Recognized as an

American National Standard (ANSI)

ANSI/IEEE Std 802.5-1998E(R2003)
(Revision of ANSI/IEEE 802.5-1995)

(Adopted by ISO/IEC and redesignated as
ISO/IEC 8802-5:1995)

Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—

Part 5: Token ring access method and Physical Layer specifications

Adopted by the ISO/IEC and redesignated as
ISO/IEC 8802-5: 1998E

Sponsor

LAN MAN Standards Committee
of the
IEEE Computer Society

Abstract: This Local and Metropolitan Area Network standard, ISO/IEC 8802-5 : 1998, is part of a family of local area network (LAN) standards dealing with the physical and data link layers as defined by the ISO/IEC Open Systems Interconnection Basic Reference Model. Its purpose is to provide compatible interconnection of data processing equipment by means of a LAN using the token-passing ring access method. The frame format, including delimiters, addressing, and priority stacks, is defined. The medium access control (MAC) protocol is defined. The finite state machine and state tables are supplemented with a prose description of the algorithms. The physical layer (PHY) functions of symbol encoding and decoding, symbol time, and latency buffering are defined. The services provided by the MAC to the station management (SMT) and the services provided by the PHY to SMT and the MAC are described. These services are defined in terms of service primitives and associated parameters. The 4 and 16 Mbit/s, shielded twisted pair attachment of the station to the medium, including the medium interface connector (MIC), is also defined. The applications environment for the LAN is intended to be commercial and light industrial. The use of token ring LANs in home and heavy industrial environments, while not precluded, has not been considered in the development of the standard. A Protocol Implementation Conformance Statement (PICS) proforma is provided as an annex to the standard.

Keywords: data processing interconnection, local area network (LAN), medium access control (MAC), token ring

---

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

ISBN 1-55937-960-X

26 May 1998 *SH94536*

IEEE Standards documents are developed within the Technical Committees of the IEEE Societies and the Standards Coordinating Committees of the IEEE Standards Board. Members of the committees serve voluntarily and without compensation. They are not necessarily members of the Institute. The standards developed within IEEE represent a consensus of the broad expertise on the subject within the Institute as well as those activities outside of IEEE that have expressed an interest in participating in the development of the standard.

Use of an IEEE Standard is wholly voluntary. The existence of an IEEE Standard does not imply that there are no other ways to produce, test, measure, purchase, market, or provide other goods and services related to the scope of the IEEE Standard. Furthermore, the viewpoint expressed at the time a standard is approved and issued is subject to change brought about through developments in the state of the art and comments received from users of the standard. Every IEEE Standard is subjected to review at least every five years for revision or reaffirmation. When a document is more than five years old and has not been reaffirmed, it is reasonable to conclude that its contents, although still of some value, do not wholly reflect the present state of the art. Users are cautioned to check to determine that they have the latest edition of any IEEE Standard.

Comments for revision of IEEE Standards are welcome from any interested party, regardless of membership affiliation with IEEE. Suggestions for changes in documents should be in the form of a proposed change of text, together with appropriate supporting comments.

Interpretations: Occasionally questions may arise regarding the meaning of portions of standards as they relate to specific applications. When the need for interpretations is brought to the attention of IEEE, the Institute will initiate action to prepare appropriate responses. Since IEEE Standards represent a consensus of all concerned interests, it is important to ensure that any interpretation has also received the concurrence of a balance of interests. For this reason IEEE and the members of its technical committees are not able to provide an instant response to interpretation requests except in those cases where the matter has previously received formal consideration.

Comments on standards and requests for interpretations should be addressed to:

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

> Note: Attention is called to the possibility that implementation of this standard may require use of subject matter covered by patent rights. By publication of this standard, no position is taken with respect to the existence or validity of any patent rights in connection therewith. The IEEE shall not be responsible for identifying patents for which a license may be required by an IEEE standard or for conducting inquiries into the legal validity or scope of those patents that are brought to its attention.

Authorization to photocopy portions of any individual standard for internal or personal use is granted by the Institute of Electrical and Electronics Engineers, Inc., provided that the appropriate fee is paid to Copyright Clearance Center. To arrange for payment of licensing fee, please contact Copyright Clearance Center, Customer Service, 222 Rosewood Drive, Danvers, MA 01923 USA; (508) 750-8400. Permission to photocopy portions of any individual standard for educational classroom use can also be obtained through the Copyright Clearance Center.

Introduction to ANSI/IEEE Std 802.5, 1998 Edition

(This introduction is not a part of ANSI/IEEE Std 802.5, 1998 Edition or of ISO/IEC 8802-5 : 1998.)

This standard is part of a family of standards for local and metropolitan area networks. The relationship between the standard and other members of the family is shown below. (The numbers in the figure refer to IEEE standard numbers.)

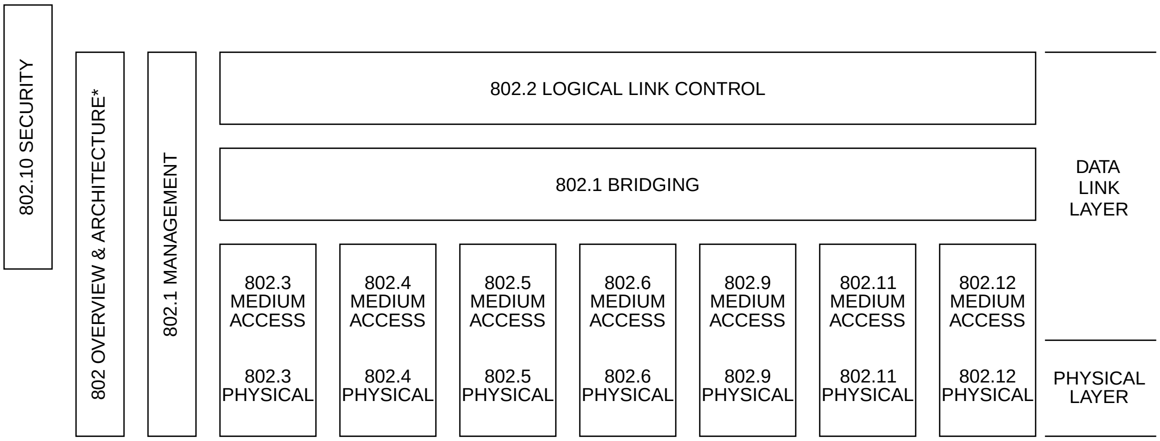

* Formerly IEEE Std 802.1A.

This family of standards deals with the Physical and Data Link layers as defined by the International Organization for Standardization (ISO) Open Systems Interconnection (OSI) Basic Reference Model (ISO/IEC 7498-1 : 1994). The access standards define seven types of medium access technologies and associated physical media, each appropriate for particular applications or system objectives. Other types are under investigation.

The standards defining the technologies noted above are as follows:

• IEEE Std 802	*Overview and Architecture.* This standard provides an overview to the family of IEEE 802 Standards.
• ANSI/IEEE Std 802.1B and 802.1k [ISO/IEC 15802-2]	*LAN/MAN Management.* Defines an OSI management-compatible architecture, and services and protocol elements for use in a LAN/MAN environment for performing remote management.
• ANSI/IEEE Std 802.1D [ISO/IEC 10038]	*Media Access Control (MAC) Bridges.* Specifies an architecture and protocol for the interconnection of IEEE 802 LANs below the MAC service boundary.
• ANSI/IEEE Std 802.1E [ISO/IEC 15802-4]	*System Load Protocol.* Specifies a set of services and protocol for those aspects of management concerned with the loading of systems on IEEE 802 LANs.
• ANSI/IEEE Std 802.1G [ISO/IEC 15802-5]	*Remote Media Access Control (MAC) Bridging.* Specifies extensions for the interconnection, using non-LAN communication technologies, of geographically separated IEEE 802 LANs below the level of the logical link control protocol.
• ANSI/IEEE Std 802.2 [ISO/IEC 8802-2]	*Logical Link Control*
• ANSI/IEEE Std 802.3 [ISO/IEC 8802-3]	*CSMA/CD Access Method and Physical Layer Specifications*

ANSI/IEEE Std 802.4 [ISO/IEC 8802-4]	*Token Passing Bus Access Method and Physical Layer Specifications*
ANSI/IEEE Std 802.5 [ISO/IEC 8802-5]	*Token Ring Access Method and Physical Layer Specifications*
ANSI/IEEE Std 802.6 [ISO/IEC 8802-6]	*Distributed Queue Dual Bus Access Method and Physical Layer Specifications*
ANSI/IEEE Std 802.9 [ISO/IEC 8802-9]	*Integrated Services (IS) LAN Interface at the Medium Access Control (MAC) and Physical (PHY) Layers*
ANSI/IEEE Std 802.10	*Interoperable LAN/MAN Security*
IEEE Std 802.11 [ISO/IEC DIS 8802-11]	*Wireless LAN Medium Access Control (MAC) and Physical Layer Specifications*
ANSI/IEEE Std 802.12 [ISO/IEC DIS 8802-12]	*Demand Priority Access Method, Physical Layer and Repeater Specifications*

In addition to the family of standards, the following is a recommended practice for a common Physical Layer technology:

IEEE Std 802.7	*IEEE Recommended Practice for Broadband Local Area Networks*

The following additional working group has authorized standards projects under development:

IEEE 802.14	*Standard Protocol for Cable-TV Based Broadband Communication Network*

Conformance test methodology

An additional standards series, identified by the number 1802, has been established to identify the conformance test methodology documents for the 802 family of standards. Thus the conformance test documents for 802.3 are numbered 1802.3.

ANSI/IEEE Std 802.5 [ISO/IEC 8802-5 : 1998]

This standard specifies that each octet of the information field shall be transmitted most significant bit (MSB) first. This convention is reversed from that used in the CSMA/CD and Token Bus standards, which are least significant bit (LSB) first transmission. While the transmission of MSB first is used for token ring, this does not imply that MSB transmission is preferable.

The IEEE 802.5 Working Group maintains a web site including notices to implementors and up-to-date information, at **http://stdsbbs.ieee.org/groups/802/5/**.

Participants

Voting members of the IEEE 802.5 Working Group who participated in developing this Third Edition of the standard were as follows:

Robert D. Love, *Chair*
David W. Wilson, *Technical Editor*

Jim Carlo
Jeremy Curtis
William Douglass
Paul Gessert
Martin Gurthrie
Sharam Hakimi
Michael Hanrahan
Neil Jarvis
Ivar Jeppesen
Allen Kasey
Ken Kutzler
George Lin
Joseph Mazor
Kelly P. McClellan
John Messenger
John E. Montague
Ivan Oakley
Syou-Chin Peng
Kirk Preiss
Bob Ross
Tony Rowell
Bill Sarles
Steve Scandalis
Showi-min Shen
Michael Siegel
Kevin G. Smith
Trevor Warwick
Kenneth T. Wilson
Edward C. Wong

The following persons were on the balloting committee:

William B. Adams
Don Aelmore
Paul Amer
Jack S. Andresen
Kit Athul
William E. Ayen
Frederic Bauchot
Manuel J. Betancor[*]
Kathleen L. Briggs
Peter K. Campbell
James T. Carlo
David E. Carlson
Alan M. Chambers
Robert S. Crowder
Christos Douligeris
Sourav K. Dutta
Paul S. Eastman
Philip H. Enslow
Changxin Fan
John W. Fendrich
Michael A. Fischer
Harvey A. Freeman
Robert J. Gagliano
Gautam Garai
Julio Gonzalez Sanz
Martin Guthrie
Kenneth C. Heck
Richard J. Iliff
Neil A. Jarvis
Henry D. Keen
Peter M. Kelly
Gary C. Kessler
Yongbum Kim
Stephen B. Kruger
Lanse M. Leach
Walter Levy
Randolph S. Little
Robert D. Love
Joseph G. Maley
John L. Messenger
Bennett Meyer
Richard H. Miller
David S. Millman
Warren Monroe
John E. Montague
Kinji Mori
David J. Morris
James R. Moulton
Wayne D. Moyers
Shimon Muller
Charles Oestereicher
Donal O'Mahony
Roger Pandanda
Lucy W. Person
Ronald C. Petersen
Thomas L. Phinney
Fernando Ramos
Ivan Reede
Edouard Y. Rocher
James W. Romlein
Floyd E. Ross
Michael Salzman
Williams F. Sarles
Stephen Scandalis
Norman Schneidewind
Rich Seifert
Lee A. Sendelbach
Donald A. Sheppard
Joseph S. Skorupa
Rosemary Slager
Fred J. Strauss
Efstathios D. Sykas
Geoffrey O. Thompson
Robert C. Tripi
Mark-Rene Uchida
Yun-Che Wang
Frank J. Weisser
Raymond P. Wenig
Paul A. Willis
David W. Wilson
Qian-Li Yang

[*] Deceased

ISO/IEC 8802-5 : 1998 [ANSI/IEEE Std 802.5, 1998 Edition] was approved by the American National Standards Institute (ANSI) on 15 April 1998.

Contents

ANNEXES

Information technology—
Telecommunications and information exchange between systems—
Local and metropolitan area networks—
Specific requirements—

Part 5: Token ring access method and Physical Layer specifications

1. Overview

1.1 Scope

For the purpose of compatible interconnection of data processing equipment via a local area network (LAN) using the token ring access method, this part of ISO/IEC 8802

a) Provides a general description of the token ring local area network (LAN) architecture (clause 2);

b) Defines the frame format, including the delimiters, address fields, information field, and frame-check sequence (FCS). Defines the Medium Access Control (MAC) frames, timers, and error counters (clause 3);

c) Defines the MAC protocols including finite state machines and state tables (clause 4);

d) Defines the system level Physical layer (PHY) signaling specifications that are specific to a ring station (clause 5);

e) Defines the managed objects necessary to manage the service and protocol elements that are involved in the management of a token ring station (clause 6);

f) Defines the PHY station attachment specification for 4 and 16 Mbit/s operation. This includes the transmitter, receiver, medium interface connector, and transmission channel for both shielded twisted pair (STP) and unshielded twisted pair (UTP) media (clause 7);

g) Defines the concentrator, incorporating multiple trunk coupling units (TCUs), for the attachment of a group of stations to the ring (clause 8);

h) Includes the protocol implementation conformance statement (PICS) proforma in compliance with the relevant requirements, and in accordance with the relevant guidance, given in ISO/IEC 9646-2 : 1994[1] (annex A);

i) Includes channel design examples and formulas for calculating cabling and concentrator system configurations (annex B);

j) Describes jitter components and provides an example of jitter buildup using a phase lock loop recovery circuit (annex C);

k) Provides informative transmitter filter design example (annex D);

l) Provides recommended guidelines for safety and operating environments (annex E);

m) Illustrates the MAC finite state machines in a notation similar to that used in ISO/IEC 8802-5 : 1992 (annex F);

n) Describes major improvements made after the first edition of the standard (annex G);

o) Provides a sample algorithm for the parsing of MAC frames (annex H);

p) Provides recommendations for the use of token ring access priorities to support multimedia traffic (annex I);

q) Provides bit error rate (BER) criteria for lobe media testing (annex P).

