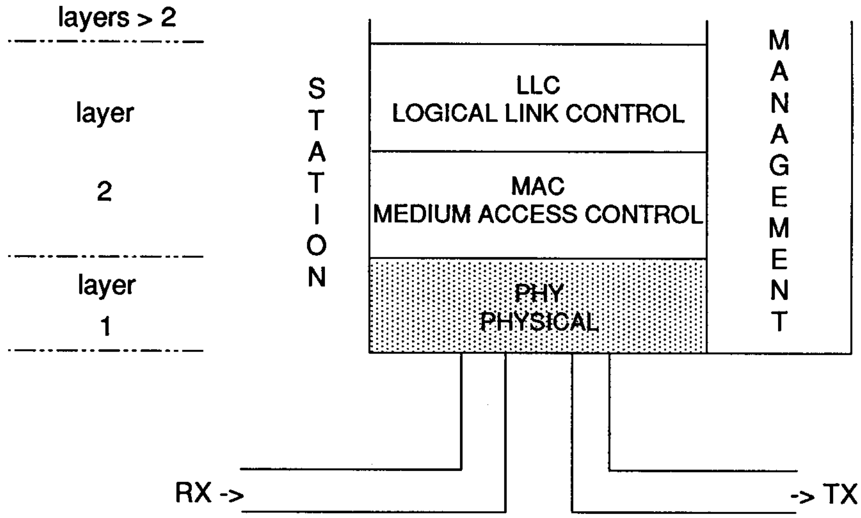

Fig 16-2
Physical Hardware Partitioning

16.1 Nomenclature. Terms used in this section whose meanings within the section are more specific than indicated by glossaries referred to in 1.2 are as follows:

cable plant interface connector (CPIC). The point at which test and conformance measurements are made; the interface between the station and the cable plant. The CPIC is keyed to correspond to the receiver operating range that is supported by the attached PLE. This connector is in the form of a duplex connector whose intermateability geometry is as defined in ISO/IEC 7498-4 : 1989 [14].

center wavelength. The wavelength that is the arithmetic mean of the half-maximum spectral intensity points of the transmitter. If the spectral intensity distribution is symmetric and singly peaked, the center wavelength is at maximum intensity.

dBm. A measure of optical power. In terms of SI units (International System of Units), dBm is defined as dB (1 mW).

detected bit error. An error which is reported as *bad_signal*. *Bad_signal* reported during the preamble or during the four symbols following the last end delimiter of a transmission is not included.

effective launch power. The effective power coupled into the core of a fiber optic waveguide by the transmitter. This power is measured with a standard test fiber connected to the CPIC (see 16.7.3.2).

effective power. The difference, expressed in dBm, between the absolute optical power measured in milliwatts at the midpoint in time of the {*H*} PHY-symbol to the absolute optical power measured in milliwatts at the midpoint of the {*L*} PHY-symbol.

NOTE: Effective power is believed to give a more accurate measurement of the conditions that affect receivers than traditional measurements, such as peak and average power. Methods for measuring effective power are for further study.

extinction ratio. The ratio of the absolute optical power measured in milliwatts at the midpoint in time of the {*H*} PHY-symbol to the absolute optical power measured in milliwatts at the midpoint of the {*L*} PHY-symbol.

NOTES: (1) In computing effective power and extinction ratio, the absolute power of an {*L*} cannot be assumed to be lower than the maximum power of an {*off*}.

(2) The following gives an example of the computation of effective power and extinction ratios:

If the midpoint of {*H*} is measured as	105 µW,
and if the midpoint of {*L*} is measured as	5 µW,
then the difference is	100 µW.

Therefore, the effective power is $10 \log_{10}$ (100 µW/1 mW), which equals –10 dBm. The extinction ratio is 105/5, which equals 21:1.

fiber optic cable. A cable containing one or more fiber optic waveguides. Jacketing material is provided to facilitate handling and to protect the fiber.

fiber optic receiver. The optics and electronics in the station that accept the optical signal received by the station through the CPIC.

fiber optic receiver operating range. The range of effective power that must be present at the CPIC to ensure that the bit error rate specifications are met.

fiber optic transmitter. A device that emits optical signals for propagation onto a fiber optic waveguide through the CPIC.

fiber optic waveguide. A flexible optically transparent strand that is used to transport optical signals from one geographic point to another geographic point.

full width half maximum (FWHM). A measure of the width of a singly peaked function. After determining the maximum (peak) value of the function, one determines the two arguments of the function that yield half of this maximum (peak) value. The FWHM value is the magnitude of the arithmetic difference of these two arguments.

inband optical power. The optical spectral power density integrated over the range from 720–990 nm.

jitter. The offset of the 50% transition points of pulse edges from their ideal position as a result of all causes.

Manchester encoding. A means by which separate data and clock signals can be combined into a single, self-synchronizable data stream, suitable for transmission on a serial channel.

optical fall time. The time it takes for a pulse that is within the limits of Fig 16-3 to go from 90% effective power to 10% effective power. When expressed as a percentage, the fall time is specified as a percent of the PHY-symbol_time.

optical rise time. The time it takes for a pulse that is within the limits of Fig 16-3 to go from 10% effective power to 90% effective power. When expressed as a percentage, the rise time is specified as a percent of the PHY-symbol_time.

optical silence power. When the absolute power imposed on the CPIC from the cable plant is less than this power, the PLE will report *silence* to the MAC sublayer entity. This power is measured with a standard test fiber connected to the CPIC.

PHY-symbol_time. One half the MAC-symbol_time.

standard test fiber. A silica fiber optic waveguide with the following nominal characteristics: core diameter of 62.5 μm, outside diameter of 125 μm, an effective numerical aperture of 0.275, and terminated in a CPIC compatible connector.

NOTE: Operation using a 50 μm alternate test fiber is described in Appendixes 16.10 and 17.9.

undetected bit error. An error that is not reported as *bad_signal* by the PLE.

16.2 Object. The object of this specification is to

(1) Provide the physical means necessary for communication between local network stations employing the token-passing bus access method defined in this standard and a fiber optic PLE,
(2) Provide for high network availability,
(3) Provide for ease of installation and service in a wide range of environments,
(4) Provide for different topologies and three signaling rates to accommodate a variety of user requirements,
(5) Use components that are readily available in the commercial marketplace,
(6) Provide a communication channel capable of high bandwidth and low bit error rate performance.

16.3 Compatibility Considerations. This standard applies to PLEs that are designed to operate on fiber optic cables. For compatibility, fiber optic PLEs are characterized by the keying of the CPIC (and the implied receiver operating range) and the signaling rate. Specific implementations based on this standard may be conceived in different ways provided compatibility at the CPIC is maintained.

This specification assumes the use of 62.5 μm standard test fiber. Implementations using a 50 μm alternate test fiber are specified in Appendix 16.10. For the purposes of compatibility, all measurements listed in this section are referenced to the use of the standard test fiber. Vendors may claim conformance using the standard test fiber, the alternate test fiber, or both test fibers.

A station is implemented using one type of fiber optic transmitter and either of two types of fiber optic receiver: a moderate sensitivity receiver or a high sensitivity receiver. This standard allows multiple compliant topologies to be implemented providing the user flexibility in network design without sacrificing interoperability. The transmitter and moderate sensitivity receiver may be used in active star or low station count passive star topologies. The transmitter and high sensitivity receiver may be used in high station count passive star topologies and in topologies with high media attenuation. The high sensitivity receiver also may be used in place of the moderate sensitivity receiver through the use of manual or automatic attenuation mechanisms.

16.4 Operational Overview. The medium specified in Section 17 consists of a pair of fiber optic waveguides providing bidirectionality by use of a separate fiber for each direction of signal propagation. These dual fibers connect to the CPIC of a station. The exposed point of attachment to the station shall use the CPIC specified in 16.7.12 and the signal characteristics at this interface are to be as specified in this section.

Subsection 17.11.1 provides examples of some of the topologies that are considered compatible within the scope of this standard. To be considered compatible, a topology shall function properly under the control of the MAC machine described in Sections 5, 6, and 7 and utilize PLEs and media detailed in Sections 16 and 17 of this document.

16.5 Physical Layer Overview

16.5.1 General Description of Functions. The functions performed by the fiber optic PLE are described informally here. Jointly, these functions provide a means whereby symbols presented at the MAC interface of one PLE can be conveyed to all of the PLEs via the media for presentation to their respective MAC interfaces.

16.5.1.1 Symbol Transmission and Reception Functions. Successive symbols presented to the PLE at its MAC-PHY interface are applied to an encoder that produces as output a PHY-symbol pair from the set of PHY-symbols {*H*}, {*L*}, and {*off*}. The output is then applied to an optical transmitter, and the resultant optical signal is coupled to the fiber optic broadcast medium and conveyed by the medium to all the receivers.

Each receiver is also coupled to the broadcast medium. It detects the received signal and infers the transmitted PHY-symbol from the presence or absence of optical power. It then decodes that inferred PHY-symbol by an approximate inverse of the encoding process and presents the resultant decoded MAC-symbols at its MAC-PHY interface. The receiver reports *silence* at the MAC-PHY interface when the optical power present at the CPIC is less than the optical silence power.

For all MAC-symbols except *pad_idle*, this decoding process is an exact inverse of the encoding process in the absence of errors. The *pad_idle* symbols, which are referred to collectively as *preamble*, are transmitted at the start of each MAC frame to provide a training signal for receivers and to provide a nonzero minimum separation between consecutive frames. In fiber optic broadcast systems the MAC-symbol encoding for successive *pad_idles* is identical with the encoding for an alternating series of *ones* and *zeros*, and receivers are permitted to decode the transmitted representation of successive *pad_idles* as an alternating series of *ones* and *zeros* and report it as such to the MAC entity.

16.5.1.2 Regenerative Repeater Functions. Regenerative repeaters may be used to extend the network size beyond the maximum that would otherwise be dictated by signal attenuation and timing jitter considerations. Regenerative repeaters do so by connecting two or more medium segments and repeating any thing "heard" on one segment, excepting its own transmissions, to the other segments. For the purposes of this standard, regenerative repeaters are considered stations, whether or not they have functionality beyond that of a repeater.

16.5.1.3 Jabber-Inhibit Function. To protect the LAN from most faults in a station, each station contains a jabber-inhibit function. This function serves as a "watchdog" on the transmitter; if the station does not turn off its transmitter after a prolonged time (roughly one-half second), then the transmitter output shall be automatically disabled for at least the remainder of the transmission.

16.5.1.4 Local Administrative Functions (OPTIONAL). These functions are activated either manually, or by way of the PLE's management interface, or both. They can include

(1) Enabling or disabling each of the station's outputs to the CPICs, and
(2) Selecting the received signal source: either medium or loopback point.

NOTE: If a loopback point is selected, all transmitter outputs to the media shall be inhibited. (See 16.7.9.)

16.5.2 Basic Functions and Options. Symbol transmission and reception functions and jabber-inhibit functions are required in all implementations. All other functions are optional.

16.6 Application of Management. The following constraints are imposed on the parameters and actions specified in Section 9:

(1) In the capabilities group, the dataRates parameter sequence shall specify one or more of the values 5, 10, and 20.
(2) In the local management information, the value of minPostSilencePreambleLength shall be 3 for 5 and 10 Mb/s, and 6 for 20 Mb/s.

16.7 Functional, Optical, Electrical, and Mechanical Specifications. Unless otherwise stated, optical power levels are specified as effective, inband optical power in dBm (dBm is defined as dB [1 mW] in 16.1).

NOTE: Standard test procedures, test fiber, and fixture for measuring launch power, optical silence power, and fiber optic receiver operating range are under study by EIA and IEEE 802.8, the Fiber Optic Technical Advisory Group.

16.7.1 Data Signaling Rates. The standard data signaling rates for single channel fiber optic systems are 5 Mb/s, 10 Mb/s, and 20 Mb/s. The permitted tolerance for each signaling rate is ±0.01% for an originating station, and ±0.015% for repeating stations while performing repeating or regenerating functions.

16.7.2 Symbol Encoding. The PLE transmits symbols presented to it by the MAC sublayer entity at the MAC-PHY interface. The possible MAC-symbols are Manchester-encoded into the PHY-symbols {*H*} and {*L*}. The encoding action to be taken for each of the input MAC-symbols is as follows:

(1) *Silence.* Each *silence* symbol shall be encoded as {*off off*}.

(2) *Pad_idle*. *Pad_idle* symbols are always originated in octets. Each pair of consecutive *pad_idle* symbols shall be encoded as the sequence {*L H*} {*H L*}.

(3) *Zero*. Each *zero* symbol shall be encoded as {*H L*}.

(4) *One*. Each *one* symbol shall be encoded as {*L H*}.