A particular emphasis of this standard is to specify the externally visible characteristics needed for interconnection compatibility, while avoiding unnecessary constraints upon and changes to internal design and implementation of the heterogeneous processing equipment to be interconnected.

The applications environment for the LAN is intended to be commercial and light industrial. The use of token ring LANs in home and heavy industrial environments, while not precluded, has not been considered in the development of this standard.

This standard, the Third Edition of this part of ISO/IEC 8802, provides enhancements and corrections to the Second Edition. The Second Edition provided greater specificity and improved clarity to the First Edition (1992-06-12) to ensure interoperability of the various components in the token ring network. The intent of this edition is to maintain interoperability with stations designed to this specification and stations designed to prior editions of the standard. However, interoperability with prior implementations (particularly in regard to clause 7) cannot be guaranteed due to nonspecificity within the 1992 edition. Annex G lists the specific differences between the second edition and the first edition.

The following items are subjects for future study:

a) Controlled bit altering by any device except a station.

b) Methodology to assure handling of joining of multiple rings as may be used by managed concentrators to assure normal insertion process protection mechanisms (such as duplicate address test or ring parameter server notification).

c) Ring data rate determination to allow managed data rate adaptation between stations.

d) Alternative active concentrators or repeaters that provide increased cabling distances and/or enhanced operation of rings containing devices built to the first edition of this standard, ISO/IEC 8802-5 : 1992.

e) Converters that allow the interconnection of stations on different media types.

f) Methodologies to provide enhanced transmission reliability over the trunk cable.

---

[1] Information on references can be found in 1.2.

g) Definition of concentrator managed objects.

1.2 Normative references

The following standards contain provisions which, through reference in this text, constitute provisions of this part of ISO/IEC 8802. At the time of publication, the editions indicated were valid. All standards are subject to revision, and parties to agreements based on ISO/IEC 8802-5 : 1998 are encouraged to investigate the possibility of applying the most recent editions of the standards indicated below. Members of ISO and IEC maintain registers of currently valid International Standards.

CISPR Publication 22 : 1985, Limits and Methods of Measurement of Radio Interference Characteristics of Information Technology Equipment.[2]

IEC 6073 : 1996, Basic and safety principles for man-machine interface, marking and identification—Coding principles for indicating devices and actuators.[3]

IEC 60603-7 : 1996, Connectors for frequencies below 3 MHz for use with printed boards—Part 7: Detail specification for connectors, 8-way, including fixed and free connectors with common mating features, with assessed quality.

IEC 60950 : 1991, Safety of information technology equipment.

ISO/IEC 7498-1 : 1994, Information technology—Open Systems Interconnection—Basic Reference Model: The Basic Model.[4]

ISO/IEC 7498-4 : 1989, Information processing systems—Open Systems Interconnection—Basic Reference Model—Part 4: Management framework.

ISO/IEC 8802-2 : 1994 [ANSI/IEEE Std 802.2, 1994 Edition], Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 2: Logical link control.[5]

ISO/IEC 8824 : 1990, Information technology—Open Systems Interconnection—Specification of Abstract Syntax Notation One (ASN.1).

ISO/IEC 9646-1 : 1994, Information technology—Open Systems Interconnection—Conformance testing methodology and framework—Part 1: General concepts.

ISO/IEC 9646-2 : 1994, Information technology—Open Systems Interconnection—Conformance testing methodology and framework—Part 2: Abstract Test Suite specification.

ISO/IEC 10038 : 1993 [ANSI/IEEE Std 802.1D, 1993 Edition], Information technology—Telecommunication and information exchange between systems—Local area networks—Media access control (MAC) bridges.

---

[2] CISPR and IEC publications are available from the International Electrotechnical Commission, 3, rue de Varembé, Case Postale 131, CH-1211, Genève 20, Switzerland/Suisse. These publications are also available in the United States from the Sales Department, American National Standards Institute, 11 West 42nd Street, 13th Floor, New York, NY 10036, USA.

[3] See footnote 2.

[4] ISO/IEC publications are available from ISO, Case Postale 56, 1, rue de Varembé, CH-1211, Genève 20, Switzerland/Suisse. These publications are also available in the United States from the Sales Department, American National Standards Institute, 11 West 42nd Street, 13th Floor, New York, NY 10036, USA.

[5] ISO/IEC [ANSI/IEEE] are available from ISO. They are also available from the Institute of Electrical and Electronics Engineers, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331, USA.

ISO/IEC 10165-2 : 1992, Information technology—Open Systems Interconnection—Structure of Management Information: Definition of management information.

ISO/IEC 10742 : 1994, Information technology—Telecommunications and information exchange between systems—Elements of management information related to OSI Data Link Layer standards.

ISO/IEC 11801 : 1995, Information technology—Generic cabling for customer premises.

ISO/IEC TR 11802-2 : 1995, Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Technical reports and guidelines—Part 2: Standard Group MAC Addresses.

ISO/IEC TR 11802-4 : 1994, Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Technical reports and guidelines—Part 4: Token ring access method and physical layer specifications—Fibre optic station attachment.

ISO/IEC 15802-1 : 1995, Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Common specifications—Part 1: Medium Access Control (MAC) service definition.

ISO/IEC 15802-2 : 1995 [ANSI/IEEE Std 802.1B, 1995 Edition], Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Common specifications—Part 2: LAN/MAN management, service, protocol.

1.3 Definitions

1.3.1 abort sequence: A sequence transmitted by an originating ring station that terminates the transmission of a frame prematurely. It also causes the ring station receiving this frame to terminate the frame’s reception.

1.3.2 accumulated jitter: The jitter at a PHY entity in the ring measured against the transmit clock of the active monitor. It is the total jitter accumulated by all the stations from the active monitor to the measurement point. It is typically used to determine the required size of the elastic buffer.

1.3.3 active monitor: A station on the ring that is performing certain functions to ensure proper operation of the ring. These functions include 1) establishing clock reference for the ring; 2) assuring that a usable token is available; 3) initiating the neighbor notification cycle; 4) preventing circulating frames and priority tokens. In normal operation only one station on a ring may be the active monitor at any instance in time.

1.3.4 active retimed concentrator: A type of token ring concentrator that performs an embedded repeater function in the lobe port’s data path, thereby providing ring segment boundaries at the concentrator lobe port connector (CMIC).

1.3.5 adjusted NEXT loss: The NEXT loss in decibels of a channel plus $15\log F/F_{\mathrm{ref}}$, where F is the measured frequency and F_{ref} is a reference frequency (4 MHz at 4 Mbit/s and 16 Mbit/s). It is used to determine the NEXT to interference (NIR) ratio of a channel.

1.3.6 alignment error: Alignment error is the deviation of the recovered clock from the ideal recovered clock embedded by the transmitter. The deviation from the ideal sampling point may be caused by static timing errors in the timing recovery circuit, internal jitter generated in the timing recovery circuit, and the inability to track exactly the jitter on the received data signal.

1.3.7 backup path: Secondary transmission path in trunk cabling and concentrator, normally used for token ring signal transmission only when there is a failure on the main ring path.

1.3.8 beaconing: A ring state that occurs when a station on the ring has detected a ring failure. The frame transmitted by the station to alert the other stations on the ring of the failure is called a beacon frame.

1.3.9 bit error rate (BER): A measurement of error rate stated as a ratio of the number of bits with an error to the total number of bits passing a given point on the ring. A BER of 10^{-6} indicates that an average of one bit per million bits is in error.

1.3.10 broadcast: The act of sending a frame addressed to all stations.

1.3.11 burst4: Exactly four consecutive signal elements of the same polarity.

1.3.12 burst5: Five or more consecutive signal elements of the same polarity.

1.3.13 burst6: Six or more consecutive signal elements of the same polarity.

1.3.14 channel: The data path from any transmitting MIC to the next downstream receiving MIC.

1.3.15 claiming: A ring state that occurs when a station detects that the active monitor functions are not being performed and at least one station is contending to become active monitor.

1.3.16 concentrator: A device that contains multiple interconnected trunk coupling units (TCUs). The concentrator contains two ports, referred to as *ring in* and *ring out*, to interface trunk cable.

1.3.17 configuration report server (CRS): A function that monitors and controls the stations of the ring. It receives configuration information from the stations on the ring and either forwards it to the network manager or uses it to maintain a configuration of the ring. It can also, when requested by a network manager, check the status of stations on the ring, change operational parameters of stations on the ring, and request that a station remove itself from the ring.

1.3.18 converter: A type of repeater that converts the data signal from one media to another.

1.3.19 correlated jitter: The portion of the total jitter that is related to the data pattern. Since every PHY receives the same pattern, this jitter is correlated among all similarly configured PHYs receiving the same data pattern and therefore may grow in a systematic way along the ring. Also referred to as *pattern jitter* or *systematic jitter*.

1.3.20 crosstalk: Crosstalk is undesired energy appearing in one signal path as a result of coupling from other signal paths.

1.3.21 cumulative latency: The time it takes for a signal element to travel from the active monitor's transmitter output to its receiver input.

1.3.22 differential Manchester encoding: A signaling method used to encode clock and data bit information into bit symbols. Each bit symbol is split into two halves, or signal elements, where the second half is the inverse of the first half. A 0 bit is represented by a polarity change at the start of the bit time. A 1 bit is represented by no polarity change at the start of the bit time. Differential Manchester encoding is polarity-independent.

1.3.23 elastic buffer: A variable delay element inserted in the ring by the active monitor to ensure that ring latency remains constant when the cumulative latency changes.

1.3.24 fill: A sequence of data symbols of any combination of 0 and 1 data bits (as opposed to non-data-J and non-data-K bits) whose primary purpose is to maintain timing and spacing between frames and tokens.

1.3.25 frame: A transmission unit that carries a protocol data unit (PDU) on the ring.

1.3.26 insertion loss: The signal loss that results when a channel is inserted between a transmitter and a receiver, which is the ratio of the signal level delivered to a receiver before a channel is inserted, to the signal level after the channel is inserted.

1.3.27 jitter: The time varying phase of a pulse train relative to the phase of a reference pulse train. For the specifications in this document, jitter is usually measured as the difference in edge times of the receiver's recovered clock or transmitter data output to a reference clock or data signal, typically the preceding station's transmitter clock or data output. The specifications are measured in nanoseconds.

1.3.28 latency: The time, expressed in number of symbols, it takes for a signal to pass through a ring component. See **ring latency**; **cumulative latency**.

1.3.29 LLC frame: A token ring frame containing an LLC PDU exchanged between peer entities using the MAC services.

1.3.30 lobe cabling: The cabling used to interconnect the station MICs to the TCUs. This cabling includes all work area cabling, horizontal cabling, and patch cables. The lobe cable only carries ring signals when the station is actively connected on the ring (inserted). When the station is not inserted in the ring, the lobe cable may contain local test signals.

1.3.31 Logical Link Control (sublayer) (LLC): That part of the Data Link layer that supports media-independent data link functions and uses the services of the MAC to provide services to the network layer.

1.3.32 MAC frame: A token ring frame containing a MAC PDU exchanged between MAC entities used to convey information that is used by the MAC protocol or management of the MAC sublayer.

1.3.33 main ring path: Principal transmission path in the trunk cabling. The main ring path carries the data in the primary direction. (*Contrast with:* **backup path**.)

1.3.34 medium: The material on which the data may be transmitted. STP, UTP, and optical fibers are examples of media.

1.3.35 Medium Access Control (sublayer) (MAC): The portion of the data station that controls and mediates the access to the ring.

1.3.36 medium interface connector (MIC): A connector interface at which signal transmit and receive characteristics are specified for attaching stations and concentrators. One class of MICs is the connection between the attaching stations and the lobe cabling. A second set is the attachment interface between the concentrator and its lobes. A third set is the interface between the concentrator and the trunk cabling. Two types of connectors are specified: one for connecting to STP media and one for connecting to UTP media.

1.3.37 monitor functions: The functions that recover from various error situations and are contained in each ring station. In normal operation only one of the stations on a ring is the active monitor at any point in time. The monitor functions in all other stations on the ring ensures that the active monitor function is being performed.

1.3.38 multiple frame transmission: A transmission where more than one frame is transmitted when a token is captured.

1.3.39 near end crosstalk (NEXT): Near end crosstalk is crosstalk that is propagated in a distributed channel in the direction opposite to the direction of the propagation of the signal in the disturbing channel and is measured at or near the source of the disturbing signal.

1.3.40 passive concentrator: A type of token ring concentrator that contains no active elements in the signal path of any lobe port. Embedded repeater functions may be provided by the ring in and ring out port.

1.3.41 phantom signaling: A technique where a dc power source is superimposed on the transmit and receive signal pairs in a transparent or “phantom” fashion (7.2.1) such that its application does not affect the data bearing signals on either pair. This dc power source is normally applied to request a concentrator to insert a lobe into the ring.

1.3.42 Physical (layer) (PHY): The layer responsible for interfacing with the transmission medium. This includes conditioning signals received from the MAC for transmitting to the medium and processing signals received from the medium for sending to the MAC.

1.3.43 physical media components (PMC): The sublayer of the PHY responsible for interfacing with the transmission medium. The functions of the PMC include receive, transmit, clock recovery, and ring access control.

1.3.44 physical signaling components (PSC): The sublayer of the PHY responsible for processing the signal elements received from the ring by the PMC for sending symbols to the MAC and for conditioning the symbols received from the MAC for inclusion as signal elements in the repeated data stream to the PMC.

1.3.45 port: A signal interface provided by token ring stations, passive concentrator lobes, active concentrator lobes, or concentrator trunks that is generally terminated at a media interface connector (MIC). Ports may or may not provide physical containment of channels.

1.3.46 protocol data unit (PDU): Information delivered as a unit between peer entities that contains control information and, optionally, data.

1.3.47 protocol implementation conformance statement (PICS): A statement of which capabilities and options have been implemented for a given Open Systems Interconnection (OSI) protocol. (See annex A.)

1.3.48 purging: A ring state that occurs when the active monitor has detected a ring error and is returning the ring to an operational state by transmitting purge frames.

1.3.49 recovery: The process of restoring the ring to normal operation. When the ring is beaconing, claiming, or purging, the ring is in a state of recovery.

1.3.50 repeat: The action of receiving a bit stream (for example, frame, token, or fill) and placing it on the medium. Stations repeating the bit stream may copy it into a buffer or modify control bits as appropriate. Contrast with **transmit**.

1.3.51 repeater: Physical layer coupler of ring segments. Provides for physical containment of channels, dividing the ring into segments. A repeater can receive any valid token ring signal and retransmit it with the same characteristics and levels as a transmitting station.

1.3.52 ring error monitor (REM): A function that collects ring error data from ring stations. The REM may log the received errors, or it may analyze this data and record statistics on the errors.

1.3.53 ring in: A port that receives signals from the main ring path on the trunk cable and transmits signals to the backup path on the trunk cable, and provides connectivity to the immediate upstream ring out port.

1.3.54 ring latency: In a token ring, the time (measured in bit times) it takes for a signal to propagate once around the ring. The ring latency time includes the signal propagation delay through the ring medium plus the sum of the propagation delays through each station or other element in the data path connected to the token ring.

1.3.55 ring out: A port that transmits the output signals to the main ring path on the trunk cable and receives from the backup ring path on the trunk cable, and provides connectivity to the immediate downstream ring in port.

1.3.56 ring parameter server (RPS): A function that is responsible for initializing a set of operational parameters in ring stations on a particular ring.

1.3.57 ring segment: Section of transmission path bounded by repeaters or converters. Ring segment boundaries are critical for determining the transmission limits that apply to the devices within the segment.

1.3.58 routing information: A field, carried in a frame, used by source routing transparent bridges that provides source routing operation in a bridged LAN.

1.3.59 service data unit (SDU): Information delivered as a unit between adjacent entities that may also contain a PDU of the upper layer.

1.3.60 shielded twisted pair (STP): Normally refers to those shielded cables with individual pairs of conductors twisted, or with a group of four conductors in a quad configuration, with any characteristic impedance. When used in this document, the term specifically refers to those shielded cables whose pairs have a high-frequency characteristic impedance of 150 Ω and with two pair of conductors shielded from any other individual pairs.

1.3.61 signal element: The logical signal during one half of a bit time which may take on the values of Logic_1 or Logic_0.

1.3.62 source routing: A mechanism to route frames through a bridged LAN. Within the source routed frame, the station specifies the route that the frame will traverse.

1.3.63 standby monitor: A station on the ring that is not in active monitor mode. The function of the standby monitor in normal ring operation is to assure that an active monitor is operating.

1.3.64 station (or data station): A physical device that may be attached to a shared medium LAN for the purpose of transmitting and receiving information on that shared medium. A data station is identified by a destination address (DA).

1.3.65 static timing error: The constant part of the difference in time between the ideal sampling point for the received data and the actual sampling point.

1.3.66 station management (SMT): The conceptual control element of a station that interfaces with all of the layers of the station and is responsible for the setting and resetting of control parameters, obtaining reports of error conditions, and determining if the station should be connected to or disconnected from the medium.

1.3.67 stripping: The action of a station removing the frames it has transmitted from the ring.

1.3.68 symbol: In this standard a symbol consists of two signal elements. Four symbols are defined: data_zero, data_one, non-data_J, and non-data_K.

1.3.69 token: A signal sequence passed from station to station that is used to control access to the medium.

1.3.70 tracking error: The portion of alignment error due to failure to track receiver jitter.

1.3.71 transferred jitter: The amount of jitter in the recovered clock of the upstream PHY which is subsequently transferred to the downstream PHY which in turn is transferred to the next downstream PHY. Transferred jitter is important because each PHY must both limit the amount of jitter it generates and track the jitter delivered by the upstream PHY.

1.3.72 transmit: The action of a station generating a frame, token, abort sequence, or fill and placing it on the medium. *Contrast with:* **repeat**.

1.3.73 transparent bridging: A bridging mechanism in a bridged LAN that is transparent to the end stations.

1.3.74 trunk cable: The transmission medium for interconnection of concentrators providing a main signal path and a back-up signal path, exclusive of the lobe cabling.

1.3.75 trunk coupling unit (TCU): A device that couples a station to the main ring path. A TCU provides the mechanism for insertion of a station into the ring and removal of it from the ring.

1.3.76 uncorrelated jitter: The portion of the total jitter that is independent of the data pattern. This jitter is generally caused by noise that is uncorrelated among stations and therefore grows in a non-systematic way along the ring. Uncorrelated jitter is also called *noise jitter* or *non-systematic jitter*.

1.3.77 upstream neighbor's address (UNA): The address of the station functioning upstream from a specific station.

1.3.78 unshielded twisted pair (UTP): Normally refers to those cables with individual pairs of conductors twisted, or with a group of four conductors in a star quad configuration, with any characteristic impedance. When used in this document, the term specifically refers to those cables whose pairs have a high-frequency characteristic impedance of 100 Ω. Shielded cables with the same high-frequency characteristic impedance are included within this definition.

1.3.79 verified frame: A valid frame as defined in 4.3.2, addressed to the station, for which the information field has met the validity requirements as specified in 3.3.5.2.

1.4 Conventions

When used in this standard, the term *reserved* shall be understood to mean *reserved for future standardization*. Use of reserved values to implement functions not covered by the standard may lead to interoperability problems and therefore is not recommended.

1.5 Abbreviations and acronyms

NOTE—All acronyms and notations used in state machine descriptions are summarized in 4.2.3 and 4.3.5.

A bit = address-recognized bit
AC = access control (field)
ACR = attenuation to crosstalk ratio
AD = abort delimiter
AJ = accumulated jitter
AUJA = accumulated uncorrelated jitter alignment
AMP = active monitor present
APS = accumulated phase slope
ARC = active retiming concentrator
BER = bit error rate
BN = beacon

C bit = frame-copied bit
CATT = concentrator maximum flat attenuation
CMIC = concentrator medium interface connector
CRS = configuration report server
CSQA = concentrator square-root-frequency attenuation
CT = claim token
DA = destination address
DAT = duplicate address test
DC = destination class
DFAPS = delta filtered accumulated phase slope
E bit = error-detected bit
ED = ending delimiter
EFS = end-of-frame sequence
ETR = early token release
FA = functional address
FAJ = filtered accumulated jitter
FAI = functional address indicator
FAPS = filtered accumulated phase slope
FC = frame control (field)
FCS = frame-check sequence
FR = frame
FS = frame status (field)
FSM = finite state machine
HFEJ = high-frequency edge jitter
I bit = intermediate frame bit
IFG = interframe gap
INIT = initialization
ISI = inter-symbol interference
JTOL = jitter tolerance
JTOLX = jitter tolerance in the presence of crosstalk
LAN = local area network
LLC = logical link control (sublayer)
LTH bit = length bit
M bit = monitor bit
MA = my (station’s) address
MAC = Medium Access Control (sublayer)
MGT = management
MIC = medium interface connector
MIC_S = shielded twisted pair MIC
MIC_U = unshielded twisted pair MIC
nFAPS = normalized filtered accumulated phase slope
NDT = net delay time
NEXT = near end crosstalk
NIR = NEXT loss to insertion loss ratio
NN = neighbor notification
OSI = Open Systems Interconnection
PDU = protocol data unit
PHY = Physical (Layer)
PICS = protocol implementation conformance statement
PLL = phase locked loop
Pm = priority of queued PDU
Pr = last priority value received
PMC = physical media components
P = priority bits (of the access control field)

PSC = physical signaling components
Px = the greater value of Pm or Rr
R = reservation bits (of the access control field)
RAC = ring access control
REM = ring error monitor
RI = routing information (a field in the frame format)
RI = ring in (trunk in port on a concentrator)
RIMIC = ring in MIC
RII = routing information indicator
RO = ring out (trunk out port on a concentrator)
ROMIC = ring out MIC
RP = ring purge
RPS = ring parameter server
RPT = repeat
RQ = request
Rr = last reservation value received
RRL = receiver return loss
RS = ring station
RUA = reporting station’s upstream neighbor’s address
RX = receive
SA = source address
SC = source class
SCNR = signal to crosstalk noise ratio
SD = starting delimiter
SDU = service data unit
SFS = start-of-frame sequence
SIR = signal to interference ratio
SMP = standby monitor present
Sr = highest stacked received priority
SRT = source routing transparent
STP = shielded twisted pair
SUA = stored upstream neighbor’s address
SV = subvector
SVI = subvector identifier
SVL = subvector length
SVV = subvector value
Sx = highest stacked transmitted priority
TAM = timer, active monitor
TBR = timer, BN receive
TBT = timer, BN transmit
TCATT = total channel attenuation
TCT = timer, claim token
TCU = trunk coupling unit
TDCD = transmitter duty cycle distortion
TDOV = transmitter differential output voltage
TER = timer, error report
TID = timer, insert delay
TJR = timer, join ring
TK = token
TNT = timer, no token
TQP = timer, queue PDU
TRI = timer, ring initialization
TRH = timer, remove hold
TRP = timer, transmit ring purge

TRR = timer, return to repeat
TRW = timer, remove wait
TSL = timer, signal loss
TSM = timer, standby monitor
TVX = timer, valid transmission
TWF = timer, wire fault
TWFD = timer, wire fault delay
TX = transmit
TXRL = transmit return loss
UNA = upstream neighbor's address
UTP = unshielded twisted pair
VC = vector class
VI = vector identifier
VL = vector length

1.6 Conformance requirements—station

The supplier of a protocol implementation that is claimed to conform to this standard shall complete a copy of the PICS proforma in annex A and shall provide the information necessary to identify both the supplier and the implementation.

1.6.1 Static conformance requirements

1.6.1.1 MAC sublayer

An implementation claiming conformance to this standard shall

a) Implement the token format, the frame format, associated address formats and fields, and MAC frame vectors and subvectors as defined in 3.1, 3.2, and 3.3.
b) Support 48-bit addressing and use either a universally administered individual address or a locally administered individual address.
c) Exhibit external behavior corresponding to the timer system timing parameters as specified in 3.4.
d) Exhibit external behavior corresponding to the station policy flags as specified in 3.5.
e) Implement capabilities corresponding to the counters as defined in 3.6.
f) Recognize the first bit of the source address as the indication of the presence of the routing information field in the frame format. Note that the ability to generate or respond to frames with source routing information is optional.

1.6.1.2 PHY layer

An implementation claiming conformance to this standard shall

a) Use a data signaling rate at 4 Mbit/s or 16 Mbit/s as defined in 5.2. Implementations may support either or both data rates.
b) Encode and decode symbols as defined in 5.4 and 5.6.
c) Achieve signal timing synchronization as defined in 5.7.
d) Add latency as defined in 5.8.
e) Meet multiple station accumulated jitter requirements defined in 7.1.

f) Support either STP or UTP cabling at the MIC as defined in 7.2. Implementations may support either or both media types using the appropriate MIC connector (MIC_S for STP and MIC_U for UTP).
g) Provide phantom signaling and wire fault detection as defined in 7.2.1.
h) Meet transmit specifications as defined in 7.2.2.
i) Meet receiver jitter tolerance specifications as defined in 7.2.3.
j) Operate with the channels specified in 7.2.4.

1.6.2 Dynamic conformance requirements

An implementation claiming conformance to this standard shall perform the following actions:

a) Receive MAC and LLC frames and perform the actions defined in 4.3.
b) Receive and utilize tokens for the transmission of queued PDUs as defined in 4.3.
c) Transmit queued PDUs as frames as defined in 4.3.
d) Strip frames transmitted by the station as defined in 4.3.
e) Transmit tokens in accordance with 4.3 after a frame(s) has been transmitted.
f) Perform priority operation as defined in 4.3.
g) Queue DAT_PDU, AMP_PDU, SMP_PDU, and RQ_INIT_PDU as defined in 4.3.
h) Autonomously transmit fill and the Ring Purge, Claim Token, and Beacon frames as defined in 4.3.
i) Enter BYPASS as defined in 4.3.
j) Assume ACTIVE MONITOR STATE as defined in 4.3.
k) Assume STANDBY MONITOR STATE as defined in 4.3.

1.7 Conformance requirements—concentrator

The supplier of a protocol implementation which is claimed to conform to this standard shall complete a copy of the PICS proforma in annex A and shall provide the information necessary to identify both the supplier and the implementation.

1.7.1 Static conformance requirements

a) Support a data signaling rate at 4 Mbit/s or 16 Mbit/s as defined in 5.2. Implementations may support either or both data rates.
b) Support either STP or UTP cabling at the MIC as defined in 8.1.1 for the lobe connector. Implementations may support either or both medium types using the appropriate MIC connector (MIC_S for STP and MIC_U for UTP).
c) Support either STP or UTP cabling at the MIC as defined in 8.1.1 for the trunk connector. Implementations may support either or both media types using the appropriate MIC connector (MIC_S for STP and MIC_U for UTP).
d) Provide the RAC function specified in 8.3 for the lobe connectors. This RAC function is complementary to the station phantom signaling and wire fault detection as defined in 7.2.1.
e) Support either a passive concentrator function or an active retimed concentrator function. For a passive concentrator function, the signal attenuation and crosstalk parameters are specified in 8.4. For an active retimed concentrator function, the parameters are specified in 8.5. The active retimed concentrator function parameters are closely related to the station PHY function parameters defined in 7.2.1 through 7.2.4.

1.8 Local regulations

The supplier of a protocol implementation that is claimed to conform to this standard should meet local safety and environmental regulations. Annex E provides some limited guidance in this area.

2. General description

2.1 Architectural view

There are two important ways to view local area network design: architectural, which emphasizes the logical divisions of the system and how they fit together; and implementational, which emphasizes the actual components, their packaging, and their interconnection.

This standard presents the architectural view, emphasizing the large-scale separation of the system into two parts: the MAC of the Data Link layer and the PHY. These layers are intended to correspond closely to the lowest layers of the ISO Basic Reference Model (ISO/IEC 7498-1 : 1994). As shown in figure 1, the LLC and MAC together encompass the functions intended for the data link layer of the OSI model. The MAC to LLC service definition is specified in ISO/IEC 15802-1 : 1995.

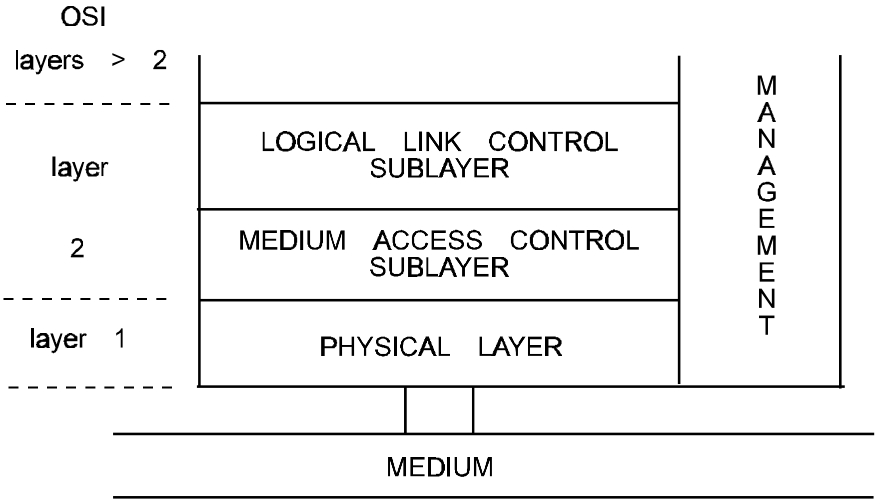

Figure 1—Relation of OSI Reference Model to the LAN Model

An architectural organization of the standard has the advantages of clarity (a clean overall division of the design along architectural lines makes the standard clearer) and flexibility (segregation of the access-method-dependent aspects of the MAC and PHY allows the LLC to apply to a variety of LAN access methods).

The exact relationship of the layers described in this standard to the layers defined by the OSI Reference Model is a subject for further study.