(5) *Non_data*. *Non_data* symbols are transmitted by the MAC sublayer entity in pairs. Each such pair of consecutive *non_data* symbols shall be encoded as the sequence {*L L*} {*H H*} when the immediately preceding PHY-symbol is an {*L*}, and shall be encoded as the sequence {*H H*} {*L L*} when the immediately preceding PHY-symbol is an {*H*}. Thus, the (start delimiter) subsequence

non_data non_data zero non_data non_data zero

shall be encoded as the subsequence

{*L L*} {*H H*} {*H L*} {*L L*} {*H H*} {*H L*}

and the (end delimiter) subsequence

non_data non_data one non_data non_data one

shall be encoded as the subsequence

{*L L*} {*H H*} {*L H*} {*H H*} {*L L*} {*L H*}

after an immediately preceding {*L*}, and as the subsequence

{*H H*} {*L L*} {*L H*} {*H H*} {*L L*} {*L H*}

after an immediately preceding {*H*}.

16.7.3 Line Signal (at the output of the station). The PHY-symbols resulting from the encoding of 16.7.2 shall be converted to its line representation, as described in 16.7.3.1, and the resultant signaling shall be imposed on the single channel fiber optic broadcast medium. The rate at which these symbols are imposed on the medium is a function of the data rate in operation at the time. Table 16-1 summarizes the relationship between data rate, media transmission rate, and PHY-symbol_time.

Table 16-1
Data Rate Versus Media Transmission Rate

Data Rate (in Mb/s)	Media Transmission Rate (in Mbaud)	PHY-Symbol Time (in ns)
5	10	100
10	20	50
20	40	25

16.7.3.1 Line Representations. The line signal representations of the {*H*}, {*L*}, and {*off*} PHY-symbols shall be as follows:

(1) An {*H*} shall be represented nominally as the absolute power level corresponding to 100% effective power for the full period of a PHY-symbol time as shown in Fig 16-3.

(2) An {*L*} shall be represented nominally as 0% effective power for the full period of a PHY-symbol_time as shown in Fig 16-3. The absolute power levels corresponding to transmitted {*H*} and {*L*} PHY-symbols may be any values which, taken together, meet the output power (16.7.3.2) and extinction ration (16.7.3.3) specifications.

(3) An {*off*} shall be represented nominally as less than –45 dBm for the full period of a PHY-symbol time and shall be transmitted as specified in 16.7.3.3.

Fig 16-3
Optical Wave Shape Template

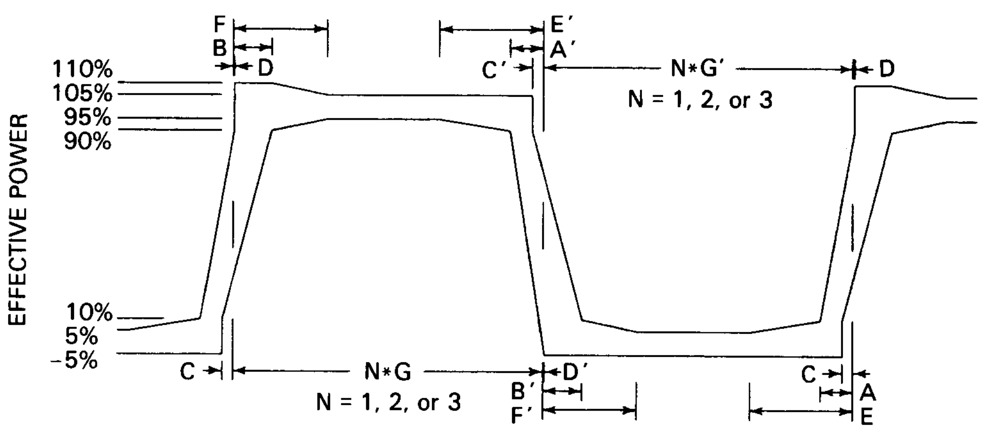

0% EFFECTIVE POWER IS THE {*L*} STATE POWER LEVEL
100% EFFECTIVE POWER IS THE {*H*} STATE POWER LEVEL

Transmit (in ns)				Receive (in ns)		
5 Mb/s	10 Mb/s	20 Mb/s	Parameter	5 Mb/s	10 Mb/s	20 Mb/s
10.46	5.23	2.62	A=A′	17.44	8.72	4.36
13.54	6.77	3.88	B=B′	22.56	11.28	5.64
3.5	1.75	0.88	C=C′	3.5	1.75	0.88
0.5	0.25	0.125	D=D′	0.5	0.25	0.125
33.24	16.62	8.3	E=E′	55.46	28.03	14.02
30.16	15.08	7.54	F=F′	50.86	25.43	12.22
N*100	N*50	N*25	G	N*100	N*50	N*25

NOTE: Where N can be 1, 2, or 3.

16.7.3.2 Output Power. The effective launch power shall be between –7 dBm and –11 dBm (except when transmitting *silence*). This is a measure of the effective power launched into the core of a standard test fiber from the station CPIC at a distance of 1 meter down the fiber on the cable plant side.

16.7.3.3 Transmitter Characteristics. The transmitter shall have its center wavelength between 800 and 910 nm. The power distribution shall exhibit a FWHM power not to exceed 60 nm. The transmitter extinction ratio shall be a minimum of 20:1. Average optical power measured over one PHY-symbol time shall be less than –45 dBm four MAC-symbol_times and longer after the last {*H*} or {*L*} PHY-symbol of a transmission. Average in-band optical power measured over one PHY-symbol time shall be less than –60 dBm 1 µs and longer after the last {*H*} or {*L*} PHY-symbol of a transmission.

The rise times and fall times of the transmit eye pattern shall be each less than 24% of the PHY-symbol_time, and the difference between rise and fall times shall be less than 8% of the PHY-symbol_time. Overshoot of transitions shall be less than 10% of the effective power. The transmit jitter shall not exceed ±4% of the PHY-symbol_time. Figure 16-3 illustrates the template of these parameters (transmit only) as they would appear while transmitting continuous *pad_idle* symbols.

NOTE: 1300 nm operation may be desirable in certain applications and is a subject of future study.

16.7.4 Jabber-Inhibit. Each PLE shall have self-interrupt capability to inhibit optical energy from reaching the LAN medium. Hardware within the PLE (with no external indication other than the prolonged detection of a non-*silence* condition within the transmitter) shall provide a nominal window of one-half second ±25% during which time a normal data link transmission may occur. If a transmission is in excess of this duration, the jabber-inhibit function shall operate to inhibit any further output from reaching the medium. Reset of this jabber-inhibit function is implementation-dependent.

16.7.5 Coupling to the Medium. The PLE functions are intended to operate satisfactorily over a medium consisting of fiber optic waveguides and splitter combiners in any configuration consistent with Section 17 of this document. The mechanical coupling of the station to the medium shall be by a connector as specified in 16.7.13.

16.7.6 Receiver Characteristics. The physical entity shall provide an undetected bit error rate of 10^{-9} or lower and a detected bit error rate of 10^{-8} or lower, when receiving signals from the standard test fiber whose effective, in-band power is within the specified receiver operating range. This bit error rate applies when receiving a properly transmitted frame or valid silence and conveyed by a medium conforming to the requirements of Section 17. The receiver shall operate within the above specification when receiving signals of any shape falling within the template of Fig 16.3 (receive only) and with received jitter which is less than ±10% of the PHY-symbol_time and with any extinction ratio of at least 10 to 1. The signal reported to the MAC sublayer shall be *silence* four MAC-symbol_times and longer after the last {*H*} or {*L*} PHY-symbol of a transmission whenever the absolute optical power is less than the specified silence power. Receivers shall be designed so that non-*silence* is reported to the MAC-PHY interface whenever there is optical energy present at the CPIC above the minimum required receiver operating power defined in the following subsections.

16.7.6.1 Receiver Operating Range—Moderate Sensitivity Receiver. The specified receiver sensitivity range is –11 dBm to –31 dBm effective power with an optical silence power of –40 dBm.

16.7.6.2 Receiver Operating Range—High Sensitivity Receiver. The specified receiver sensitivity range is –21 dBm to –41 dBm effective power with an optical silence power of –50 dBm.

NOTE: It is permissible to extend the receiver operating range so long as the receiver still operates within the specified range and conforms to the silence power specifications. The receiver should be labeled to indicate the guaranteed operating range.

16.7.7 Symbol Timing. During the recovery of *pad_idle* and *silence* symbols, it is permissible to vary the MAC-symbol reporting time by up to one nominal MAC-symbol_time. The time of each reported MAC-symbol shall nonetheless be within 90–210% of the nominal MAC-symbol time at all times. Further, the time of each reported MAC-symbol after *pad_idle* symbol recovery and until either silence or *bad_signal* is reported shall be within 90–110% of the nominal MAC-symbol_time.

16.7.8 Symbol Decoding. After determination of each received PHY-symbol pair, that PHY-symbol pair shall be decoded by a process inverse to that described in 16.7.2 and the decoded MAC-symbols shall be reported at the MAC-PHY interface. Whenever a signaled sequence is received for which the encoding process has no inverse, those PHY-symbol pairs shall be decoded as an appropriate number of *bad_signal* MAC-symbols and reported as such at the MAC-PHY interface.

NOTES: (1) Receivers are permitted to decode each transmitted *pad_idle* MAC-symbol as a *zero* or a *one* MAC-symbol.

(2) When measuring bit error rate, multiple errors in a frame may be counted as a single error.

16.7.9 Transmitter Enable/Disable and Received Signal Source Selection (OPTIONAL). The ability to enable and disable transmission onto the medium as directed by the management entity is recommended but optional.

The ability to select the source of received signaling, either a loopback point within the PLE or (one of) the (possibly redundant) media, as directed by the management entity, is recommended but optional. When such an option is invoked and the selected source is other than (one of) the media, transmission to all connected media shall be disabled automatically while such selection is in force.

16.7.10 Redundant Media Considerations. Embodiments of this standard that can function with redundant media are not precluded, provided that the embodiment, as delivered, functions correctly in a non-redundant environment. Where redundant media are employed, the provisions of 16.7.4 and 16.7.9 shall apply separately and independently to each single medium interface, and much of 16.7.9 shall be mandatory. Specifically, separate jabber-inhibit monitoring shall exist for each medium (although common inhibition is permissible), receiver signal source selection shall be provided capable of selecting any one of the redundant media, and it shall be possible to enable or disable each single transmitter independently of all other redundant transmitters when the source of received signaling is one of the redundant media.

16.7.11 Reliability. The PLE shall be designed so that its probability of causing a communication failure among other stations connected to the medium is less than 10^{-6} per hour of operation. Connectors and other components whose failure may cause interruption of the fiber optic media shall be designed to minimize the probability of total network failure.

16.7.12 Regenerative Repeater Considerations. The PLE of a regenerative repeater can be considered to be a composite entity, with separate electrical and mechanical transmit and receive functions for each connected trunk segment (that is, each port), that all work under a common encoding, decoding, timing-recovery and control function.

In performing its repeating functions, the regenerative repeater serves as a relay station. When a PHY-symbol other than *off* is received, the composite PLE determines which trunk carried the PHY-symbol, and it then selects that trunk as the source of reported PHY-symbols. Concurrently, the MAC entity begins to retransmit onto the other trunks all MAC-symbols reported. When a collision or noise is detected (for example, *bad_signal* is reported), the repeater's MAC entity sends an abort sequence (see 4.1.8) in lieu of the received MAC-symbols.

The basic mode of operation, originating or repeating, shall be determined by the superior MAC entity and conveyed by the PHY-MODE invoke primitive (see 8.2.3).

When originating, the repeater PLE shall originate the symbol timing provided to the MAC entity and transmit the encoded MAC-symbols onto all connected trunk segments. The repeater PLE shall use either loopback or any one of the attached trunks as the source of PHY-symbols that are decoded and reported by way of the PHY-DATA indicate primitive.

When repeating and when switching to repeating, the repeater PLE shall delay for an implementation-dependent amount of time (typically a few symbol-durations) to prevent repeating the reflection of the just-prior transmission, and shall then scan the connected ports for one on which signaling is being received. During the delay period, and while this scan for signal is unsuccessful, the repeater PLE shall indicate silence symbols to its MAC-entity using its locally originated symbol timing. Upon detecting signaling at one or more ports, the repeater entity shall select one of those active ports as the source of its received signaling. It shall then temporarily disable that selected port's transmitter function, decode the received signaling, retime that signaling before or after decoding, and indicate the decoded MAC-symbols to its associated MAC entity. It should then vary the frequency of the MAC-symbol timing, within the bounds of this subsection (16.7), as necessary to maintain the proper relationship with the frequency of the received PHY-symbol timing.

In summary, when the MAC entity is originating

(1) The PLE alone shall determine the MAC-symbol timing.

(2) Transmission occurs on all attached trunks (unless disabled by the provisions of 16.7.9).

(3) Loopback or any one of the attached trunks can be used as the source of PHY-symbols that are decoded and reported by way of the PHY-DATA indicate primitive.

When the MAC entity is repeating

(a) The PLE first delays long enough to ensure that the prior transmission is not repeated, then scans all attached trunks for signaling and selects one of those trunks with signaling as the source of received signaling.
(b) Transmission to the selected trunk is temporarily inhibited.
(c) The received signaling from the selected trunk is decoded and indicated to the MAC entity.
(d) The frequency of MAC-symbol timing is varied as necessary (at most ±0.015%) to track the frequency of the peer transmitter's MAC-symbol timing until the MAC entity requests transmission of *silence*, and then repeats the whole procedure.