2.2 Station functional organization and data flow

Figure 2 is an illustration of a station’s data flow indicating which clauses of this standard address the various functions of a token ring station. The operation of a station is specified in clauses 3 through 7, and operation of the concentrator is specified in clause 8. Internal service interfaces have been defined solely for the purpose of specifying operation between the clauses and are not requirements for a physical interface.

a) The PMC/PSC internal service interface (PM_UNITDATA.request; PM_UNITDATA.indication) defines the information exchange between the physical media components (PHY) specified in clause 7 and the physical signaling components (PSC) specified in clause 5. This service interface is defined in clause 5.

b) The PSC/MAC internal service interface (PS_CONTROL.request, PS_STATUS.indication, PS_UNITDATA.indication, PS_UNITDATA.request, PS_EVENT.response) defines the control mechanism of, the mechanism for indicating the status of, and the information exchange between, the physical signaling components (PSC) specified in clause 5 and the MAC sublayer specified in clauses 3 and 4. Clause 3 defines frame formats and station facilities. Clause 4 specifies the MAC protocol that uses the formats and facilities defined in clause 3 to receive and transmit information. This service interface is defined in clause 5.

c) The MAC/PHY internal service interface (PM_CONTROL.request and PM_STATUS.indication) provides the control mechanism of the PHY functions by the MAC protocol, and the mechanism for indicating the status of the PHY functions to the MAC protocol. This service interface is defined in clause 5.

d) The MAC/LLC service interface (MA_UNITDATA.indication, MA_UNITDATA.request) is specified in ISO/IEC 15802-1 and defines the information exchange between the MAC sublayer and the LLC sublayer.

e) The MAC/Bridge service interface (M_UNITDATA.indication, M_UNITDATA.request, M_UNITDATA.response) is specified in ISO/IEC 10038 : 1993 and defines the information exchange between the MAC and the internal bridging sublayer.

f) The MAC/MGT service interface (MGT_UNITDATA.indication, MGT_UNITDATA.request, MGT_CONTROL.request, MGT_STATUS.indication) defines the control mechanism of, the mechanism for indicating the status of, and the information exchange between, the MAC protocol specified in clause 4 and management (MGT). The managed objects are specified in clause 6. The MGT_UNITDATA.indication and MGT_UNITDATA.request primitives are specified in clause 4 and are used to convey MAC management frames between the MAC and the appropriate management function (e.g., RPS, CRS, REM).

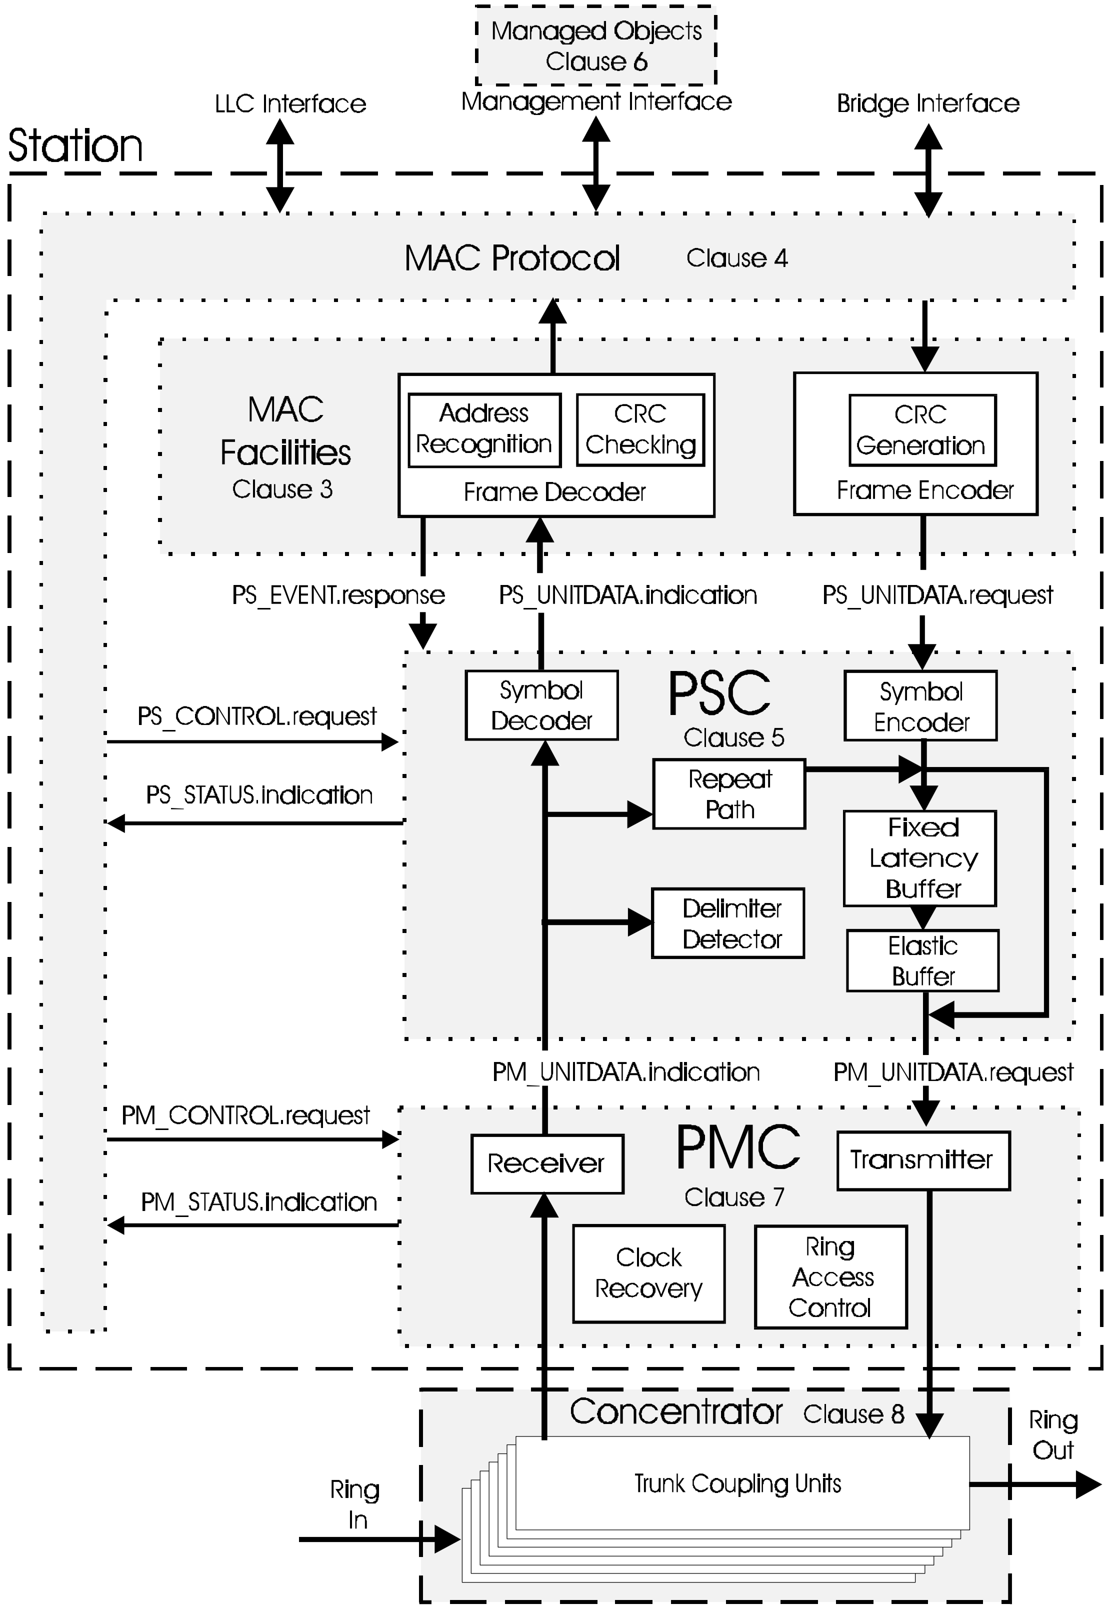

Figure 2—Station functional organization and data flow

2.3 Physical structure of a token ring network

A token ring consists of a star wired system of stations with each station connected by a lobe cable to a concentrator TCU, as shown in figure 2. A concentrator contains multiple TCUs. Concentrators are serially interconnected using the concentrator ring in and ring out ports, forming the main ring path and the back-up ring path. Interconnection on the main ring and back-up ring paths may include the use of repeaters and converters that divide the total signal path into separate segments to bound the total channel path length for each station.

Networks may include repeaters and/or converters at each trunk interface where the transmission medium is changed, i.e., between STP and fiber, between STP and UTP, between UTP and fiber. Thus, each of these interfaces becomes a ring segment boundary. Evaluation and control of jitter buildup across these segments is a subject for future study.

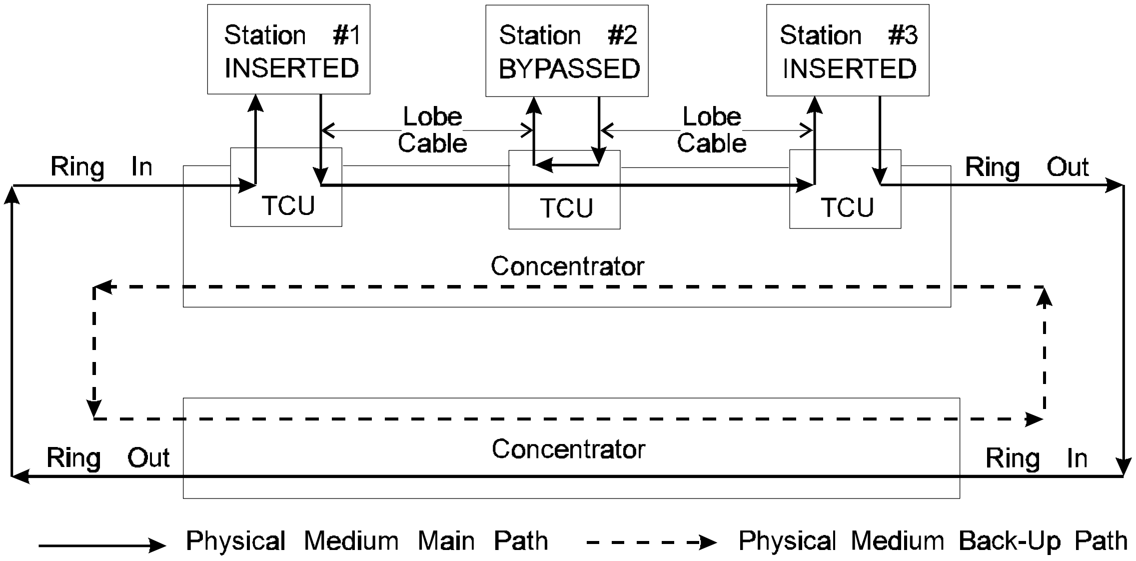

Figure 3—Typical token ring wiring example

Each TCU provides attachment (“insertion” or “bypass”) of the station into or from the trunk ring. This insertion/bypass mechanism is usually controlled by the attaching station’s ring access control (RAC) function. A station performs self-tests of the attached cable lobe and the station's MAC and PHY layer circuitry by operating a one-node network while the TCU function is in the bypass mode (e.g., Station #2 in figure 3.) When the TCU receives a dc signal from the station via the lobe cable, the TCU switches from bypass mode to inserted mode, allowing the inserting station’s MAC to receive a symbol stream from its upstream neighbor, and provide an output symbol stream to its downstream neighbor.

The maximum station count on a single token ring is limited by both MAC and PHY considerations. If the recommended system timing parameters (see 3.4) are used, operation of up to 250 stations on the ring is possible. Modifications of this station count may occur based on media type, data rate, use of repeaters or converters, and concentrator type.

Information on the token ring is transferred sequentially, bit by bit, from one inserted station to the next. In figure 3, Station #1 would transmit to Station #3 and Station #3 would transmit to Station #1. Each station generally regenerates and repeats each bit and serves as the means for attaching one or more devices (terminals, work stations) to the ring for the purpose of communicating with other devices on the network. A given station transfers information onto the ring, where the information circulates from one station to the next. The addressed destination station(s) copies the information as it passes. Finally, the station that transmitted the information strips the information from the ring.

A station gains the right to transmit its information onto the medium when it detects a token passing on the medium. The token is a control signal comprised of a unique signaling sequence that circulates on the medium following each information transfer. Any station, upon detection of an appropriate token, may capture the token by modifying it to a start of frame sequence and appending appropriate control and status fields, address fields, routing information field, information field, check sum, and the ending frame sequence. At the completion of its information transfer and after appropriate checking for proper operation, the station initiates a new token, which provides other stations the opportunity to gain access to the ring.

The maximum period of time a station may transmit on the medium before releasing the token is controlled by counter CTO (used as a timer).

Multiple levels of priority are available for independent and dynamic assignment depending upon the relative class of service required for any given message; for example, real-time voice, interactive, immediate (network recovery). The allocation of priorities is by mutual agreement among users of the network.

Error detection and recovery mechanisms are provided to restore network operation in the event that transmission errors or medium transients (for example, those resulting from station insertion or removal) cause the access method to deviate from normal operation. Detection and recovery for these cases utilize a network monitoring function that is performed in a specific station (the active monitor), with back-up capability in all other stations that are attached to the ring (standby monitors).

2.4 Data stations, servers, and system management

Each ring in a token ring network may have a set of server stations (servers) that provide a means through which a system manager manages the stations in a token ring system. Such arrangement is depicted in figure 4.

Servers are data collection points on each ring where reports from the data stations are gathered. Servers then communicate the necessary information to the system manager for the purpose of managing a token ring system.

Data stations communicate with the servers by

a) Reporting errors that are detected, such as lost token, FCS error, or lost frames;
b) Requesting operating parameters when inserting into the ring;
c) Reporting changes in configuration due to insertion or removal of stations (UNA changes);
d) Responding to requests for various status information; and
e) Removing from the ring when requested.

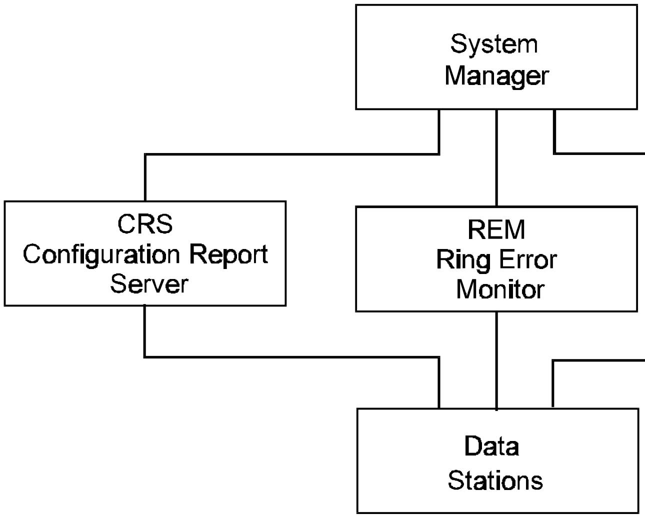

Figure 4—Relationship of data stations, servers, and system manager

The station-to-server message format and content is specified in clause 3; the protocol for message exchange is specified in clause 4.

The RPS sets MAC entity parameters and can limit access to the ring. It does this by responding to the Request Initialization MAC frame with an Initialize Station MAC frame.

The CRS receives management information about the current status and configuration of the ring, and may issue MAC frames to stations on the ring in response to a request from a management entity.

The REM receives from MAC entities, and reports to the requester, error statistics regarding ring and station operation.

Specification of the operation of the RPS, CRS, and REM is beyond the scope of this standard.

The specification of the format and protocol for the information interchange between the servers and the systems manager is not covered by this standard. However, the objects (parameters, events, and actions) specified in clause 6 are the elements of such communication.

A source routing station may send and receive frames that include a routing information field as specified in ISO/IEC 10038 : 1993. The routing information field may be determined through the use of protocols specified in ISO/IEC 8802-2 : 1994. A station that does not support the source routing information field is prohibited from setting the RII bit (first bit in the source address field).

A source routing transparent bridge copies frames based on the routing information field as well as the destination address as specified in ISO/IEC 10038 : 1993. A source routing transparent bridge may set the address-recognized bits in a frame that it copies with intent to forward. Otherwise, stations are prohibited from setting the address-recognized bits unless the frame is addressed to the station's individual address or to a group address for which the station is enabled.

3. Formats and facilities

This clause defines token ring formats (3.1), token and frame field descriptions (3.2), MAC frames (3.3), system timing parameters (3.4), station policies (3.5), and error counters (3.6).

The figures depict the formats of the fields *in the sequence they are transmitted on the medium*, with the left-most bit or symbol transmitted first.

Processes that require comparison of fields or bits perform that comparison upon those fields or bits *as depicted*, with the left-most bit or symbol compared first and, for the purpose of comparison, considered most significant.

Binary values are represented by the binary digits 0, 1 , or x, which are enclosed by B' and '. The lowercase x represents a digit that may be 0 or 1. An example of a binary notation is B'001x'.

Hexadecimal values are represented by hexadecimal digits 0–9, A–F, or x, which are enclosed by X' and '. The lowercase x represents a digit that may be any hexadecimal value. An example of a hexadecimal notation is X'01 6F'.

All values are stated most significant digit first. Spaces within a value are for clarity and may be ignored.

3.1 Formats

The symbol sequences used in token ring are the token, frame, abort sequence, and fill.

3.1.1 Token format

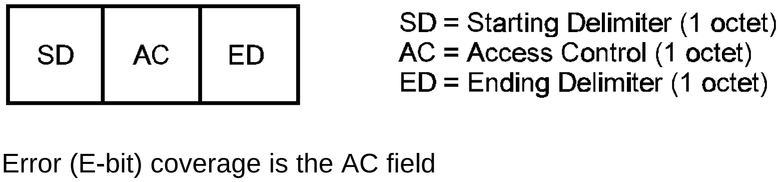

Figure 5—Token format

The token is the means by which the right to transmit (as opposed to the normal process of repeating) is passed from one station to another. The token may occur anywhere in the data stream; that is, receiving stations shall be able to detect a token on any signal element boundary.

Error checking of token formats is defined in 3.2.8.2, and validity requirements are specified in 4.3.1.

3.1.2 Frame format

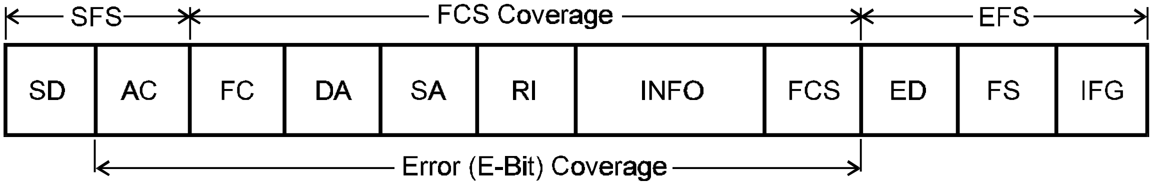

SFS = Start-of-Frame Sequence
SD = Starting Delimiter (1 octet)
AC = Access Control (1 octet)
FC = Frame Control (1 octet)
DA = Destination Address (6 octets)
SA = Source Address (6 octets)
RI = Routing Information (0 to 30 octets)
INFO = Information (0 or more octets) (see 4.2.4.1 for length limitations)
FCS = Frame Check Sequence (4 octets)
EFS = End-of-Frame Sequence
ED = Ending Delimiter (1 octet)
FS = Frame Status (1 octet)
IFG = Interframe Gap (see 3.2.10)

Figure 6—Frame format

The frame format shall be used for transmitting both MAC and LLC messages to the destination station(s). It may or may not contain an information (INFO) field. It may or may not contain a routing information (RI) field (see 3.3.5.1 regarding restrictions in certain MAC frames). The frame format may occur anywhere in the data stream; that is, receiving stations shall be able to detect a frame on any signal element boundary.

Error checking of frame formats is defined in 3.2.8.2 and validity requirements are specified in 4.3.2.

3.1.3 Abort sequence

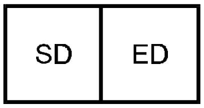

SD = Starting Delimiter (1 octet)
ED = Ending Delimiter (1 octet)

Figure 7—Abort sequence

The abort sequence is transmitted by a station when an error condition causes the station to prematurely terminate the transmission of a frame. The abort sequence causes stations receiving the frame to recognize that it is not a valid frame. The abort sequence may occur anywhere in the data stream; that is, receiving stations shall be able to detect an abort sequence on any signal element boundary.

3.1.4 Fill

A station shall transmit fill as any combination of Data_zero and Data_one symbols in accordance with the protocol described in clause 4.

3.2 Field descriptions

The following is a detailed description of the individual fields used for the abort sequence, token, and frame formats.

3.2.1 Starting delimiter (SD) field

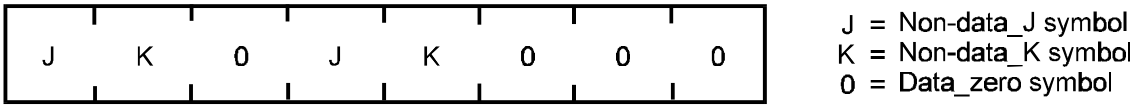

(For a discussion of non-data symbols, see 5.3.)

Figure 8—Starting delimiter field

A frame, token, or abort sequence always starts with a starting delimiter. Receiving stations shall consider the SD valid only if all eight symbols of the SD (J K 0 J K 0 0 0) are received correctly.

3.2.2 Access control (AC) field

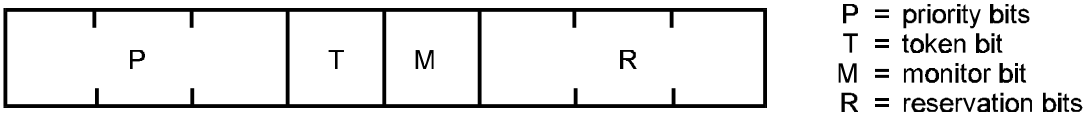

Figure 9—Access control field

3.2.2.1 Priority bits

The priority bits indicate the priority of a token and, therefore, when stations are allowed to use the token. Frames carry the same access priority as the token used to gain access. In a multiple-priority system, the priorities of tokens that a station uses depend on the priority of the PDU to be transmitted. The eight levels of priority increase from the lowest (B'000') to the highest (B'111') priority.

3.2.2.2 Token bit

The token bit is a 0 in a token and a 1 in a frame. When a station with a PDU to transmit detects a token that has a priority equal to or less than the PDU to be transmitted, it may change the token to a SFS and transmit the PDU.

3.2.2.3 Monitor (M) bit

The M bit is used to prevent a token that has a priority greater than B'000' or any frame from continuously circulating on the ring. All frames and tokens shall be transmitted with the M bit set to 0. As a token with priority greater than zero or any frame is being repeated by the active monitor, the M bit is set to 1. All other stations repeat this bit as received. If an active monitor detects a token or a frame with the M bit equal to 1, it initiates the ring purge process. During beaconing, the station downstream from the beaconing station will optionally assume this role of setting and inspecting the M bit.

3.2.2.4 Reservation bits

The reservation bits allow stations with high priority PDUs to request (in frames or tokens as they are repeated) that the next token be issued at the requested priority. The precise protocol for setting these bits is described in clause 4. The eight levels of reservation increase from B'000' to B'111'.

3.2.3 Frame control (FC) field

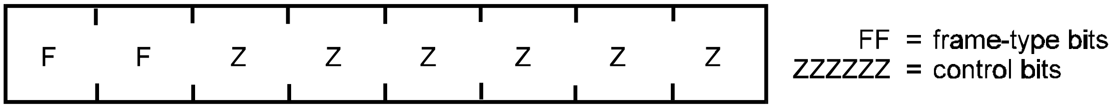

Figure 10—Frame control field

The FC field identifies the frame type and indicates MAC frame handling or LLC frame priority.

3.2.3.1 Frame-type bits

The frame-type bits indicate the type of the frame as follows:

a) B'00' = MAC frame (contains a MAC PDU)
b) B'01' = LLC frame (contains an LLC PDU)
c) B'1x' = undefined format (reserved for future standardization)

In this standard, the term MAC frame is defined as the type of frame used to convey information used by the MAC protocol or management of the MAC sublayer.

3.2.3.1.1 Undefined frame format

The value B'1x' is reserved for frame types that may be defined in the future. However, although currently undefined, any future frame formats are required to adhere to the following conditions to ensure compatibility:

a) The frame format shall be delimited by the SFS fields and the EFS fields. No additional fields may be included in or follow the EFS field.
b) The position of the AC and FC field shall be unchanged.
c) The SFS and EFS of the format shall be separated by an integral number of octets. This number shall be at least 1 (that is, the FC field) and the maximum length is subject to the constraints of the token holding time.
d) All symbols between the SD and ED shall be Data_zero and Data_one symbols.

3.2.3.2 ZZZZZZ bits

3.2.3.2.1 MAC frames

If the frame-type bits indicate a MAC frame, the ZZZZZZ control bits specify MAC frame handling (see 3.3.1 and 3.3.5).

3.2.3.2.2 LLC frames

If the frame-type bits indicate an LLC frame, the ZZZZZZ bits are designated as “rrrYYY”. The “rrr” bits are reserved for future standardization and shall be transmitted as B'000' in all transmitted frames and may be ignored upon reception. The “YYY” bits may be used to carry the user priority of the PDU from the source LLC entity to the target LLC entity or entities. Note that P (the priority in the AC field of a frame) does not convey user priority.

3.2.3.2.3 Undefined frame format

The ZZZZZZ bits are undefined and reserved for future standardization.

3.2.4 Destination address (DA) and source address (SA) fields

Each frame contains two address fields: the destination (station) address and the source (station) address, in that order. Representation of addresses used in this standard are different than the representation utilized in ISO/IEC 15802-1 : 1995.

3.2.4.1 Destination address (DA) field

The frame's DA identifies the station or stations for which the information field of the frame is intended. Included in the frame's DA are two bits that 1) indicate whether the DA is a group or individual address and 2) indicate whether the DA is a locally or universally administered address.

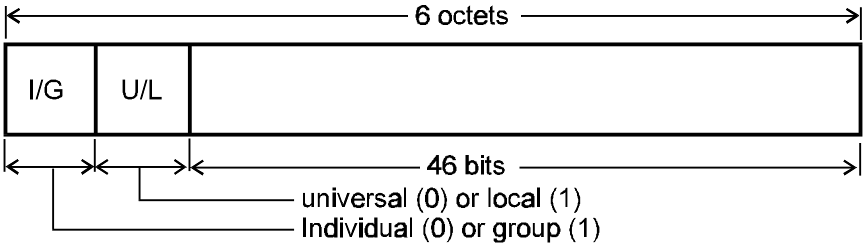

Figure 11—Destination address field

3.2.4.1.1 Individual address

When the first bit of the frame's DA field is equal to a zero (0), the address is an individual address. An individual address identifies a particular station and must be distinct from all other individual station addresses on the same LAN (in the case of local administration) or from the individual addresses of other LAN stations on a global basis (in the case of universal administration).

3.2.4.1.2 Group address

When the first bit of the frame's DA field is equal to a one (1), the address is a group address. A group address is used to address a frame to multiple destination stations. Group addresses may be associated with zero or more stations on a given LAN. In general, a group address is an address associated by convention with a group of logically related stations. It has been observed that several protocols in general usage today require support for one or more general group addresses, and that failure to provide station support for multiple group address recognition can be a serious impediment to use of that station in multi-protocol environments. It is recommended that stations support recognition of multiple general group addresses.

3.2.4.1.3 Broadcast address

The group address X'FF FF FF FF FF FF' is defined as a broadcast address and denotes the set of all stations on a given LAN. X'C0 00 FF FF FF FF' is also a broadcast address for MAC frames. Note that this requirement generally places the address X'C0 00 FF FF FF FF' in the set of addresses recognized by the token ring station and therefore the use of X'C0 00 FF FF FF FF' for non-MAC frames should be avoided.

3.2.4.1.4 Null address

The individual address X'00 00 00 00 00 00' is defined as the null address. It indicates the frame is not addressed to any station. No station shall be assigned the null address and therefore frames addressed to the null address are not expected to be copied by the station.

3.2.4.1.5 Functional addresses (FAs)

FAs are contained within a subset of locally administered group addresses. They are bit-significant addresses used to identify well-known functional entities. There are 31 functional addresses that are indicated by the least significant 31 bits of the destination address when the most significant 17 bits indicate functional addressing. The functional address indicator (FAI) identifies an FA (versus a general group address) when the first 2 octets of the DA are X 'C000' and the first bit of the third octet is B'0' (i.e., B'1100 0000 0000 0000 0'). The remaining 31 bits of the address field each represent a functional address.

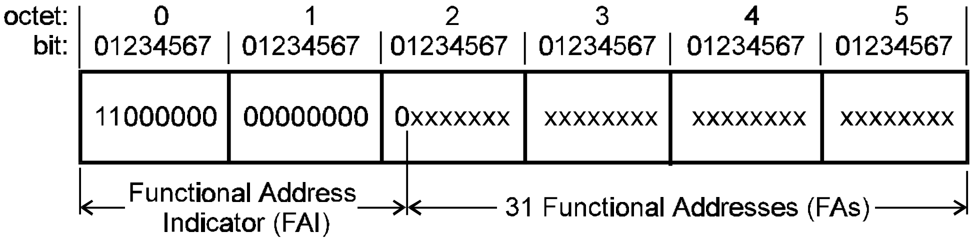

Figure 12—Functional address format

The following functional addresses are defined for use in MAC frames:

Table 1—MAC functional address

Functional address (FA)	Function name	6 octet address
B'xxx xxxx xxxx xxxx xxxx xxxx xxxx xxx1'	Active monitor	X'C0 00 00 00 00 01'
B'xxx xxxx xxxx xxxx xxxx xxxx xxxx xx1x'	Ring parameter server (RPS)	X'C0 00 00 00 00 02'
B'xxx xxxx xxxx xxxx xxxx xxxx xxxx 1xxx'	Ring error monitor (REM)	X'C0 00 00 00 00 08'
B'xxx xxxx xxxx xxxx xxxx xxxx xxx1 xxxx'	Configuration report server (CRS)	X'C0 00 00 00 00 10'

The 6 octet address in the table indicates the address for that single function. Multiple functions are addressed by combining the functional addresses. An example is X'C0 00 00 00 00 12' which is addressed to both the RPS and the CRS functions.