16.7.13 Connectors. The CPIC is specified in accordance with the intermateability geometry given in ISO/IEC JTC1 DIS 9314-3 [17]. No specifications are given for non-CPIC connectors. The receptacle portion of the CPIC is considered part of the station. This receptacle uses key B for high sensitivity receivers and key A for moderate sensitivity receivers. A station that functions in either operating range uses the universal key. This CPIC can either be integral to the station or can be mounted at the end of a pigtail that is a part of the station.

NOTE: Steps have been undertaken by the IEC to standardize the CPIC. At the time this standard was written, a draft document IEC SC 86B (USA) 7 existed.

16.8 Environmental Specifications

16.8.1 Electromagnetic Emanation. Equipment shall comply with local and national requirements for limitation of electromagnetic interference. Where no local or national requirements exist, equipment shall comply with CISPR Publication 22 (1985) [3].

16.8.2 Safety Requirements. All stations meeting this standard shall comply with relevant local, national, and international safety codes and standards such as IEC Publication 950 (1986) [7]. The recommendations of IEC Publication 825 (1984) [6] shall be followed in determining the optical source safety and user warning requirements.

16.8.3 Electromagnetic and Electric Environment. The PLE embodiment shall meet its specifications when operating in an ambient continuous wave plane wave field of

(1) 2 V/m for other frequencies from 10 kHz through 30 MHz, and
(2) 5 V/m for other frequencies from 30 MHz through 1 GHz.

16.9 Labeling. It is recommended that each embodiment (and supporting documentation) of a PLE conformant to this standard be labeled in a manner visible to the user with at least these parameters:

(1) Data rate capabilities in Mb/s.

(2) Worst-case round trip delay this equipment introduces on a one-way transmission exchange between stations as specified in 6.1.9.
(3) Any delays that add to transmission_path_delay.
(4) Operating modes and selection capabilities as defined in 16.7.9 and 16.7.10.
(5) Whether the CPIC is keyed as high sensitivity using key B, moderate sensitivity using key A, or universal with no key.
(6) Fiber size or sizes that this equipment is capable of operating.

Additionally, when the station has multiple CPICs, the role of each such CPIC shall be clearly designated on the station in the vicinity of that CPIC.

Appendix

(This Appendix is not a part of this International Standard or of IEEE Std 802.4-1990, but is included for information only.[16])

16.10 Appendix—Alternative Fiber Optic Medium. This appendix defines parameters values for claiming conformance to the alternative test fiber. These values are specified in Table 16-2.

Table 16-2
Alternate Test Fiber Parameter Values

Parameter	As Specified in Subsection	As Specified for Alternate Test Fiber
Fiber		
Core diameter	16.7.1	50 µm
Cladding diameter	16.7.1	125 µm
Numerical aperture	16.7.1	0.21
Transmitter		
Silence transmit and {*off*} sent as	16.7.3.1 (3)	–48 dBm
Output power	16.7.3.2	–10 to –14 dBm
Silence power after 4 MAC-symbol times	16.7.3.3	–48 dBm
Silence power after 1 µs	16.7.3.3	–63 dBm
Moderate sensitivity receiver		
Operating range	16.7.6.1	–14 to –34 dBm
Silence level	16.7.6.1	–43 dBm
High sensitivity receiver		
Operating range	16.7.6.2	–24 to –44 dBm
Silence level	16.7.6.2	–53 dBm

[16]Although a strong attempt was made to align IEEE and ISO/IEC versions of this Standard, the version approved by the IEEE Standards Board included this Appendix as a normative Appendix. Therefore, this Appendix is now under consideration by IEEE as an informative Appendix.

Appendix

(This Appendix is not a part of this International Standard or of IEEE Std 802.4-1990, but is included for information only.)

16.11 Appendix—Comparison of Parameter Values of Standard and Alternative Fiber Optic Medium. This appendix tabulates the parameters values for the standard test fiber and the alternative test fiber.

The parameter values are reproduced in Table 16-3 for clarity only; the values that appear in 16.7 or Table 16-2 take precedence.

Table 16-3
Parameter Values for Standard and Alternate Test Fiber

Parameter	As Specified in Subsection	As Specified for Standard Test Fiber	As Specified for Alternate Test Fiber
Fiber			
Core diameter	16.7.1	62.5 μm	50 μm
Cladding diameter	16.7.1	125 μm	125 μm
Numerical aperture	16.7.1	0.275	0.21
Transmitter			
Silence transmit and {*off*} sent as	16.7.3.1 (3)	–45 dBm	–48 dBm
Output power	16.7.3.2	–7 to –11 dBm	–10 to –14 dBm
Silence power after 4 MAC-symbol times	16.7.3.3	–45 dBm	–48 dBm
Silence power after 1 μs	16.7.3.3	–60 dBm	–63 dBm
Moderate sensitivity receiver			
Operating range	16.7.6.1	–11 to –31 dBm	–14 to –34 dBm
Silence level	16.7.6.1	–40 dBm	–43 dBm
High sensitivity receiver			
Operating range	16.7.6.2	–21 to –41 dBm	–24 to –44 dBm
Silence level	16.7.6.2	–50 dBm	–53 dBm

17. Fiber Optic Medium

The functional, electrical, optical, and mechanical characteristics of one form of medium of this standard are specified in this section. This specification includes medium embodiments of a fiber optic media Local Area Network (LAN). The relationship of this section to other sections of this standard and to LAN specifications is illustrated in Fig 17-1. The relationship of this section to the fiber optic bus PLE and medium is illustrated in Fig 17-2.

This standard specifies the medium only insofar as necessary to ensure

(1) The interoperability of PLEs conforming to Section 16 when connected to a medium conformant to this section, and

(2) The protection of the LAN itself and those using it.

Fig 17-1
Relation to LAN Model

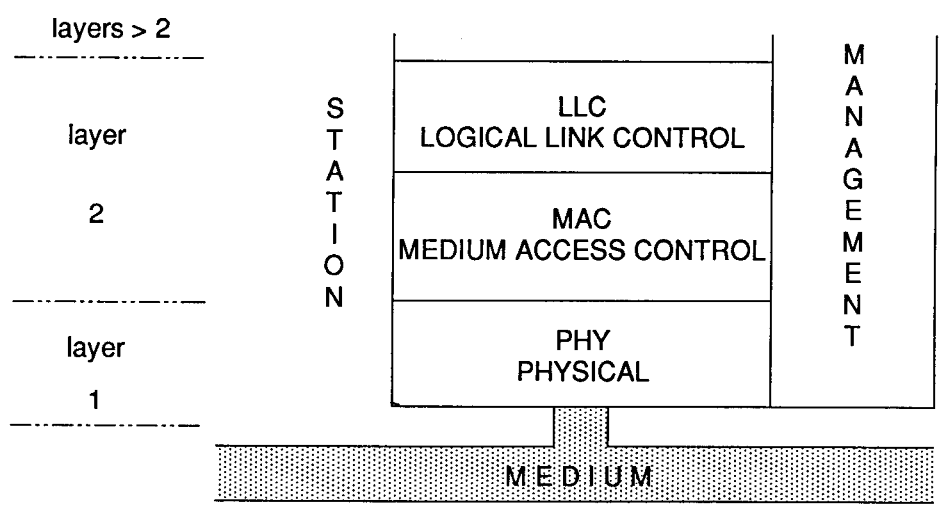

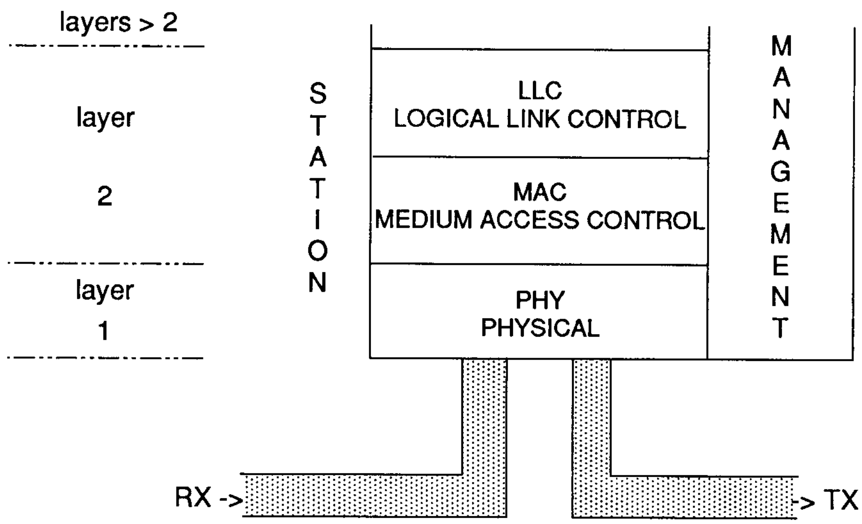

Fig 17-2
Physical Hardware Partitioning

17.1 Nomenclature. Terms used in this section whose meanings within the section are more specific than indicated by glossaries referred to in 1.2 are as follows:

cable plant interface connector (CPIC). The point at which test and conformance measurements are made; the interface between the station and the cable plant. The CPIC is keyed to correspond to the receiver operating range which is supported by the attached PLE. This connector is in the form of a duplex connector whose intermateability geometry is as defined in ISO/IEC JTC1 DIS 9314-3 [17].

center wavelength. The wavelength that is the arithmetic mean of the half-maximum spectral intensity points of the transmitter. If the spectral intensity distribution is symmetric and singly peaked, the center wavelength is at maximum intensity.

dBm. A measure of optical power. In terms of SI units (International System of Units), dBm is defined as dB (1 mW).

detected bit error. An error that is reported as *bad_signal*. *Bad_signal* reported during the preamble or during the four symbols following the last end delimiter of a transmission is not included as a bit error.

effective launch power. The effective power coupled into the core of a fiber optic waveguide by the transmitter. This power is measured with a standard test fiber connected to the CPIC (see 16.7.3.2).

extinction ratio. The ratio of the absolute optical power measured in milliwatts at the midpoint of the {*H*} PHY-symbol to the absolute optical power measured in milliwatts at the midpoint of the {*L*} PHY-symbol.

fiber optic cable. A cable containing one or more fiber optic waveguides. Jacketing material is provided to facilitate handling and to protect the fiber.

fiber optic receiver. The optics and electronics in the station that interpret the optical signal received by the station through the CPIC.

fiber optic receiver operating range. The range of effective power that must be present at the CPIC to ensure that the bit error rate specifications are met.

fiber optic transmitter. A device that emits optical energy signals for propagation onto a fiber optic waveguide through the CPIC.

fiber optic waveguide. A flexible transparent strand that is used to transport optical signals from one geographic point to another geographic point.

full width half maximum (FWHM). A measure of the width of a singly peaked function. After determining the maximum (peak) value of the function, one determines the two arguments of the function which yield half of this maximum (peak) value. The FWHM value is the magnitude of the arithmetic difference of these two arguments.

inband optical power. The optical spectral power density integrated over the range from 720–990 nm.

jitter. The offset of the 50% transition points of pulse edges from their ideal position as a result of all causes.

optical fall time. The time it takes for a pulse that is within the limits of Fig 16-3 to go from 90% effective power to 10% effective power. When expressed as a percentage, the fall time is specified as a percent of the PHY-symbol_time.

optical rise time. The time it takes for a pulse that is within the limits of Fig 16-3 to go from 10% effective power to 90% effective power. When expressed as a percentage, the rise time is specified as a percent of the PHY-symbol_time.

optical silence power. When the absolute power imposed on the CPIC is less than this power, the PLE will report *silence* to the MAC sublayer entity. This power is measured with a standard test fiber connected to the CPIC.

path loss. The difference between the launch power at the transmitting CPIC and the effective power at the receiving CPIC.

PHY-symbol_time. One half the MAC-symbol_time.

reflections. Optical energy that is inadvertently deflected from following the intended path and, hence, arrives at a later time as a spurious signal.

splitter/combiner primitive. A device that couples power received on any input port to all of its output ports. Examples of splitter/combiner primitives are active and passive star couplers.

star coupler, active. A device that couples optical power received on its input ports to all of its output ports. These devices are active and, hence, may regenerate the signals to their original power level. These devices require electrical power to operate.

star coupler, passive. A device that couples optical power received on its input ports to all of its output ports with an associated power loss that is a function of the number of ports and the excess loss parameter. These devices require no electrical power to operate.

standard test fiber. A silica fiber optic waveguide with the following nominal characteristics: core diameter of 62.5 µm, outside diameter of 125 µm, an effective numerical aperture of 0.275, and terminated in a CPIC compatible connector.

NOTE: Operation using silica fiber waveguide with the following nominal characteristics: a core diameter of 50 µm, outside diameter of 125 µm, and an effective numerical aperture of 0.21 is described in Appendix 16.10.