When a station "sets" a functional address, it will recognize all frames sent to that function by setting the A bits and attempting to copy all frames whose DA contains the FAI and has the corresponding FA bit set to B'1'. The active monitor functional address shall only be set as specified in 4.3.4. All other functional addresses are set by management.

Use of functional addresses, octet 5 bit 1 through octet 5 bit 7, that are not defined above are reserved. Functional addresses, octet 2 bit 1 through octet 5 bit 0, are available for general use.

Refer to ISO/IEC TR 11802-2: 1995 for locally administered MAC group (functional) addresses used by this part of ISO/IEC 8802. Also note that functional address X'C0 00 00 00 00 80' is widely used by higher layer protocols and that functional address X'C0 00 00 00 01 00' is widely used for bridging. Because a functional address may be used simultaneously by multiple protocols, a user cannot assume that it designates a single function. Other frame fields should be used to identify the destination function.

The MAC protocol depends on receiving a proper INITIALIZE STATION MAC frame when the RPS functional address is recognized by any station on the ring. The join ring process may not complete properly (e.g., fail joining the ring) if the RPS function does not function as expected (see 4.3.4, JS=RQI). Therefore, if a station does not perform the RPS function, it shall not recognize the RPS functional address and shall not set A or C bits in a frame addressed exclusively to the RPS functional address.

3.2.4.1.6 Address administration

There are two methods of administering station addresses: locally or through a universal authority. The second bit transmitted of the DA indicates whether the address has been assigned by a universal or local administrator:

0 = universally administered 1 = locally administered

3.2.4.1.7 Universal administration

With this method, all individual addresses are distinct from the individual addresses of all other LAN stations on a global basis. The procedure for administration of these addresses is not specified in this standard. Information concerning the registration authority and its procedures may be obtained on request from the Secretary General, ISO Central Secretariat, Case Postale 56, CH-1211 Genève, Switzerland. For information on global address administration contact the Registration Authority for ISO 8802-5, c/o The Institute of Electrical and Electronics Engineers, Inc., 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331, USA.

3.2.4.1.8 Local administration

Individual station addresses are administered by a local (to the LAN) authority. Network administrators should be careful when assigning locally administered addresses due to the structure of hierarchical addressing (defined in previous versions of this standard). The least significant 14 bits of the first two octets of the address have a special meaning (ring ID). Two values of ring ID are assigned for use with all stations (null ring—all zeros, and any ring—all ones). Because of this, the following sets of addresses should be used with care: X'40 00 xx xx xx xx', X'7F FF xx xx xx xx', X'C0 00 xx xx xx xx', X'FF FF xx xx xx xx' (where x is any hexadecimal value). When using locally administered addresses in these ranges, it is possible for the duplicate address test (DAT) to fail when another station is using an address that only differs by the first two octets. In addition, locally administered group addresses with the last 4 octets X'FF FF FF FF' may also be recognized by hierarchical addressing as a broadcast address. Note that the hierarchical addressing “ring ID” is not associated with the “local ring number” subvector.

Hierarchical addressing for locally administrated addresses is not required nor specified by this standard. Future implementations are discouraged from implementing hierarchical addressing.

3.2.4.2 Source address (SA) field

The frame’s SA identifies the station originating the frame. In contrast to the DA field, no individual/group (I/G) bit is encoded in the SA since the SA is constrained to be an individual address. The implied value for the I/G is always 0. In its place is the routing information indicator (RII) bit used to indicate the presence or absence of a routing information (RI) field in the frame. If the RII bit is 0, then no

RI field is present, and if the RII bit is 1, then a RI field is present. The universal/local (U/L) bit still indicates whether the address is universally or locally administered, as in the DA.

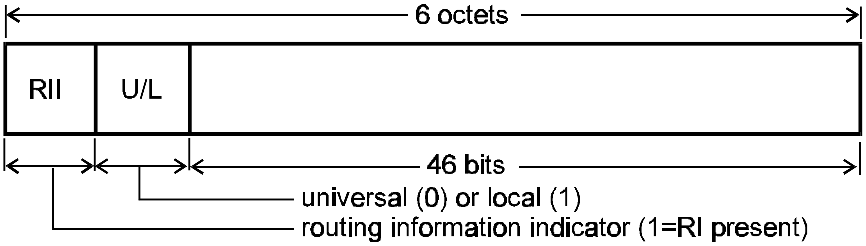

Figure 13—Source address field

3.2.5 Routing information (RI) field

When a frame's routing information indicator bit in the source address field is equal to 1 (RII=1), the RI field shall be included in the frame. The following provides sufficient information for the MAC entity to determine the size of the RI field and parse the frame properly. The detailed structure and contents for the RI field is described in ISO/IEC 10038 : 1993 and ISO/IEC 8802-2 : 1994.

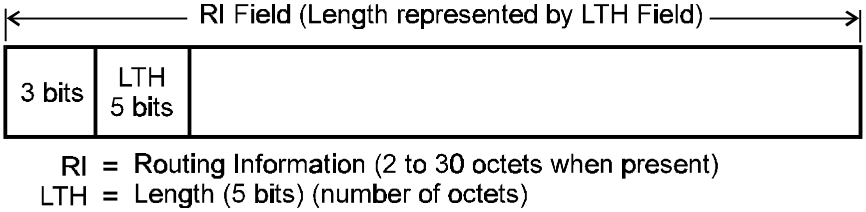

Figure 14—Routing information field

3.2.5.1 Length bits (LTH)

These five bits indicate the length (in octets) of the RI field. Length field values shall be even values between 2 and 30, inclusive.

3.2.6 Information (INFO) field

The information field contains zero, one, or more octets that are intended for MAC, MGT, or LLC. The maximum length of the information field is limited by the maximum frame length as specified in 4.2.4.1, Counter, transmitted octets. The format of the information field is indicated in the frame-type bits of the FC field. The frame types defined are MAC frame and LLC frame.

3.2.6.1 Order of bit transmission

Each octet of the information field shall be transmitted most significant bit first.

3.2.6.2 MAC frame format

The format of the information field for MAC frames is specified in 3.3.2. All stations shall be capable of receiving MAC frames whose total length is the size of the largest frame specified in table 5. However, it is recommended that all stations be capable of receiving MAC frames whose total length is up to and including 96 octets.

3.2.6.3 LLC frame format

The format of the information field for LLC frames is not specified in this standard. However, in order to promote interworking among stations, all stations shall be capable of receiving LLC frames whose information field is up to and including 133 octets in length.

3.2.7 Frame-check sequence (FCS) field

The FCS is a 32-bit sequence based on the following standard generator polynomial of degree 32.

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

A valid FCS is the ones complement of the sum (modulo 2) of the following:

a) The remainder of $X^k (X^{31} + X^{30} + X^{29} + \ldots + X^2 + X + 1)$ divided (modulo 2) by $G(X)$, where k is the number of bits in the FC, DA, SA, RI, and INFO fields, and

b) The remainder after multiplication by X^{32} and then division (modulo 2) by $G(X)$ of the content (treated as a polynomial) of the FC, DA, SA, RI, and INFO fields.

The FCS shall be transmitted commencing with the coefficient of the highest term. As a typical implementation, at the transmitter, the initial remainder of the division is preset to all 1’s and is then modified by division of the FC, DA, SA, RI, and INFO fields by the generator polynomial, $G(X)$. The ones complement of this remainder is transmitted, most significant bit first, as the FCS. At the receiver, the initial remainder is preset to all 1’s and the serial incoming bits of FC, DA, SA, INFO, RI, and FCS, when divided by $G(X)$, results, in the absence of transmission errors, in a unique nonzero remainder value. The unique remainder value is the polynomial:

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

3.2.8 Ending delimiter (ED) field

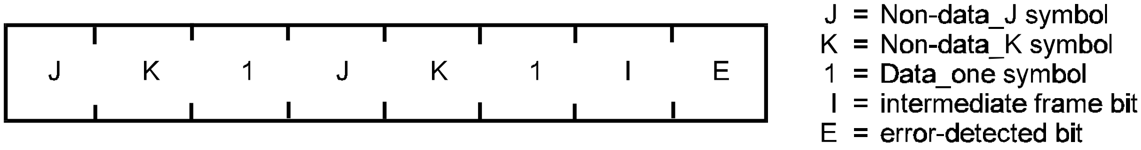

Figure 15—Ending delimiter field

The transmitting station shall transmit the delimiter as shown. Receiving stations shall consider the ED valid if the first six symbols (J K 1 J K 1) are received correctly.

3.2.8.1 Intermediate frame bit (I bit)

The I bit is optionally transmitted as 1 in intermediate (or first) frames of a multiple-frame transmission. The I bit in the last or only frame of the transmission shall be transmitted as 0.

3.2.8.2 Error-detected bit (E bit)

The E bit shall be transmitted as 0 by the station when it originates the token, abort sequence, or frame. All stations on the ring check tokens and all frame formats for errors (for example, FCS error in a MAC or LLC frame or non-data symbols within a frame or token). The E bit of a frame being repeated is set to 1 when a frame with an error (4.3.2) is detected and, optionally, the E bit of a token being repeated is set to 1 when a token with an error (4.3.1) is detected; otherwise, the E bit is repeated as received.

3.2.9 Frame status (FS) field

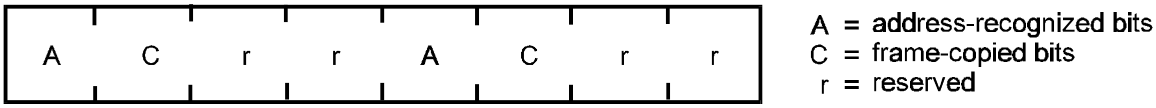

Figure 16—Frame status field

The reserved (r) bits are reserved for future standardization. They shall be transmitted as 0’s and repeated as received.

3.2.9.1 Address-recognized (A) bits and frame-copied (C) bits

The A and C bits shall be transmitted as 0 by the station when it originates the frame. If another station recognizes the DA as its own address or relevant group address, or if indicated by the bridge interface, it sets both of the A bits to 1. If it copies the frame, it also sets both of the C bits to 1. The A and C bits are set without regard to the value of the received E bit and only if the frame is valid as defined in 4.3.2.

3.2.10 Interframe gap (IFG) field

The interframe gap is composed of fill (3.1.4) and shall be a minimum of 1 octet in length (a length of 2 octets is recommended) for 4 Mbit/s transmission and shall be a minimum of 5 octets in length for 16 Mbit/s transmission. Note that the size and content of the IFG may be altered (see 5.8.3.3). Receiving stations can not be assured of receiving the full IFG (5.8.3.3). Note that additional fill may follow the interframe gap.

3.3 Medium Access Control (MAC) frames

The following subclauses give descriptions of various MAC frames that are used in the management of the token ring. Definitions and requirements for PDU priority (Pm), FC, DA, INFO field content (Vector Class—VC, Vector Identifier—VI, Subvector Identifier—SVI, and Subvector Value—SVV) are also given.

3.3.1 MAC frame control

Frames with the following FC values are handled as follows:

a) If the value of the FC of the frame is X'00' and it is addressed to the station, it will be copied only if there is sufficient free buffer available for copying.

b) If the frame has an FC field value X'01' through X'3F', it is addressed to the station, and it has a Destination Class B'0000' (station), then every effort should be made to copy the frame. The MAC protocol (4.3.4) has been designed to handle the case when a MAC frame is not copied. FC values of X'01' through X'06' shall only be used for the corresponding MAC frames as defined in table 4.

3.3.2 MAC frame information field format

Figure 17 shows the format of the information field for MAC frames.

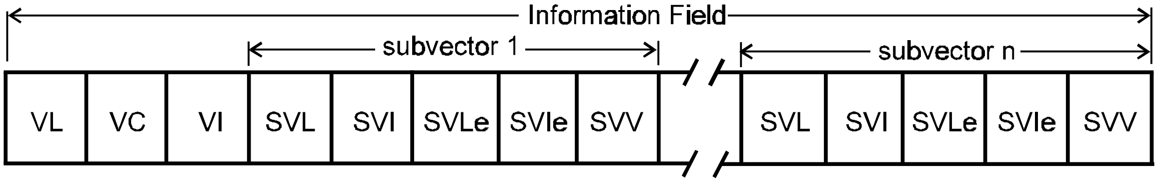

Figure 17—MAC frame information field structure

Vector. The fundamental unit of MAC and MGT information. A vector contains its length, an identifier of its function, and zero or more subvectors. Only one vector is permitted per MAC frame.

VL (vector length). A 16-bit binary number that gives the length, in octets, of the vector. The length includes the length of the VL field and can have values such that X'0004' ≤ VL ≤ X'FFFF', subject to the other constraints in clause 4 on the length of the INFO field.

VC (vector class). A one-octet code point that identifies the function class of both the source and destination entities. The function classes are defined as follows:

Table 2—Vector class values

Function class	Class value
Ring station (RS)	X'0'
Configuration report server (CRS)	X'4'
Ring parameter server (RPS)	X'5'
Ring error monitor (REM)	X'6'
Network management (NM)	X'8'

The organization of the vector class field is shown in figure 18.

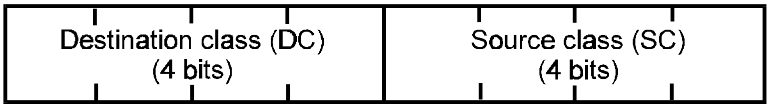

Figure 18—Vector class field

The VC field is divided into two 4-bit fields specifying the destination class (DC) and the source class (SC).

The DC field provides a means to route the frame to the appropriate management function within the station. Table 5 defines the MAC frames sent to destination class 0 that are processed by the ring station as specified in 4.3.4. All other MAC frames sent to destination class 0 may be ignored.

The SC field allows the station’s MAC and MGT functions to reject vector identifiers from invalid sources.

VI (vector identifier). A one- or three-octet code point that uniquely identifies the vector. This field is one octet in length if the first octet is not X'FF' and three octets in length if the first octet is X'FF'. A value of X'FF' means the vector code is contained in the next two octets. A vector code of less than X'FF' shall be transmitted using one octet. A vector code of greater than or equal to X'FF' shall be transmitted using three octets. The definition of all undefined vector codes used with a source class or destination class of zero (ring station) is reserved for future standardization.

SV (subvector). Vectors require all data or modifiers to be contained within subvectors. A separate subvector is required to contain each piece of data or modifier that is being transported. A subvector is not position-dependent within a vector, but rather, each subvector is identified by its subvector identifier.

SVL (subvector length). A one-octet field that uniquely identifies the length of the subvector depending on its value. If the SVL value is less than X'FF' then SVL contains the length, in octets, of the subvector including the length of SVL, SVI, SVIe (if present), and SVV fields. If the value of SVL is X'FF', then the subvector length is specified in the SVLe field.

SVI (subvector identifier). A one-octet field that identifies the subvector depending on its value. A value of X'FF' indicates the subvector code is contained in the SVIe field.

Definitions of the SVIe subvector codes are beyond the scope of this standard and are reserved for future standardization. The subvectors with code points from X'00' through X'7F' are defined to have the same definition for all vectors so that certain specific strings that are common to many vectors can be formatted and labeled in a standard manner. This standardization is intended to facilitate sharing of data between MAC and MGT applications and make the data as application-independent as possible. The subvectors with code points from X'80' through X'FE' are for specific definition within a particular vector-by-vector identifier. For example, the subvector X'90' can have an entirely different definition in every different vector. The subvector X'40' has only one definition across all vectors and applications. Subvectors themselves may contain other subvectors and other types of vectors and optional fields that are unique only to the particular subvector to which they belong.

SVLe (extended subvector length). A zero- or two-octet field depending on the value of SVL. If the SVL value is less than X'FF' then the length of SVLe is zero, When the SVL value is X'FF', then the length of the SVLe field is 2 octets and its value identifies the length, in octets, of the subvector including the length of the SVL, SVI, SVLe, SVIe (if present), and SVV fields.

SVIe (extended subvector identifier). A zero- or two-octet field depending on the value of SVI whose value identifies the subvector. If the value SVI is less than X'FF' then the length of SVIe is zero. When the SVI value is X'FF', the SVIe field is 2 octets in length and contains the code point identifying the subvector. The definition of the extended subvector identifier is provided for future expansion and definition of the SVIe subvector codes is reserved for future standardization.

SVV (subvector value). This field contains the subvector information as described in 3.3.4.

3.3.3 Vector descriptions

Following is a list of vectors used by the ring station in numerical order of the vector identifier.

X'00'—Response (RSP). The RSP MAC frame is sent by a station as specified by the MAC protocol to acknowledge receipt of, or to report an error in, a received MAC frame. The RSP MAC frame is never sent to Destination Class 0 (ring station). The RSP MAC frame is not required to be sent if the received frame initiating the response was sent to a broadcast address and the address-recognized and frame-copied bits in the FS field were all ones.

X'02'—Beacon (BN). The BN MAC frame is transmitted by a station which detects a continuing interruption to the data flow on the ring.

X'03'—Claim Token (CT). The CT MAC frame is used by stations when restoring the active monitor function to the ring.

X'04'—Ring Purge (RP). The RP MAC frame is transmitted by the active monitor to verify the ring data path before restoring the token.

X'05'—Active Monitor Present (AMP). The AMP MAC frame is transmitted periodically by the active monitor to indicate its presence on the ring and to initiate the Neighbor Notification Process.

X'06'—Standby Monitor Present (SMP). The SMP MAC frame is transmitted by standby monitor stations as part of the Neighbor Notification Process.

X'07'—Duplicate Address Test (DAT). The DAT MAC frame is used during the Join Ring Process to verify that the station's individual address is unique.

X'08'—Lobe Media Test (TEST). The TEST MAC frame is optionally used by a station during the Join Ring Process to test its lobe connection prior to inserting into the ring or at any time the station needs to test the transmission path such as during the Beacon Test state. When the TEST MAC frame is used, it is addressed to the null address which is specifically addressed to no station. The station determines the quality of the transmission path by inspecting the received frame for errors (see 4.3.2).

X'0B'—Remove Ring Station (REMOVE). The REMOVE MAC frame is sent to any of the station's non-broadcast addresses to force its removal from the ring. The REMOVE MAC frame is sent by a configuration report server (CRS) or optionally by network management (NM).

X'0C'—Change Parameters (CHG_PARM). The CHG_PARM MAC frame is sent by the CRS to a ring station to set its ring operating parameters. The ring station normally replies with an RSP MAC frame if the SET_APPR_PARMS action is taken.

X'0D'—Initialize Station (INIT). The INIT MAC frame is sent by the RPS to a ring station to initialize its ring operating parameters. The ring station normally replies with an RSP MAC frame if the SET_APPR_PARMS action is taken.

X'0E'—Request Station Addresses (RQ_ADDR). The RQ_ADDR MAC frame is sent by any non-zero function class (such as CRS) to a station(s) to request the addresses recognized by the station.

X'0F'—Request Station State (RQ_STATE). The RQ_STATE MAC frame is sent by any non-zero function class (such as CRS) to a station(s) to request information on the state of the station.

X'10'—Request Station Attachments (RQ_ATTCH). The RQ_ATTCH MAC frame is sent by any non-zero function class (such as CRS) to a station(s) to request information on the functions active in the station.

X'20'—Request Initialization (RQ_INIT). The RQ_INIT MAC frame is sent by a station to the RPS as part of the join ring process. It informs the RPS that a station has been inserted and will accept modified parameters from either the RPS or CRS.

X'22'—Report Station Addresses (RPRT_ADDR). The RPRT_ADDR MAC frame is sent by a station in response to a Request Station Addresses MAC frame.

X'23'—Report Station State (RPRT_STATE). The RPRT_STATE MAC frame is sent by a station in response to a Request Station State MAC frame.

X'24'—Report Station Attachments (RPRT_ATTCH). The RPRT_ATTCH MAC frame is sent by a station in response to a Request Station Attachments MAC frame.

X'25'—Report New Active Monitor (NEW_MON). The NEW_MON MAC frame is sent by a station to the CRS to report that it has become the active monitor.

X'26'—Report SUA Change (SUA_CHG). The SUA_CHG MAC frame is sent by a station to the CRS when a change is made in the station's SUA as the result of the neighbor notification process.

X'27'—Report Neighbor Notification Incomplete (NN_INCMP). The NN_INCMP MAC frame is sent by the active monitor to the REM to report an error in the neighbor notification process.

X'28'—Report Active Monitor error (ACT_ERR). The ACT_ERR MAC frame is sent to the REM by an active monitor if it detects more than one station is active monitor on the ring or if another station indicates a problem with the active monitor function by initiating the Claim Token process.

X'29'—Report Error (RPRT_ERR). The RPRT_ERR MAC frame is sent to the REM by a station to report errors detected since the last error report.

The following VIs are reserved:

X'01'- Re-IMPL
X'09'-Transmit forward
X'0A'-Resolve
X'2A'-Report transmit forward

3.3.4 Subvector descriptions

Following is a list of subvectors, in numerical order of their respective subvector identifier values. The lengths specified are only for the SVV field and do not include the length of SVL and SVI fields.

X'01'—Beacon Type. This subvector has a value field 2 octets long and is used to indicate the type of fault detected. It has one of the following values:

X'0001'—Issued by a Dual Ring station during reconfiguration. Refer to IEEE Std 802.5c-1991 [B4].[6]
X'0002'—Signal loss.
X'0003'—Timer TCT expired during claiming token; no claim token frames received.
X'0004'—Timer TCT expired during claiming token; claim token frames received.

---

[6] The numbers in brackets preceded by the letter B correspond to those of the bibliography in annex J.

X'02'—Upstream Neighbor's Address (UNA). This subvector value field is 6 octets long and contains the address of the upstream neighbor of the sending station. When transmitting, this is the value of SUA. When received, the value of this subvector is referred to as Reported Upstream Address (RUA).

X'03'—Local Ring Number. This subvector has a value field 2 octets long. It indicates the local ring number of the sending station. This value is assigned to the station by the Initialize Station MAC frame or the Change Parameters MAC frame and is used in the subvector field of appropriate report frames. This value does not affect the MAC address assigned to the station.

X'04'—Assign Physical Drop Number. This subvector has a value field 4 octets long. It specifies the physical location of the target station. The value of this subvector is not defined by this standard. This value is assigned to the station by the Initialize Station MAC frame or the Change Parameters MAC frame and is used in the subvector field X'0B' of appropriate report frames.

X'05'—Error Report Timer Value. This subvector has a value field 2 octets long. It states the value of TER in increments of 10 ms.

X'06'—Authorized Function Classes. This subvector has a value field 2 octets long. It indicates the functional classes that are allowed to be active in the station. Valid range is B'0000 0000 0000 0000' to B'1111 1111 1111 1111' where each bit 0 to 15 corresponds to function class X'0' to X'F'. Defined function classes are the following:

Table 3—Function class values

Subvector value	Function class	Class value
B'xxxx 1xxx xxxx xxxx'	Configuration report server (CRS)	X'4'
B'xxxx x1xx xxxx xxxx'	Ring parameter server (RPS)	X'5'
B'xxxx xx1x xxxx xxxx'	Ring error monitor (REM)	X'6'
B'xxxx xxxx 1xxx xxxx'	Network Management (NM)	X'8'

Other function classes are reserved for future standardization.

X'07'—Authorized Access Priority. This subvector has a value field 2 octets long. The value is set by the system administrator and it may establish the maximum frame access priority (Pm) at which a station queues non-MAC frames.

X'09'—Correlator. This subvector has a value field 2 octets long and is used to associate responses with requests. The value of the correlator SV for a frame sent in response to a received frame shall be the value of the received correlator SV. If the received frame does not contain a correlator, the MAC may send any value or omit the correlator SV in the transmitted frame.

X'0A'—SA of Last AMP or SMP Frame. This subvector has a value field 6 octets long and is used in the Report Neighbor Notification Incomplete MAC frame. Its value is the same as the value of LMP which is specified in table 7. A value of null indicates that the AMP frame was not received. A value equal to the station's address indicates either that the AMP frame was not received or that no SMP frames were received. Any other value indicates the SA of the last SMP frame received.

X'0B'—Physical Drop Number. This subvector has a value field 4 octets long. It reports the physical location of the sending station (see X'04' Assign Physical Drop Number).

X'20'—Response Code. This subvector has a value field 4 or 6 octets long and is used in the Response MAC frame. It consists of a 2-octet response code followed by the 1-octet vector class (VC) and the 1- or 3-octet vector identifier (VI) from the received MAC frame that caused the station to send the Response MAC frame. The response code values are as follows:

X'0001'—Positive acknowledgment. The MAC frame was accepted by the station.
X'8001'—MAC frame information field is incomplete. The MAC frame was too short to contain the vector length, vector class and vector ID.
X'8002'—Vector length error. Vector length does not agree with the length of the frame or a subvector was found that did not fit within the vector.
X'8003'—Unrecognized vector ID. The vector ID is not recognized by the station.
X'8004'—Inappropriate source class. The source class is not valid for the VI.
X'8005'—Subvector length error. The length of a recognized subvector conflicts with its expected length or is less than 2.
X'8006'—Reserved.
X'8007'—Missing subvector. A subvector required to process the MAC frame is not in the MAC frame.
X'8008'—Subvector unknown. A subvector received in the MAC frame is not known by the adapter.
X'8009'—MAC frame too long. The received frame was rejected because it exceeded maximum length.
X'800A'—Function requested was disabled. The received MAC frame was rejected because the function requested was disabled.

X'21'—Individual Address Count. This subvector has a value field 2 octets long. The default value of X'00 00' means more than one individual address is not supported. A non-zero value specifies the number of individual addresses in use by this station.

X'22'—Product Instance ID. This subvector’s value is used by a station manufacturer to identify a station’s characteristics, such as serial number, machine type, model number, plant of manufacture, etc. The length of this subvector is not defined by this standard. It is recommended that this subvector be the “ResourceTypeID” managed object as specified by ISO/IEC 10742 : 1994.

X'23'—Ring Station Version Number. This subvector is used in the Request Initialization and Report Station State MAC frames. The length and value of this subvector is not defined by this standard.

X'26'—Wrap Data. The length and function of this subvector are product implementation choices. The subvector is used in the Lobe Media Test MAC frame.

X'28'—Station Identifier. This subvector has a value field 6 octets long and is used in the Report Station State MAC frame. It should uniquely identify the station. It is recommended that this value be a universally administered individual address.

X'29'—Ring Station Status. This subvector is used in the Report Station State MAC frame. The length and contents of this subvector are not specified by this standard. However, an application receiving this subvector may be able to determine its format by examination of the Product Instance ID subvector (X'22').

X'2B'—Group Addresses. This subvector has a value field of 4 or 6 octets in length. It contains a group address of the reporting station. When 4 octet values are used, the field will contain the low order 4 octets of the address and no assumptions can be made about the first 2 octets of the address. A value of all zeros is used to denote that a station does not support a group address or that a group address is not assigned. If more than one group address is recognized, then any of those addresses may be reported.

X'2C'—Functional Addresses. This subvector has a value field 4 octets long and specifies the functional addresses that are active in the reporting station (see 3.2.4.1).

X'2D'—Isolating Error Counts. This subvector has a value field 6 octets long and is used in the Report Error MAC frame. It indicates the number of each type of error detected since the last error report. Refer to 3.6 for a description of each of the counters. Isolating errors are errors that are known to occur between the reporting station and its upstream station. If a counter has not been incremented or is marked reserved, its value shall be reported as X'00'.

Octet 0—line error (CLE)
Octet 1—internal error (CIE)
Octet 2—burst error (CBE)
Octet 3—AC error (CACE)
Octet 4—abort sequence transmitted (CABE)
Octet 5—reserved (X'00')

X'2E'—Non-isolating Error Counts. This subvector has a value field 6 octets long and is used in the Report Error MAC frame. It indicates the number of each type of error detected since the last error report. Refer to 3.6 for a description of each of the counters. The origin of Non-isolating errors are unknown. If a counter has not been incremented or is marked reserved, its value shall be reported as X'00'.

Octet 0—lost frame error (CLFE)
Octet 1—receive congestion (CRCE)
Octet 2—frame-copied error (CFCE)
Octet 3—frequency error (CFE)
Octet 4—token error (CTE)
Octet 5—reserved (X'00')

X'30'—Error Code. This subvector has a value field 2 octets long and is used in the Report Active Monitor Error MAC frame. It has one of the following values:

X'0001' — Active monitor error, used when the active monitor receives a Claim Token MAC frame.
X'0002' — Duplicate active monitor, used when the active monitor receives a RP or an AMP MAC frame that it did not transmit, indicating the presence of another active monitor.
X'0003' — Duplicate address, used when a station in claim token receives a Claim Token MAC frame in which the SA equals the station’s individual address (SA equal to MA) but the RUA is different from the station’s SUA. This indicates that another station on the ring has the same individual address.

The following SVIs are reserved:

X'27'—Frame Forward
X'2A'—Transmit Forward Status Code
X'2F'—Function Request ID

3.3.5 MAC frame tables

Tables 4 and 5 are tabulations of MAC frame requirements for transmitting and receiving stations, respectively.

3.3.5.1 MAC frames transmitted

A station shall support the frames shown in table 4. The Beacon, Claim Token, Ring Purge, Active Monitor Present, and Standby Monitor Present MAC frames shall be transmitted as shown. Other frames may be transmitted with additional subvectors.

MAC frames shall be queued for transmission with the Pm values specified in table 4. It is recommended that all MAC frames transmitted on a token be queued for transmission with a Pm of 7.