17.2 Object. The object of this specification is to

(1) Provide the physical medium necessary for communications between local network stations employing the token-passing bus access method defined in this standard and a fiber optic PLE,
(2) Provide for high network availability,
(3) Provide for ease of installation and service in a wide range of environments,
(4) Permit the use of commercially available components,
(5) Allow for incorporation of future advances in fiber optic technology where possible,
(6) Allow the use of low-cost solutions wherever possible, and
(7) Take advantage of the capabilities of fiber as a transmission medium and to exploit its desirable features wherever possible.

17.3 Compatibility Considerations. This standard applies to medium entities that are designed to operate with fiber optic broadcast systems. This specification applies to systems in which two-way alternate communication is accomplished through the use of fiber optic waveguides and active or passive splitter/combiner primitives in the media. This specification assumes the use of 62.5 µm fiber. Implementations using other fiber sizes are possible. For the purposes of compatibility, all measurements reference the use of the standard test fiber.

This specification assumes the use of 62.5 µm standard test fiber. Implementations which are tested using a 50 µm alternate test fiber are specified in Appendix 17.9. For the purposes of compatibility, all measurements listed in this section are referenced to the use of the standard test fiber.

All implementations of medium entities that have the same signaling rate and are conformant to this standard shall be compatible at their CPIC interface. Specific implementations based on this standard may be conceived in different ways provided compatibility at the CPIC interface is maintained.

17.4 Overview

17.4.1 General Description of Functions. The functions performed by the fiber optic medium are described informally here. Jointly, these functions provide a means whereby signaling presented at the station interfaces to the fiber optic waveguides can be conveyed to all of the stations attached to the medium.

17.4.1.1 Operational Overview. Media topologies are generated by with fiber optic waveguides and *n* by *m* splitter/combiner (see Fig 17-3). The splitter/combiner has *n* optical inputs and *m* optical outputs; the signals arriving at the inputs are distributed to the outputs. The splitter/combiner can be either active or passive. A topology is acceptable if

(1) Optical energy associated with transmission is presented only once at each CPIC receive port,

NOTE: This implies that there is effectively only one transmission path between each input and output pair of CPICs.

(2) The exiting of energy from the network occurs independent of the intelligent activity of any media element, and

(3) The line representation of PHY-symbols presented to the media is not modified by an intelligent activity of any media element.

17.4.1.2 Basic Topologies. A number of basic topologies can be built. These range in scope from a single passive star to a tree topology. The basic observations about the nature of a medium that appear in Section 5 serve as a guideline in the design of primitive elements.

17.5 Functional, Optical, and Mechanical Specifications. The fiber optic broadcast medium is an entity whose sole function is signal transport between stations of a fiber optic LAN. Consequently, only those characteristics of the medium entities that impinge on station-to-station transport, or on human and equipment safety, are specified in this standard.

Fig 17-3
***n* by *m* Splitter/Combiner Primitive**

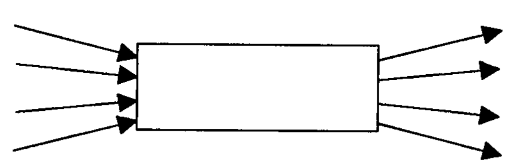

An implementation of the medium entity shall be deemed conformant to this standard if it provides the specified signal transport services and characteristics for the stations of a fiber optic LAN, and if it meets the relevant safety and environmental codes.

All measurements specified in 17.5 are to be made at the CPIC interface of station attachment to the cable plant on the receiving side of the connector (so that connector losses are attributed to the cable plant). Unless otherwise stated, all optical power levels are specified as effective, in-band optical power in dBm (dBm is defined as dB [1 mW] in 16.1).

NOTE: Standard test procedures for measuring launch power, optical silence power, and fiber optic receiver operating range are under study by EIA and IEEE 802.8, the Fiber Optic Technical Advisory Group.

17.5.1 Coupling to the Station. The PLE functions are intended to operate satisfactorily over a medium consisting of fiber optic waveguides and splitter/combiners in any configuration consistent with this section. The mechanical coupling of the medium to the station shall be by a connector as specified in 16.7.12.

17.5.2 Power Levels. When a signal that complies with the specifications of 16.7.3.2 and 16.7.3.3 is impressed on any of the medium's CPIC *transmit* ports, the signals delivered to all of the medium's CPIC *receive* ports, shall comply with the specifications of 16.7.6 and 16.7.6.1 for ports designated as *moderate sensitivity* or with 16.7.6 and 16.7.6.2 for ports designated as *high sensitivity.*

17.5.3 Data Signaling Rate. The media must present the data signaling rate with a tolerance of ±0.015% at all of the medium's CPIC receive ports when presented with a data signaling rate with a tolerance of ±0.015% on any of the medium's CPIC *transmit* ports.

17.5.4 Noise and Reflections. The power of extraneous optical energy in the receiver sensitivity band present at the interface to the medium shall not exceed –60 dBm averaged over one PHY-symbol time, when the fiber connectors are fully seated.

The reflections present at the receive port of any moderate sensitivity CPIC shall be less than –42 dBm power averaged over one PHY-symbol time. The reflections present at the receive port of any high sensitivity CPIC shall be less than –52 dBm power averaged over one PHY-symbol time.

17.5.5 Compatibility with the Stations. An embodiment of a single channel fiber optic broadcast medium entity is deemed to support a specific LAN if the requirements of 17.5.1 through 17.5.4 are met when measured from each CPIC, independent of which CPIC is chosen for test signal origination.

17.5.6 Path Loss. A path loss between 4 dB and 20 dB is required from the transmit port of any CPIC to the receive port of any moderate sensitivity CPIC. A path loss between 14 dB and 30 dB is required from the transmit port of any CPIC to the receive port of any high sensitivity CPIC. In accordance with the definition of launch power, losses due to the connector at the receive port of the CPIC shall be included in the path loss and losses at the transmit port shall be excluded.

17.5.7 Signal Distortion. When a signal complying with 16.7.3.3 is imposed on any of the medium's CPIC transmit ports, signals complying with 16.7.6 shall be received on all of the medium's CPIC receive ports. The transmit ports of active media elements shall comply with 16.7.3.3.

17.5.8 Jitter. The medium itself, consisting of fiber optic waveguides, connectors, and splitter/combiners, when it is presented with signals transmitted in accordance with Section 16, shall deliver signals at all receiving CPICs with total jitter less than ±10% of the PHY-symbol_time.

17.5.9 Signal Conveyance. The medium shall convey each transmitted PHY-symbol, including the first PHY-symbol of each transmission, to all receive ports when it is presented with signals transmitted in accordance with Section 16. Where active elements exist in the medium, detected errors should be flagged by octet-aligned abort sequences, as described in Section 4.

17.5.10 Redundancy Considerations. Where redundant media are employed, the provisions of 17.5.1 to 17.5.7 shall apply separately to each single nonredundant medium interface.

17.5.11 Reliability. All medium equipment shall be designed so that the aggregate probability of that equipment causing a communication failure at any station connected to the medium is less than 10^{-6} per hour of operation.

17.6 Safety Requirements. All media meeting this standard shall comply with relevant local, national, and international safety codes and standards such as IEC Publication 950 (1986) [8]. The recommendations of IEC Publication 825 (1984) [6] shall be adhered to in determining the optical source safety and user warning requirements.

17.7 Transmission_path_delay Considerations. When specifying an embodiment of a medium that conforms to the specification of this section, a vendor shall state the transmission_path_delay as the maximum one-way delay that the fiber optic broadcast medium could be expected to induce on a transmission from the CPIC of any connected station through the medium and any intervening fiber optic regenerative repeaters to the CPIC of any other station. The delays induced by the transmitting and receiving stations themselves should not be included in the transmission_path_delay.

For each potentially worst-case path through the medium, a path delay is computed as the sum of the medium-induced delay and repeater-induced delay, if any, in propagating a signal from one station to another. The transmission_path_delay used for determining the network's slot time (see 6.1.9) shall be the largest of these path delays for the cable system.

These path delays shall take into account all optical and electronic circuitry delays in all relevant fiber optic regenerative repeaters, combiner/splitters, etc., and all signal propagation delays within the cable segments up to the CPICs.

The transmission_path_delay shall be expressed in terms of the network's symbol signaling rate on the medium. When not an integral number of signaled symbols, it shall be rounded up to such an integral number. When uncertain of the exact value of the delay, vendors shall state an upper bound for the value.

17.8 Documentation. It is recommended that each vendor of an embodiment of a medium conformant to this standard provide to the user supporting documentation with at least these parameters:

(1) Specific sections of the standard to which the embodiment conforms,

(2) Transmission_path_delay introduced on a one way transmission exchange between CPIC's as specified in 6.1.9 and 17.7, and

(3) CPIC keyed as *high sensitivity* using key B or *moderate sensitivity* using key A.

Additionally, when the cable plant has multiple connectors, the role of each such connector shall be designated clearly by marking the station and cable plant in the vicinity of these connectors.

Appendix

(This Appendix is a binding part of IEEE Std 802.4-1990; it is not a part of this International Standard .[17])

17.9 Appendix—Alternative Fiber Optic Medium. This appendix defines parameters values for claiming conformance to the alternative test fiber. These values are specified in Table 17-1.

Table 17-1
Alternate Test Fiber Parameter Values

Parameter	As Specified in Subsection	As Specified for Alternate Test Fiber
Fiber		
Core diameter	16.7.1	50 µm
Cladding diameter	16.7.1	125 µm
Numerical aperture	16.7.1	0.21
Moderate sensitivity CPIC Allowable reflection level	17.5.4	–45 dBm
High sensitivity receiver Allowable reflection level	17.5.4	–55 dBm

[17]Although a strong attempt was made to align IEEE and ISO/IEC versions of this Standard, the version approved by the IEEE Standards Board included this Appendix as a normative Appendix. Therefore, this Appendix is now under consideration by IEEE as an informative Appendix.

Appendix

(This Appendix is not a part of this International Standard or of IEEE Std 802.4-1990, but is included for information only.)

17.10 Appendix—Comparison of Parameter Values of Standard and Alternative Fiber Optic Medium. This appendix tabulates the parameters values for the standard test fiber and the alternative test fiber.

The parameter values are reproduced in Table 17-2 for clarity only; the values that appear in 16.7 and 17.5 or Table 17-1 take precedence.

Table 17-2
Parameter Values for Standard and Alternate Test Fiber

Parameter	As Specified in Subsection	As Specified for Standard Test Fiber	As Specified for Alternate Test Fiber
Fiber			
Core diameter	16.7.1	62.5 µm	50 µm
Cladding diameter	16.7.1	125 µm	125 µm
Numerical aperture	16.7.1	0.275	0.21
Moderate sensitivity CPIC Allowable reflection level	17.5.4	–42 dBm	–45 dBm
High sensitivity CPIC Allowable reflection level	17.5.4	–52 dBm	–55 dBm

Appendix

(This Appendix is not a part of this International Standard or of IEEE Std 802.4-1990, but is included for information only.)

17.11 Appendix—Network Size and Configuration. This appendix illustrates some of the network topologies that can be constructed compliant to this section. This appendix does not imply that these are the only topologies possible.

17.11.1 Network Topology. The physical size of a network, both in geographic length and the number of desired stations, will have a significant impact on the choice of network topology. When a limited number of stations (five to ten, for instance) are required, almost any topology can be utilized. Figure 17-4 illustrates a possible topology that meets the requirement of 17.4.1.1. For situations that require a large number of stations located in widely separated geographic locations, a multi-tiered topology may be required. Topologies of this type may grow to very large numbers of stations.

In applying any of these topologies, it is the responsibility of the network designer to assure that the overall system rise time, jitter, and signal power levels are maintained and that the requirements given in 17.4.1.1 are met. Practical networks are possible using fibers that range in size from 100/140 to 50/125 μm (for example, see 17.9).

17.11.2 Loss Budgets for Passive Star Topologies. The following are two examples of loss budgets for passive star topologies. A maximum radius of 500 meters is used.

In either configuration, if the radius is increased to 1 km, an extra 3.5 dB for attenuation and 0.5 dB for dispersion should be added.

Fig 17-4
Star Topology

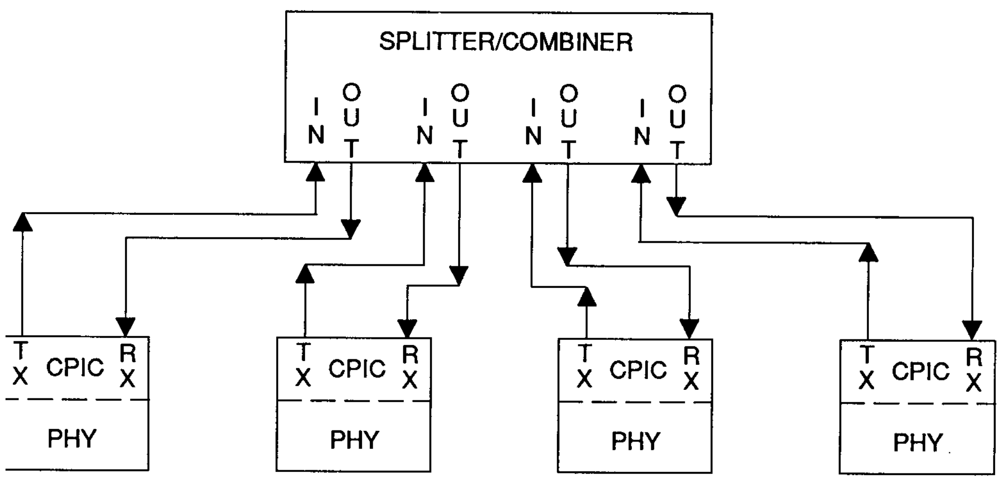

Table 17-3
Passive Star Power Loss Budget

	8-Port Passive Star	32-Port Passive Star
Receiver sensitivity	moderate	high
Fiber attenuation	3.5 dB	3.5 dB
Fiber dispersion	0.5 dB	0.5 dB
Connectors	2.0 dB	2.0 dB
Star coupler	11.0 dB	19.0 dB
System margin	3.0 dB	3.0 dB
Total	20.0 dB	28.0 dB

18. Single-Channel Phase-Continuous-FSK Bus Physical Layer Entity (PLE)

The functional, electrical, and mechanical characteristics of one specific form of PLE (single-channel phase-continuous-FSK bus) of this standard are specified in this section. This specification includes the PLE embodiments found in stations that could attach to the single-channel phase-continuous-FSK bus Local Area Network (LAN). The relationship of this section to other sections of this standard and to LAN specifications is illustrated in Fig 18-1. The relationship of this section to the single-channel phase-continuous-FSK bus PLE and medium is illustrated in Fig 18-2.

This standard specifies these PLEs only insofar as necessary to ensure

(1) The interoperability of implementations conforming to this specification, and

(2) The protection of the LAN and those using it.

Fig 18-1
Relation to LAN Model

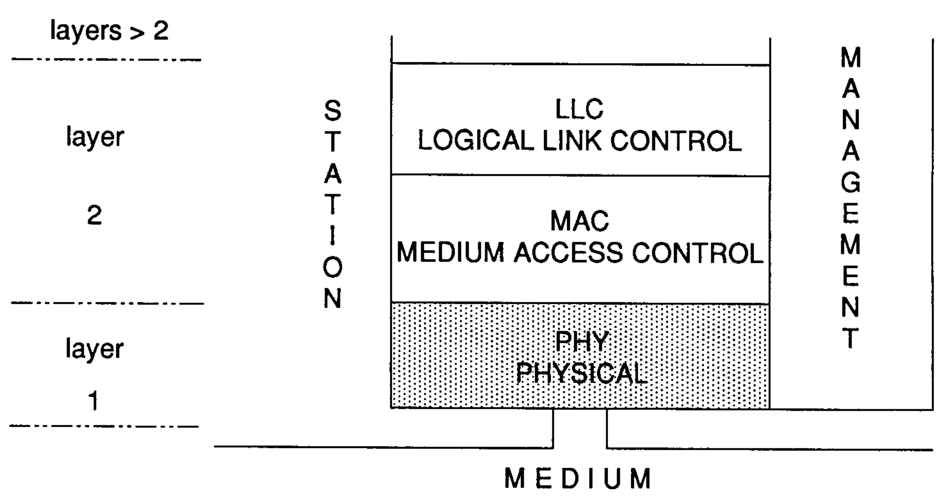

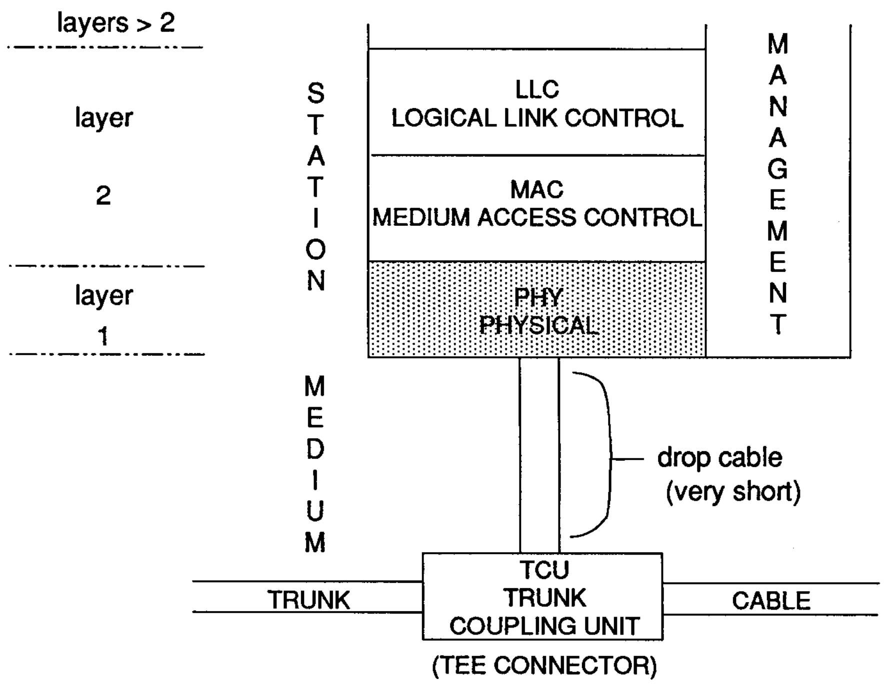

Fig 18-2
Physical Hardware Partitioning

18.1 Nomenclature. Some terms used in this section whose meanings within the section are more specific than indicated by glossaries referred to in 1.2 are as follows:

detected bit error. An error that is reported as *bad_signal*. *Bad_signal* reported during the preamble or during the four symbols following the last end delimiter of a transmission is not included.

drop cable. The very short 34–53 Ω stub cable that connects the station to a tee connector on the trunk cable.

frequency shift keying (FSK). A modulation technique whereby information is impressed upon a carrier by shifting the frequency of the transmitted signal to one of a small set of frequencies.

Manchester encoding. A means by which separate data and clock signals can be combined into a single, self-synchronizable data stream, suitable for transmission on a serial channel.

phase-continuous FSK. A particular form of FSK where the translations between signaling frequencies are accomplished by a continuous change of frequency (as opposed to the discontinuous replacement of one frequency by another, such as might be accomplished by a switch). Thus, it is also a form of frequency modulation.

regenerative repeater. A device used to extend the length, topology, or interconnectivity of a LAN beyond that imposed by the minimum transmit and receive level specifications of the station and the connectivity restrictions of the medium. Regenerative repeaters perform the basic actions of restoring signal amplitude, waveform, and timing. They also prefix enough *pad_idle* symbols to a transmission to compensate for any symbols lost in transmission from the prior station or repeater.

single-channel FSK coaxial system. A system whereby information is encoded, frequency modulated onto a carrier, and impressed on the coaxial transmission medium. At any point on the medium, only one information signal at a time can be present within the channel without disruption.

trunk cable. The main 75 Ω coaxial cable of a single-channel phase-continuous FSK coaxial cable system.

undetected bit error. An error that is not reported as *bad_signal* by the PLE.

18.2 Object. The object of this specification is to

(1) Provide the physical means necessary for communication between local network stations employing the LAN token-passing bus access method described in this standard and a single-channel phase-continuous-FSK bus medium.
(2) Define a physical interface that can be implemented independently among different manufacturers of hardware and achieve the intended level of compatibility when interconnected to a common single-channel phase-continuous-FSK bus LAN medium.
(3) Provide a communication channel capable of high bandwidth and low bit error rate performance. The resultant mean bit error rate at the MAC service interface (see Section 8) shall be less than 10^{-8}, with a mean undetected bit error rate of less than 10^{-9} at that interface.
(4) Provide for ease of installation and service in a wide range of environments.
(5) Provide for high network availability.
(6) Facilitate low-cost and low-data-rate implementations.

18.3 Compatibility Considerations. This standard applies to PLEs that are designed to operate on a 75 Ω coaxial trunk cable with the approximate configuration of an unbranched trunk, as specified in Section 19. All single-channel phase-continuous-FSK bus coaxial cable systems shall be compatible at the (drop cable) medium interface. Specific implementations based on this standard may be conceived in different ways provided compatibility at the medium is maintained.

18.4 Medium Overview. The communications medium specified in Section 19 consists of a long unbranched trunk cable that connects to stations by way of tee connectors and very short stubbed drop cables. Extension of the topology to a branched trunk usually is accomplished by way of active regenerative repeaters that are connected to span the branches.

18.5 Physical Layer Overview

18.5.1 General Description of Functions. The functions performed by the single-channel phase-continuous-FSK bus PLE are described informally here. Jointly, these functions provide a means whereby symbols presented at the MAC interface of one PLE can be conveyed to all of the PLEs on the bus for presentation to their respective MAC interfaces.

18.5.1.1 Symbol Transmission and Reception Functions. Successive symbols presented to the PLE at its MAC service interface are applied to an encoder that produces as output a pair of PHY-symbol code from the set {*H*}, {*L*}, {*off*}. The output is then applied to a two-tone FSK modulator which represents each {*H*} as the higher frequency tone, each {*L*} as the lower frequency tone, and each {*off*} as no tone. This modulated signal is then ac-coupled to the single-channel bus medium and conveyed by the medium to one or more receivers.

Each receiver is also ac-coupled to the single-channel bus medium. It bandpass filters the received signal to reduce received noise, demodulates the filtered signal, and then infers the transmitted PHY-symbol from the presence of carrier and the frequency of the received signal. It then decodes that inferred PHY-symbol by an approximate inverse of the encoding process and presents the resultant decoded MAC-symbols at its MAC service interface.

For all MAC-symbols except *pad_idle*, this decoding process is an exact inverse of the encoding process in the absence of errors. The *pad_idle* symbols, which are referred to collectively as *preamble*, are transmitted at the start of each MAC frame, to provide a training signal for receivers and to provide a nonzero minimum separation between consecutive frames. Since each transmission begins with *pad_idle* symbols, it is expected that some of these initial symbols may be "lost in transit" between the transmitter and receivers. Additionally, in phase-continuous-FSK systems, the MAC-symbol encodings for successive *pad_idles* are identical with the encoding for an alternating series of *ones* and *zeros*, and receivers are permitted to decode the transmitted representation of successive *pad_idles* as an alternating series of *ones* and *zeros* and report it as such to the MAC entity.

18.5.1.2 Regenerative Repeater Functions. Regenerative repeaters may be used to extend the network size beyond the maximum signal loss budget of an unassisted station, or the topology to that of a branched trunk. They do so by connecting two or more medium segments and repeating anything "heard" on one segment to the other segments. For the purposes of this standard, regenerative repeaters are considered stations, whether or not they have functionality beyond that of a repeater.

18.5.1.3 Jabber-Inhibit Function. To protect the LAN from most faults in a station, each station contains a jabber-inhibit function. This function serves as a "watchdog" on the transmitter; if the station does not turn off its transmitter after a prolonged time (roughly one-half second), then the transmitter output shall be automatically disabled for at least the remainder of the transmission.

18.5.1.4 Local Administrative Functions (OPTIONAL). These functions are activated either manually, or by way of the PLE's management interface, or both. They can include

(1) Enabling or disabling each transmitter output. (A redundant medium configuration has two or more transmitter outputs.)
(2) Selecting the received signal source: any single medium (if redundant media are present) or any available loopback point.

NOTE: If a loopback point is selected, then all transmitter outputs shall be inhibited.

18.5.2 Basic Functions and Options. Symbol transmission and reception functions and jabber-inhibit functions are required in all implementations. All other functions are optional.

18.6 Application of Management. The following constraints are imposed on the parameters and actions specified in Section 9:

(1) In the capabilities group, the dataRates parameter sequence shall specify the value 1.
(2) In the local management information the value of minPostSilencePreamble Length shall be 1.

18.7 Functional, Electrical, and Mechanical Specifications. Unless otherwise stated, all voltage and power levels specified are in rms and dB (1 mV, 37.5 Ω) rms, respectively, based on transmissions of arbitrary data patterns.

18.7.1 Data Signaling Rates. The standard data signaling rate for phase-continuous-FSK systems is 1 Mb/s. The permitted tolerance for this signaling rate is ±0.01% for an originating station, and is ±0.015% for a repeater station while repeating. When a composite PLE is embodied in a regenerative repeater, it shall originate signaling on all trunks at the exact same data rate.

18.7.2 Symbol Encoding. The PLE transmits MAC-symbols presented to it at its MAC interface by the MAC sublayer entity. The possible MAC-symbols are (see 8.2.1.2) *zero*, *one*, *non_data*, *pad_idle*, and *silence*. Each of these MAC-symbols is Manchester encoded into a pair of PHY-symbols from a different three-symbol {*H*}, {*L*}, {*off*} code and then transmitted. The encoding action to be taken for each of the input MAC-symbols is:

(1) *Silence*—Each *silence* symbol shall be encoded as the sequence {*off off*}.
(2) *Pad_idle*—*Pad_idle* symbols are always originated in octets. Each pair of consecutive *pad_idle* symbols shall be encoded as the sequence {*L H*} {*H L*}.
(3) *Zero*—Each *zero* symbol shall be encoded as the sequence {*H L*}.
(4) *One*—Each *one* symbol shall be encoded as the sequence {*L H*}.
(5) *Non_data*—*Non_data* symbols are transmitted by the MAC sublayer entity in pairs. Each such pair of consecutive *non_data* symbols shall be encoded as the sequence {*L L*} {*H H*} when the immediately preceding

PHY-symbol is an {*L*} and shall be encoded as the sequence {*H H*} {*L L*} when the immediately preceding PHY-symbol is an {*H*}. Thus, the (start delimiter) subsequence

non_data non_data zero non_data non_data zero

shall be encoded as the subsequence

{*L L*} {*H H*} {*H L*} {*L L*} {*H H*} {*H L*}

and the (end delimiter) subsequence

non_data non_data one non_data non_data one

shall be encoded as the subsequence

{*L L*} {*H H*} {*L H*} {*H H*} {*L L*} {*L H*}

after an immediately preceding {*L*}, and as the subsequence

{*H H*} {*L L*} {*L H*} {*H H*} {*L L*} {*L H*}

after an immediately preceding {*H*}.

18.7.3 Modulated Line Signal (at the line output of the station). The PHY-symbols resulting from the encoding of 18.7.2 shall be applied to an FSK modulator with the result that each {*H*} shall be represented by the higher of the modulator's two signaling frequencies, each {*L*} shall be represented by the lower of the modulator's two signaling frequencies, and each {*off*} shall be represented by the absence of both carrier and modulation. The resultant modulated carrier shall be coupled to the single-channel bus medium as specified in 18.7.5.

18.7.3.1 The line signal shall correspond to an FSK signal with its carrier frequency at 5.00 MHz, varying smoothly between the two signaling frequencies of 3.75 MHz ± 80 kHz and 6.25 MHz ± 80 kHz.

18.7.3.2 Each of the PHY-symbols resulting from the encoding of 18.7.2 shall be transmitted for a period equal to one-half of the inter-arrival time of the MAC-symbols which the MAC entity presents at the MAC interface. The maximum jitter in this periodicity shall be less than 1% of that MAC-symbol inter-arrival time.

18.7.3.3 When transitioning between the two signaling frequencies, the FSK modulator shall change its frequency in a continuous and monotonic manner within 100 ns, with amplitude distortion of at most 10%.

18.7.3.4 The output level of the transmitted signal at the modulated carrier frequency into a 37.5 Ω resistive load shall be between +54 dB and +60 dB (1 mV, 37.5 Ω).

NOTE: The signal level on the cable due to a single transmitter is propagated along the very short (stub) drop cable to a tee connector and thence along the trunk cable in two directions. Thus, the two half sections of the 75 Ω trunk cable appear as a 37.5 Ω load to the transmitting entity.

18.7.3.5 When in the transmitter-off state (that is, while "transmitting" the PHY-code {*off*}), the station shall present a dc shunt impedance of 50 kΩ or more at its medium interface. When in the transmitter-on state (that is, while transmitting the PHY-codes {*H*} and {*L*}), the station shall function as a voltage source with an impedance of 38 Ω or less at its medium interface. In either state, the station shall present a maximum shunt capacitance of 25 pF at that medium interface.

18.7.3.6 The residual or leakage transmitter-off output signal (that is, while "transmitting" the PHY-code {*off*} shall be no more than –26 dB (1 mV, 37.5 Ω).

18.7.4 Jabber-Inhibit. Each PLE shall have a self-interrupt capability to inhibit modulation from reaching the LAN medium. Hardware within the PLE (with no external message other than the prolonged detection of an output-on condition within the transmitter) shall provide a nominal window of one-half second ±25% during which time a normal data link transmission may occur. If a transmission is in excess of this duration, the jabber-inhibit function shall operate to inhibit any further output from reaching the medium. Reset of this jabber-inhibit function is implementation-dependent.

18.7.5 Coupling to the Medium. The Physical Layer functions are intended to operate satisfactorily with a medium consisting of a 75 Ω bidirectional coaxial trunk cable, tee connectors, and short stubbed 34–53 Ω drop cables. The mechanical coupling of the station to the medium shall be to the 34–53 Ω drop cable by way of a 50 Ω male BNC-series connector, as specified in Section 19 and defined in IEC Publication 169-8 (1978) [4].

Both the transmitter and the receiver shall be ac-coupled to the center conductor of one of the medium's drop cables, and the breakdown voltage of that ac coupling means shall be at least 500 V ac rms at 50/60 Hz. In addition to this coupling, the shield of the coaxial cable medium shall be connected to chassis ground, and the impedance of that connection shall be less than 0.1 Ω.

18.7.6 Receiver Sensitivity and Selectivity. The PLE shall be capable of providing an undetected bit error rate of 10^{-9} or lower, and a detected bit error rate of 10^{-8} or lower, when receiving signals with a level of +24 to +60 dB (1 mV, 37.5 Ω), in a system with an in-band noise floor of +4 dB (1 mV, 37.5 Ω) or less, as measured at the point of connection to the station.

18.7.7 Symbol Timing. During the recovery of *pad_idle* and *silence* symbols, it is permissible to vary the MAC-symbol reporting period by up to one nominal MAC-symbol_time. The time of each reported MAC-symbol shall nonetheless be within 90–210% of the nominal MAC-symbol_time under all conditions. Further, the time of each reported MAC-symbol after *pad_idle* symbol recovery and until either *silence* or *bad_signal* is reported, shall be within 90–110% of the nominal MAC-symbol time.

18.7.8 Symbol Decoding. After demodulation and determination of each received signaled PHY-symbol, that PHY-symbol shall be decoded by the process inverse to that described in 18.7.2, and the decoded MAC-symbols shall be reported at the MAC interface. (As noted in 18.5.1.1, receivers are permitted to decode the transmitted representation of *pad_idle* as *zero* or *one*.) Whenever a signaled PHY-symbol sequence is received for which the encoding process has no inverse, those PHY-symbols shall be decoded as an appropriate number of *bad_signal* MAC-symbols and reported as such at the MAC interface. In such cases, the receiving entity should resynchronize the decoding process as rapidly as possible.

18.7.9 Transmitter Enable/Disable and Received Signal Source Selection (OPTIONAL). The ability to enable and disable the transmission of modulation onto the single-channel bus medium as directed by the management entity is recommended but optional.

The ability to select the source of received signaling, either a loopback point within the PLE or (one of) the (possibly redundant) media, as directed by the management entity, is recommended but optional. When such an option is invoked and the selected source is other than one of the media, transmission to all connected bus media shall be disabled automatically while such selection is in force.

18.7.10 Redundant Media Considerations. Embodiments of this standard that can function with redundant media are not precluded, provided that the embodiment, as delivered, functions correctly in a nonredundant single-cable environment. Where redundant media are employed, the provisions of 18.7.4 and 18.7.9 shall apply separately and independently to each single medium interface, and much of 18.7.9 shall be mandatory. Specifically, separate jabber-inhibit monitoring shall exist for each medium (although common inhibition is permissible), receiver signal source selection shall be provided capable of selecting any one of the redundant media, and it shall be possible to enable or disable each single transmitter independently of all other redundant transmitters when the source of received signaling is one of the redundant media.

18.7.11 Reliability. The PLE shall be designed so that its probability of causing a communication failure among other stations connected to the medium is less than 10^{-6} per hour of continuous (or discontinuous) operation. For regenerative repeaters this requirement is relaxed to a probability of 10^{-5} per hour of operation. Connectors and other passive components comprising the means of connecting the station to the coaxial cable medium shall be designed to minimize the probability of total network failure.

18.7.12 Regenerative Repeater Considerations. The PLE of a regenerative repeater can be considered to be a composite entity, with separate electrical and mechanical low-level transmit and receive functions for each connected trunk segment (that is, each port), all under a common encoding, decoding, timing-recovery and control function.

The basic mode of operation, originating or repeating, shall be determined by the superior MAC entity and conveyed by the PHY-MODE invoke primitive (see 8.2.3). When originating, the repeater PLE shall originate the symbol timing provided to the MAC entity and transmit the encoded MAC-symbols onto all connected trunk segments. It shall use either internal loopback or any one of the attached trunks as the source of PHY-symbols that are decoded and reported by way of the PHY-DATA indicate primitive.

When switching to repeating, the repeater PLE shall delay for an implementation-dependent amount of time (typically, a few symbol-durations) to prevent repeating the end of the just-prior transmission, and shall then scan the connected ports for one on which signaling is being received. During the delay period, and while this scan for signal is unsuccessful, the repeater PLE shall indicate *silence* symbols to its MAC entity using its locally originated symbol timing. Upon detecting signaling at one or more ports, the repeater entity shall select one of those active ports as the source of its received signaling. It shall then temporarily disable that selected port's low-level transmitter function, decode the received signaling, and indicate the decoded MAC-symbols to its associated MAC entity. It should then vary the frequency of the MAC-symbol timing, within the bounds of this subsection (18.7), as necessary to maintain the proper relationship with the frequency of the received PHY-symbol timing.

When repeating, after decoding the MAC-symbol *silence* received from the active port, the repeater PLE shall await the MAC-entity's transmission of *silence* by way of a PHY-DATA request, shall then re-enable the temporarily disabled low-level transmitter function, and shall then again follow the procedure outlined in the prior paragraph just as if the MAC entity had just switched to repeating.

In summary, when the MAC entity is originating

(1) The PLE alone shall determine the MAC-symbol timing,
(2) Transmission occurs on all attached trunks (unless disabled by the provisions of 18.7.9),
(3) Loopback or any one of the attached trunks shall be used as the source of PHY-symbols that are decoded and reported by way of the PHY-DATA indicate primitive.

When the MAC entity is repeating

(a) The PLE first delays long enough to ensure that the prior transmission is not repeated, then scans all attached trunks for signaling and selects one of those trunks with signaling as the source of received signaling,
(b) Transmission to the selected trunk is temporarily inhibited,
(c) The received signaling from the selected trunk is decoded and indicated to the MAC entity,
(d) The frequency of MAC-symbol timing is varied as necessary (at most $\pm 0.015\%$) to track the frequency of the peer transmitter's MAC-symbol timing,
(e) Upon detecting a loss of signaling (that is, receiving *silence*) from the selected trunk, the PLE first waits until the MAC entity requests transmission of *silence*, and then repeats the whole procedure.

18.8 Environmental Specifications

18.8.1 Electromagnetic Emanation. Equipment shall comply with local and national requirements for limitation of electromagnetic interference. Where no local or national requirements exist, equipment shall comply with CISPR Publication 22 (1985) [3].

18.8.2 Safety Requirements. All stations meeting this standard shall comply with relevant local, national, and international safety codes and standards such as IEC Publication 950 (1986) [7].

18.8.3 Electromagnetic Environment. Sources of interference from the environment include electromagnetic fields, electrostatic discharge, transient voltages between earth connections, etc. Several sources of interference contribute to voltage buildup between the coaxial cable and the earth connection, if any, of the station.

The PLE embodiment shall meet its specifications when operating in an ambient plane wave field of

(1) 2 V/m from 10 kHz through 30 MHz.
(2) 5 V/m from 30 MHz through 1 GHz.

18.9 Labeling. It is recommended that each embodiment (and supporting documentation) of a PLE conformant to this standard be labeled in a manner visible to the user with at least these parameters:

(1) Data rate capability in Mb/s (that is, 1 Mb/s).
(2) Worst-case round-trip delay (for nonrepeaters) or one-way delay (for repeaters) that this equipment induces on a two-way transmission exchange between stations, as specified in 6.1.9.
(3) Operating modes and selection capabilities as defined in 18.7.9 and 18.7.10.

Additionally, when the station has multiple BNC-series connectors (for example, for redundant media), the role of each such connector shall be designated clearly by markings on the station in the vicinity of that connector.

19. Single-Channel Phase-Continuous-FSK Bus Medium

The functional, electrical, and mechanical characteristics of one specific form of medium (single-channel phase-continuous-FSK bus) of this standard are specified in this section. This specification includes the medium embodiments of a single-channel phase-continuous-FSK bus Local Area Network (LAN).

The relationship of this section to other sections of this standard and to LAN specifications is illustrated in Fig 19-1. The relationship of this section to the single-channel phase-continuous-FSK bus PLE and medium is illustrated in Fig 19-2.

This standard specifies the medium only in so far as necessary to ensure

(1) The interoperability of PLEs conforming to Section 18 when connected to a medium conformant to this section, and

(2) The protection of the LAN and those using it.

Fig 19-1
Relation to LAN Model

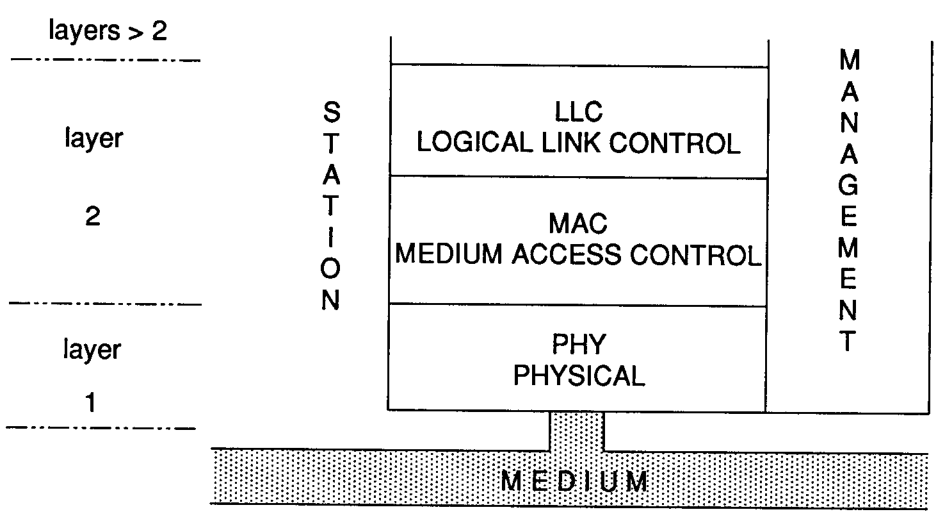

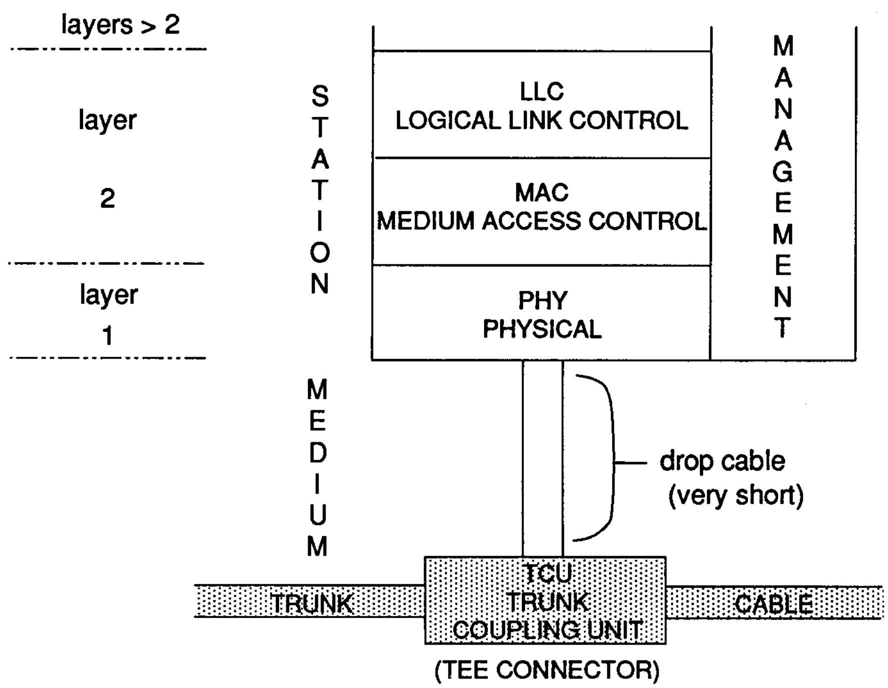

Fig 19-2
Physical Hardware Partitioning

19.1 Nomenclature. Terms used in this section whose meanings within the section are more specific than indicated glossaries referred to in 1.2 are as follows:

BNC-connector. A 50 Ω BNC-series coaxial cable connector (of the kind commonly found on RF equipment), as defined by IEC Publication 169-8 (1978) [4].

detected bit error. An error that is reported as *bad_signal*. *Bad_signal* reported during the preamble or during the four symbols following the last end delimiter of a transmission is not included.

drop cable. The very short 34–53 Ω stub coaxial cable that connects the station to a tee connector on the trunk cable.

frequency shift keying (FSK). A modulation technique whereby information is impressed upon a carrier by shifting the frequency of the transmitted signal to one of a small set of frequencies.

phase-continuous FSK. A particular form of FSK where the translations between signaling frequencies are accomplished by a continuous change of frequency (as opposed to the discontinuous replacement of one frequency by another, such as might be accomplished by a switch).

regenerative repeater. A device used to extend the length, topology, or interconnectivity of a single-channel bus LAN beyond the limits imposed by the minimum transmit and receive level specifications of the station. Regenerative repeaters perform the basic actions of restoring signal amplitude, waveform, and timing. They also prefix enough *pad_idle* symbols to a transmission to compensate for any symbols lost in transmission from the prior station or repeater.

(impedance-matching) splitter. A small module that electrically and mechanically couples one trunk cable to other trunk cables, providing a branching topology for the single-channel FSK trunk. An impedance-matching splitter combines signal energy received at its ports, splitting any signal energy received from a trunk symmetrically among the other trunks. It contains only passive electrical components (R, L, C).

single-channel FSK coaxial system. A system whereby information is encoded, frequency modulated onto a carrier, and impressed on the coaxial transmission medium. At any point on the medium, only one information signal at a time can be present within the channel without disruption.

tee connector. A small module, usually a T-shaped connector, that electrically and mechanically couples the trunk cable to a very short drop cable.

trunk cable. The main cable of a single-channel phase-continuous-FSK bus coaxial cable system.

undetected bit error. An error that is not reported as such by the PLE.

19.2 Object. The object of this specification is to

(1) Provide the physical medium necessary for communication between local network stations employing the LAN token-passing bus access method defined in this standard and a single-channel phase-continuous-FSK bus PLE,
(2) Provide for ease of installation and service in a wide range of environments,
(3) Provide for high network availability,
(4) Facilitate low-cost implementations.

19.3 Compatibility Considerations. This standard applies to medium entities that are designed to operate as nondirectional single-channel coaxial cable bus systems. Such systems generally use a long flexible trunk cable connected by way of tee connectors to very short (stubbed) drop cables to stations. This specification applies to a single trunk system in which two-way alternate communication is accomplished through the use of nondirectional tee connectors and splitters and, in large systems, multidirectional regenerative repeaters.

All implementations of medium entities conformant to this standard shall be compatible at their station interfaces. Specific implementations based on this standard may be conceived in different ways provided compatibility at the actual station interfaces is maintained.

19.4 General Overview

19.4.1 General Description of Functions. The functions performed by the single-channel phase-continuous-FSK bus medium are described informally here. Jointly, these functions provide a means whereby signaling presented at the station interfaces of very short 34–53 Ω drop cables can be combined and conveyed to all of the stations on all of the medium's drop cables. Thus stations connected to these drop cables can communicate.

19.4.1.1 Operational Overview of the Single-Channel Phase-Continuous-FSK Bus Medium. Stations are connected to the long *trunk* coaxial cable(s) of single-channel phase-continuous-FSK bus systems by very short *drop* cables and *tee connectors* (which do not impedance-match the drop cable to the trunk). These tee connectors are passive devices, usually simple connectors, which are nondirectional (that is, omnidirectional) with regard to signal propagation. The nondirectional characteristics of the tee connectors permit the station's signal to propagate in both directions along the trunk cable. The very short length of the drop cables minimizes the effect of the reflections on the drop cable due to the impedance mismatch between the 34–53 Ω drop cable and the 37.5 Ω apparent impedance of the bidirectional trunk cable.

The topology of the single-channel phase-continuous-FSK bus system is that of a very long unbranched tree trunk, with the stations connected as leaves to the very short stubs (drop cables) extending from the trunk. Branching is accomplished in the trunk itself by way of *regenerative repeaters* (described below) and *splitters*, which provide nondirectional coupling of the signaling carried on the trunk cables. The splitters employ only passive electrical components (R, L, C, only).

Regenerative repeaters provide both branching and the ability to extend a system topology beyond that permitted by the minimum transmit and receive level specifications of an unassisted station. The regenerative repeaters of the single-channel phase-continuous-FSK bus system are connected to trunk cables by tee connectors and drop cables and function as specialized stations that normally repeat signaling received on any branch to all other branches of the trunk cable system.

19.4.1.2 Regenerative Repeater Functions. In an actual single-channel phase-continuous-FSK bus system, regenerative repeaters may be used to connect trunk segments into a highly branched topology, or to extend the length of, or number of tee connectors and drops on, a trunk beyond that which the unassisted station's minimum transmit and receive level specifications would allow. Regenerative repeaters are discussed in detail in 18.7.12.

19.4.2 Basic Characteristics and Options. All signal-conveyance characteristics and station-interface characteristics are mandatory. All other characteristics are optional.

19.5 Functional, Electrical, and Mechanical Specifications. The single-channel phase-continuous-FSK bus medium entity is an entity whose sole function (relative to this standard) is signal transport between the stations of a single-channel phase-continuous-FSK bus LAN. Consequently, only those characteristics of the medium that impinge on station-to-station signal transport, or on human and equipment safety, are specified in this standard.

An implementation of the medium shall be deemed conformant to this standard if it provides the specified signal transport services and characteristics for the stations of a single-channel phase-continuous-FSK bus LAN, and if it meets the relevant safety and environmental codes.

All measurements specified in the following paragraphs are to be made at the point of station or regenerative repeater connection to the medium. Unless otherwise stated, all voltage and power levels specified are in rms and dB (1 mV, 37.5 Ω) rms, respectively, based on transmissions of arbitrary data patterns.

19.5.1 Coupling to the Station. The connection of the single-channel phase-continuous-FSK bus medium to the station shall be by way of a flexible 34–53 Ω drop cable terminated in a female BNC-series 50 Ω connector, see IEC Publication 169-8 (1978) [4]. This combination shall mate with a male BNC-series 50 Ω connector mounted on the station.

In addition to this coupling, the shield(s) of the coaxial drop cable medium shall be connected to the outer barrel of the terminating female connector and the impedance of that connection shall be less than 0.1 Ω. Also, the impedance of a connection between the outer barrel of that female connector and the shell of a mated male connector shall be less than 0.1 Ω.

NOTE: For tee-connectors that do not impedance match the drop cable to the trunk cabling, the length of the stub cable should be 350 mm or less.

19.5.2 Characteristic Impedance. The characteristic impedance of the single-channel phase-continuous-FSK bus drop cable medium shall be 34–53 Ω. The maximum VSWR at each of that medium's BNC-connectors shall be 1.5:1 or less when the BNC-connector is terminated with a 37.5 Ω resistive load.

The trunk cable shall be properly terminated at its ends. The maximum VSWR at each of those terminators shall be 1.5:1 or less when the trunk cable is terminated with a 75 Ω resistive load.

Characteristic impedance measurements shall be made over the entire single-channel phase-continuous-FSK bus cable spectrum of 3–7 MHz.

19.5.3 Signal Level. When receiving the signal of a single station or regenerative repeater whose transmit level is as specified in 18.7.3.4, the single-channel phase-continuous-FSK bus medium shall present that signaling to the connected station or regenerative repeater at an amplitude of between +24 and +60 dB (1 mV, 37.5 Ω).

19.5.4 Distortion. The maximum group delay distortion shall be 25 ns or less over the frequency range of 3–7 MHz.

19.5.5 Signal-to-Noise (S/N) Level. It is recommended that the in-band (3–7 MHz) noise floor be 0 dB (1 mV, 37.5 Ω) or less. In no case shall it be worse than +4 dB (1 mV, 37.5 Ω) as measured at the point of connection to any station or regenerative repeater.

19.5.6 Power Handling Capability. The total power over the entire cable spectrum, as presented to the station or regenerative repeater, shall be less than 0.25 W.

19.5.7 Compatibility with the Stations and Regenerative Repeaters. An embodiment of a single-channel phase-continuous-FSK bus medium is deemed to support a specific single-channel phase-continuous-FSK bus LAN if the requirements of 19.5.1 through 19.5.6 (inclusive) are met when measured from each point of station connection to the medium, independent of which one of the points of station connection is chosen for test signal origination.

19.5.8 Redundancy Considerations. As stated in 18.7.10, redundant single-channel phase-continuous-FSK bus media are not precluded from this standard. Where redundant media are employed, the provisions of 19.5.1 to 19.5.7 shall apply separately and independently to each single nonredundant medium interface.

19.5.9 Reliability. All active (powered) medium equipment shall be designed so that the aggregate probability of that equipment causing a communication failure at more than one station connected to the medium is less than 10^{-4} per hour of continuous (or discontinuous) operation. Connectors and other passive components comprising the means of connecting the station to the coaxial cable medium shall be designed to minimize the probability of total network failure.

19.6 Environmental Specifications

19.6.1 Electromagnetic Emanation. LAN cable systems shall comply with local and national requirements for limitation of electromagnetic interference. Where no local or national requirements exist, equipment shall comply with CISPR Publication 22 (1985) [3].

19.6.2 Safety Requirements. All medium meeting this standard shall comply with relevant local, national, and international safety codes and standards such as IEC Publication 950 (1986) [7].

19.6.3 Electromagnetic Environment. Sources of interference from the environment include electromagnetic fields, electrostatic discharge, transient voltages between earth connections, etc. Several sources of interference contribute to voltage buildup between the coaxial cable and the earth connection, if any, of the station.

The medium embodiment shall meet its specifications when operating in an ambient cw plane wave field of

(1) 2 V/m from 10 kHz through 30 MHz.

(2) 5 V/m from 30 MHz through 1 GHz.

19.7 Transmission_path_delay Considerations. When specifying an embodiment of a medium that conforms to the specifications of this section, a vendor shall state the transmission_path_delay as the maximum one-way delay that the single-channel phase-continuous-FSK bus medium could be expected to induce on a transmission from any connected station through any intervening regenerative repeaters to any other station. The delays induced by the transmitting and receiving stations themselves should not be included in the transmission_path_delay.

For each potentially worst-case path through the medium, a path delay is computed as the sum of the medium-induced delay and repeater-induced delay, if any, in propagating a signal from one station to another. The transmission_path_delay used for determining the network's slot_time (see 6.1.9) shall be the largest of these path delays for the cable system.

These path delay computations shall take into account all circuitry delays in all relevant regenerative repeaters and medium splitters, etc., and all signal propagation delays within the cable segments.

The transmission_path_delay shall be expressed in terms of the network's symbol signaling rate on the medium. When not an integral number of signaled symbols, it shall be rounded up to such an integral number. When uncertain of the exact value of the delay, vendors shall state an upper bound for the value.

19.8 Documentation. It is recommended that each vendor of an embodiment of a medium conformant to this standard provide to the user supporting documentation with at least these parameters:

(1) The transmission_path_delay, as specified in 19.7 and 6.1.9.

(2) Data-rate capabilities in Mb/s, where regenerative repeaters are part of the medium and so constrain the supported data rates.

19.9 Network Sizing

19.9.1 Topology Considerations. Small to medium-sized systems can be constructed using only flexible coaxial cable and tee connectors. Larger systems require both semirigid and flexible trunk cable segments, or splitters, or regenerative repeaters, or any combination of the three. Branched topologies are achieved in this medium using CATV-like splitters and regenerative repeaters.

19.9.2 Signal Loss Budget Considerations. The placement of phase-continuous-FSK regenerative repeaters and splitters shall take into account

(1) Each station's minimum transmit level and receive level specifications (see 18.7.3.4 and 18.7.6).

(2) The desired current and anticipated future placement of stations and regenerative repeaters.

(3) The presented signal level and noise floor specifications of 19.5.

Appendix

(This Appendix is not a part of this International Standard or of IEEE Std 802.4-1990, but is included for information only.)

19.10 Appendix—Guidelines for Configuring the Medium. The following recommendations for designing and installing LANs using phase-continuous-FSK transmission are the result of practical experience, and correspond to typical field conditions that are encountered.

19.10.1 Application. One of the first tasks in network design is the definition of both the current application and the possible future expanded or upgraded services. In typical situations where phase-continuous-FSK at 1 Mb/s is chosen, it is probable that planned (and unplanned) future changes will be limited to adding more drops on the existing trunk cable, and adding extensions at the ends of the cable. If the present network is small, and the possible future network is much larger, it probably will be more cost-effective, temporarily, to install a minimum network of low-cost flexible cable. Future requirements can then be based on current experience, but without any particular commitment to the current cable.

19.10.2 Network Size. Two important aspects of network size are overall length of the cable and total number of drops. Since phase-continuous-FSK PLEs (modems) simply bridge the cable by a relatively loss-free tee connection, the loss per drop is small. Thus, in typical practical networks the cable length is usually limited by cable attenuation rather than tap loss. The network's signal loss budget is then essentially equal to the dynamic range (difference between transmit and receive levels) specified in 18.7 (30 dB). Since cable loss increases with frequency, cable loss should be measured at 6.25 MHz, the highest frequency of interest.

Sections 18 and 19 recommend use of RG-6, RG-11, and semirigid (typical CATV trunk) types of 75 Ω cable. Many networks predating these specifications, but essentially conforming to them, are now installed and operating with cables of these types (and of various other 75 Ω cable types). Typical maximum network lengths achieved with several commercially available cables are the following:

Cable	Distance
RG-59	1280 m
RG-6	1600 m
RG-11 (foam)	2900 m
JT4412J	4600 m
JT4750J	7600 m

Although this specification includes the use of repeaters to extend the trunk cable beyond the length permitted by the loss budget, this has rarely been needed in practice, and such use may compromise the simplicity otherwise offered by the phase-continuous-FSK type of service.

The number of drops in the above networks has typically ranged from 2 to approximately 30. These numbers simply represent users' needs, not physical limitations. This specification does not constrain explicitly the maximum number of drops, nor is that number explicitly defined by cable parameters. The largest number of drops can be facilitated by minimizing the drop cable length.

19.10.3 Network Topology. The preferred network topology is a long unbranched trunk requiring a single trunk cable to be routed to every station site in turn. In principle, branches could be added by lossy nondirectional impedance-matching couplers, but this would usually lead to building a smaller network or requiring more expensive cable; such an approach is not recommended.

This specification does cover the use of active, regenerative repeaters for branching. Even so, the branch should be relatively long to justify its existence, and the evaluation of alternate unbranched routes is recommended first. The range of low-loss cables available makes it possible to meet most LAN needs with a completely passive medium.

19.10.4 Drop Cables. The maximum length of drop cables permitted for phase-continuous-FSK bus networks is 350 mm. Since it is not matched or terminated, the drop cable is an open stub much shorter than 1/4 wavelength, and effectively presents a shunt capacitance to the trunk cable. (The purpose of the length limitation is to limit the shunt capacitance, which causes reflections, not to limit the length per se.)

Each drop causes a small reflection on the trunk. Since the size of the reflection is proportional to the stub length, it is beneficial to keep the stub length short. Only the maximum length is specified, and the length may be zero if convenient or desirable. This minimizes reflections, and maximizes the number of drops possible in a given network.

19.10.5 Trunk Connection. The drop cable is usually coupled to the trunk cable through a simple tee connection. This is typically a matched 75 Ω connector for the trunk cable used, with a suitable adapter for the drop cable.

If use of a particular tap is permanently discontinued, the tee connector should be replaced with a corresponding straight-through (barrel) type, or the stub cable should be removed from the tee connector and replaced with a standard shielding cap.

NOTE: Since each tap introduces an impedance mismatch, spacing of taps can cause signal distortion due to accumulated reflections degrading signal-to-noise of the cable system. Network designers should consider the effects of tap placement.

19.10.6 Earthing. The trunk cable shield may be floating, single-point, or multiple-point earthed as far as signal transmission is concerned. Earths thus may be installed to eliminate electromagnetic interference (EMI) and comply with safety codes and other regulations applicable to the particular installation. This usually means grounding where the cable enters or leaves a building, and at intervals not exceeding approximately 100 m within the building. Earths should be applied carefully by a clamp that does not crush or damage the cable, because such cable damage causes serious reflections. Suitable clamps are available from suppliers of CATV system hardware.

19.10.7 Surge Protection. It is good practice to protect the cable against surges due to lightning. Suitable surge protectors that meet the requirements of IEEE C37.90.1-1989 [8] should be used at each end of the cable. The capacitive loading of surge protectors should be small to avoid affecting PLE (modem) performance, and should not exceed values permitted for a standard tap. For maximum surge protection, a low impedance, heavy-duty earth connection is required.

19.10.8 Termination. The trunk cable should be properly terminated at both ends. Shielded 75 Ω coaxial terminations with good broadband characteristics are commercially available for most coaxial cables. Since transmit levels do not exceed 60 dB (1 mV, 37.5 Ω), power ratings of 0.25 W are sufficient.

Do not terminate any drop cable. Drop cables are short unmatched stubs on the trunk cable; they must present a high dc shunt impedance to the tee connector.

19.10.9 Joining Cable Sections. In general, a trunk cable will consist of a number of separate sections of coaxial cable, each with a matched connector at each end. Some sections will be joined by the tee connectors used for taps, and others will be joined by straight-through (barrel) connectors. A good engineering practice is to maintain constant impedance between cable sections by using one cable type (such as RG-6 or RG-11) from one manufacturer for the entire trunk. This practice avoids significant reflections where cables are joined. When dissimilar types have to be joined, it is suggested that a lossy (attenuating) impedance-matched connector, such as a tap, be employed to reduce repeated reflections. In order to minimize discontinuities, cable should be obtained from the same lot (same extruder at about the same time of manufacture). It is the extruder-to-extruder differences that contribute to the mismatch. A minimal-loss splice should only be used between two pieces of cable from the same roll, or at worst from the same lot.

19.10.10 Pretested Cable. It is a good practice to pretest all trunk cable before installation. The objectives are to ensure that the attenuation does not exceed the expected values at frequencies of interest, and to ensure that concealed (that is, internal) discontinuities that can cause reflections do not exist. For a nominal charge, most cable suppliers will pretest or certify all cable before shipment. On-site testing after installation is also recommended, since any damage may degrade operating margins or cause outright failure. A recommended method for testing the installed cable for damage, improper termination, shorts, or discontinuities is to use a time domain reflectometer, which is available from various instrument manufacturers.

Appendix A
A Model Used for the Service Specification

(This Appendix is not a part of this International Standard or of IEEE Std 802.4-1990, but is included for information only.)

A1. Service Hierarchy. The services of a layer are the capabilities that it offers to a user in the next higher layer. In order to provide its service, a layer builds its functions on the services that it requires from the next lower layer. Figure A1 illustrates this notion of service hierarchy and shows the relationship of the two correspondent N-users and their associated N-layer peer protocol entities.

A2. N-Layer Interface. Services are specified by describing the information flow at the interface between the N-user and the N-layer. This information flow is modeled by discrete, instantaneous interface events, which characterize the provision of a service. Each event consists of passing a service primitive from one layer to the other through an N-layer service access point associated with an N-user. Service primitives convey the information required in providing a particular service. These service primitives are an abstraction in that they specify only the service provided rather than the means by which the service is provided. This definition of service is independent of any particular interface implementation.

Specific implementations may also include provisions for interface interactions that have no direct end-to-end effects. Examples of such local interactions include interface flow control, status requests and indications, error notifications, and layer management.

Fig A1
Service Hierarchy Relationships

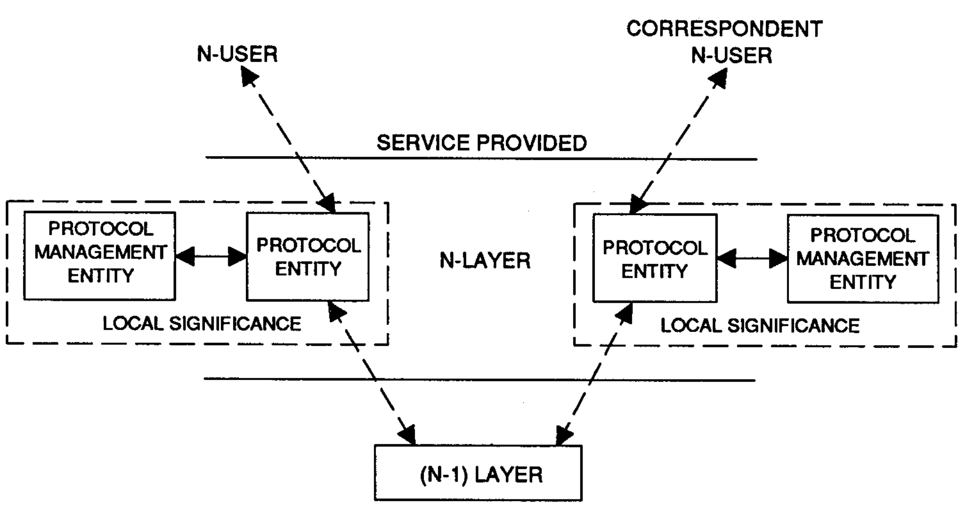

A3. Specification of Services. Services are specified here by describing the service primitives and parameters that characterize each service. A service may have one or more related primitives that constitute the interface activity that is related to the particular service. Each service primitive may have zero or more parameters which convey the information required to provide the service.

A4. Classification of N-Layer Service Primitives. Primitives are of two generic types.

A4.1 Request. The request primitive is passed from the N-layer to the (N-1) layer to request that a service be initiated.

A4.2 Indication. The indication primitive is passed from the (N-1) layer to the N-layer to indicate an internal (N-1) layer event that is significant to the N-layer. This event may be logically related to a remote service request, or may be caused by an event internal to the (N-1) layer.

A5. Interaction Behavior. Possible relationships among primitive types are illus trated by the time sequence diagrams shown in Fig A2. The figure also indicates the logical relationship of the primitive types. Primitive types that occur earlier in time and are connected by dotted lines in the diagrams are the logical antecedents of subsequent primitive types. Note that the logical and time relationship of the indication and the response primitive types are specified by the semantics of a particular service.

Fig A2
Service Primative Interactions