Table 4—MAC frame transmit definitions

Vector (VI,Name)	FC Pm	DA	VC	Designator**7 Subvector (SVI,Name)

X'00' Response	X'00' **2	SA of rcvd frame	X'*0'		X'09' Correlator **5 X'20' Response Code
X'02' **4 Beacon	X'02' **3	broadcast	X'00'		X'01' Beacon Type X'02' UNA X'0B' Physical Drop Number
X'03' **4 Claim Token	X'03' **3	broadcast	X'00'		X'02' UNA X'0B' Physical Drop Number
X'04' **4 Ring Purge	X'04' **3	broadcast	X'00'		X'02' UNA X'0B' Physical Drop Number
X'05' **4 Active Monitor Present	X'05' **1	broadcast	X'00'		X'02' UNA X'0B' Physical Drop Number
X'06' **4 Standby Monitor Present	X'06' **1	broadcast	X'00'		X'02' UNA X'0B' Physical Drop Number
X'07' **4 Duplicate Address Test	X'01' **1	DA=MA	X'00'		none
X'08' Lobe Media Test	X'00' **1	Null	X'00'		X'26' Wrap Data
X'20' **4 Request Initialization	X'00' or X'01' **1 **6	FA(RPS)	X'50'	 op_txop_tx op_tx	X'02' UNA X'21' Individual Address Count X'22' Product Instance ID X'23' Ring Station Version Number
X'22' Report Station Addresses	X'00' **2	SA of request	X'*0'	 op_tx op_tx	X'09' Correlator **5 X'02' UNA X'2B' Group Address X'2C' Functional Address(s) X'21' Individual Address Count X'0B' Physical Drop Number
X'23' Report Station State	X'00' **2	SA of request	X'*0'	 op_tx op_tx	X'09' Correlator **5 X'28' Station Identifier X'29' Ring Station Status X'23' Ring Station Version Number
X'24' Report Station Attachments	X'00' **2	SA of request	X'*0'	 op_tx op_tx	X'09' Correlator **5 X'06' Authorized Function Classes X'07' Authorized Access Priority X'2C' Functional Address(s) X'21' Individual Address Count X'22' Product Instance ID
X'25' **4 Report New Active Monitor	X'00' **1	FA(CRS)	X'40'	 op_tx op_tx	X'02' UNA X'0B' Physical Drop Number X'22' Product Instance ID
X'26' **4 Report SUA Change	X'00' **1	FA(CRS)	X'40'	 op_tx	X'02' UNA X'0B' Physical Drop Number
X'27' **4 Report Neighbor Notification Incomplete	X'00' **1	FA(REM)	X'60'		X'0A' SA of Last AMP or SMP Frame
X'28' **4 Report Active Monitor Error	X'00' **1	FA(REM)	X'60'	 op_tx	X'30' Error Code X'02' UNA X'0B' Physical Drop Number
X'29' **4 Report Error	X'00' **1	FA(REM)	X'60'	 op_tx	X'2D' Isolating Error Count X'2E' Non-isolating Error Count X'02' UNA X'0B' Physical Drop Number

Legend: **VC=X'*0'** Indicates the destination class of the requesting entity.

****1**	Indicates Pm=7 or 3.
****2**	Indicates Pm=7 or 0.
****4**	Shall not be transmitted with an RI field.
****5**	See 3.3.4 for exception to X'09' Correlator SV.
****6**	During the join process, a station may transmit multiple Request Initialization MAC frames (Vector X'20') as part of the assured delivery process. The first Request Initialization frame shall be transmitted with an FC value of X'00'. If the Ring Parameter Server sets the A bits in the frame, but subsequently does not respond to the station, the station may transmit subsequent Request Initialization MAC frames with an FC value of X'01'.
****7**	Subvectors with no preceding designator shall be transmitted.
op_tx	Indicates SV is optional and may be omitted.

3.3.5.2 MAC frame reception

A station shall support reception of the frames as defined in table 5.

A station may optionally support reception of frames not specified in table 5.

The term “Ignore MAC Frame” as used in this clause means the station has received a MAC frame, but the station operation tables ignore the MAC frame as it is not successfully verified as a correctly formatted MAC frame destined for an address recognized by the station.

Table 5—MAC frame receive definitions

Vector (VI,Name)	FC	VC	Designator**3	Subvector (SVI,Name)
X'02' Beacon	X'02'	X'00'	req req op_req	X'01' Beacon Type X'02' UNA X'0B' Physical Drop Number
X'03' Claim Token	X'03'	X'00'	req op_req	X'02' UNA X'0B' Physical Drop Number
X'04' Ring Purge	X'04'	X'00'	req op_req	X'02' UNA X'0B' Physical Drop Number
X'05' Active Monitor Present	X'05'	X'00'	req op_req	X'02' UNA X'0B' Physical Drop Number
X'06' Standby Monitor Present	X'06'	X'00'	req op_req	X'02' UNA X'0B' Physical Drop Number
X'07' Duplicate Address Test	X'01'	X'00'		none
X'0B' Remove Ring Station	X'01'	X'04' ****2**		none
X'0C' Change Parameters	X'00'	X'04' ****2**		X'09' Correlator X'03' Local Ring Number X'04' Assign Physical Drop Number X'05' Error Timer Value X'06' Authorized Function Classes X'07' Authorized Access Priority
X'0D' Initialize Station	X'00' or X'01' ****1**	X'05'		X'09' Correlator X'03' Local Ring Number X'04' Assign Physical Drop Number X'05' Error Timer Value
X'0E' Request Station Addresses	X'00'	X'0*'		X'09' Correlator
X'0F' Request Station State	X'00'	X'0*'		X'09' Correlator
X'10' Request Station Attachments	X'00'	X'0*'		X'09' Correlator

Legend:		
	VC=X'0*'	Indicates any nonzero source class.
	req	The SV is required to be present for the frame to be considered valid.
	op_req	The station may optionally require this subvector to be present.
	****1**	A station may optionally receive the Initialize Station MAC frame (Vector X'0D') with an FC value of X'01'. Management servers transmitting this frame are recommended to use FC values of both X'00' and X'01'.
	****2**	A station may optionally receive these frames with a VC of X'08'.
	****3**	Subvectors with no preceding designator may or may not be present within the received vector, and are not required for verification.

When a valid (4.3.2) MAC frame addressed to the station is copied, it is processed as follows. When the destination class (DC) is other than X'0', the frame is indicated to the management interface (MGT_UNITDATA.indication) and no other action is taken by the MAC protocol. When the DC is X'0', then the station performs a verification process and, based on the result, takes the actions specified in tables 7.1 to 7.6. Figure 18.1 shows the processing of a received MAC frame.

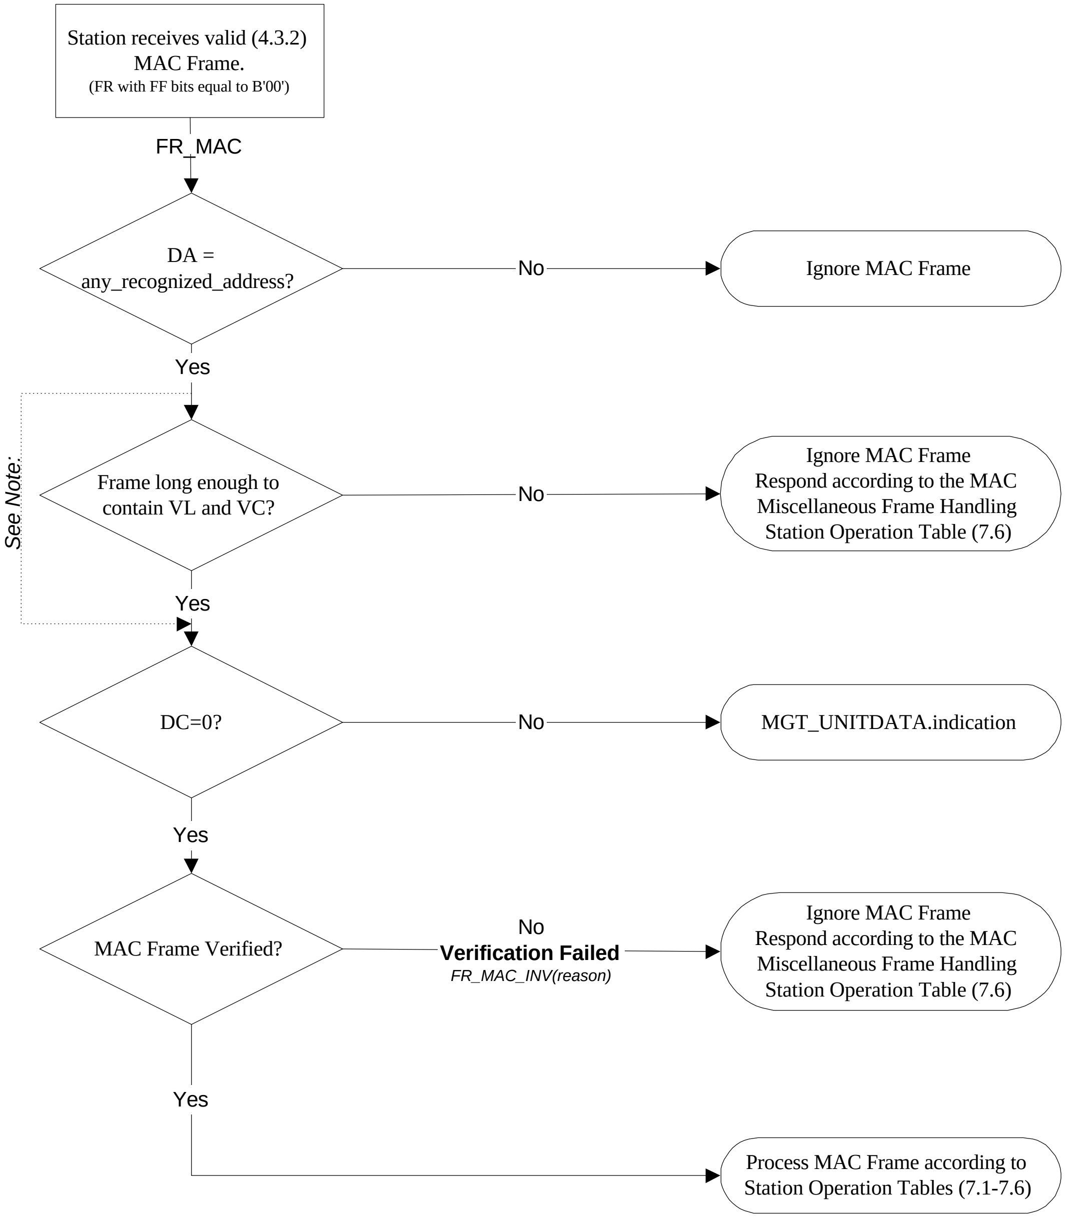

NOTE—It is not mandatory to check that the frame is long enough to contain VL and VC, but it is recommended.

Figure 18.1—Received MAC frame processing

The verification of the received MAC frame shall fail if any of the following criteria are met:

RI_INVALID: The station may ignore any MAC frame containing an RI field.

VI_UNK: Unrecognized vector ID value. The station shall reject the frame if the VI value is not **known** by the station. The station may reject or ignore the frame if the FC field does not match the value specified for the VI in table 5. Validity checking of the FC field is not required by this standard

VI_LTH_ERR: Vector length error. The station shall reject the frame if the VL does not agree with the length of the frame, or if the VL is not equal to the sum of all the SVLs plus the length of VL,VC, and VI fields.

SHORT_MAC: MAC frame information field is incomplete. The MAC frame is too short to contain the vector length, vector class, and vector ID.

LONG_MAC: MAC frame is too long. The station may reject the frame if the VL indicates a size larger than that needed for all allowed subvectors.

SV_LTH_ERR: Subvector length error. The station shall reject the frame if the subvector length is less than the minimum or greater than the maximum specified in 3.3.4.

SV_MISSING: Missing one or more required subvectors. Based on the VI, the station shall reject a frame that does not contain a subvector listed as “req” in table 5. The station may additionally reject the frame if it does not contain a subvector listed as “op_req” in table 5. The station shall not require subvectors that are not listed in table 5.

SV_UNK: SVI value is unknown. The station may reject the frame if the SVI value is not known by the station. The station shall recognize all SVI values given in table 5.

SC_INVALID: Invalid source class. The station shall reject the frame if the VC does not match the value specified in table 5.

FUNCTION_DISABLED: The requested function is disabled by management.

The order of parsing received frames is not specified by this standard. Annex H has been provided as an example of frame verification parsing.

As specified in 4.3.4, certain MAC frames require a response frame if the frame is rejected.

Setting parameters: When a valid Change Parameters or Initialize Station MAC frame is received, the local parameters are set to the value specified. These parameters are used as follows:

X'03'—Local ring number: This parameter is reserved for use by management and not otherwise used by the station.

X'04'—Assign physical drop number: This parameter is used as the X'0B' Physical Drop Number subvector in appropriate MAC frames transmitted by the station (3.3.5.1).

X'05'—Error timer value: When set, the value of this subvector becomes the time interval for the error timer TER (3.4.2.5).

X'06'—Authorized function classes: When set, this subvector defines which function classes may be active in a station. The default is that all function classes are allowed. If a function class is not enabled, the station shall not transmit frames with the corresponding source class, the station shall not set the functional address assigned to that function, and thus shall not set the A bits in frames addressed to the functional address assigned to that function.

X'07'—Authorized access priority: When set, this parameter may establish the maximum priority used for queuing non-MAC frames. The default is that any priority except 7 may be used. A station is not required to take any action based on this parameter.

3.4 System timing parameters

Except for the token holding time, system timing parameters are specified by the definition of specific timers.

3.4.1 Token holding time

Token holding time is the length of time, after capturing a token, that a station may continue to transmit frames before releasing the token. This time is specified in terms of the equivalent number of octets transmitted in a period of 9.1 ms. Refer to “Counter, Transmitted Octets (CTO)” in 4.2.4.1. The MAC protocol uses the CTO counter to determine whether the next frame can be completely transmitted in the time remaining. The difference between the 9.1 ms used for the basis of CTO and the minimum value of TVX is not dependent on ring latency.

3.4.2 Timers

In general, timers control the maximum period of time that a particular condition may exist. All timers are stopped when the bypass state (J0) of the join state machine is entered (refer to clause 4 for details) and do not start until the first time they are reset. In defining the timers, the terms *reset* and *expired* are used, as explained below.

— **reset:** The term *reset* means that timing is started for an interval equal to the value of the timer.

— **expired:** The term *expired* means that a time interval equal to the timer value has elapsed since the timer was reset. A timer that expires remains stopped until reset.

Each timer is described below and the use of each timer is specified in 4.3.4.

Implementors should consider the granularity of their timers when choosing values to ensure that the values used fall in the ranges specified.

3.4.2.1 Timer, active monitor (TAM)

Each station shall have a timer TAM. This timer is used by the active monitor function to stimulate the enqueuing of an AMP PDU for transmission. The value of TAM shall be between 6.8 s and 7.0 s.

3.4.2.2 Timer, beacon repeat (TBR)

Each station shall have a timer TBR. This timer is reset each time the station receives a beacon MAC frame while in Beacon Repeat to allow the station to detect the absence of the beacon MAC frames. The absence of beacon MAC frames is detected by the timer TBR expiring and causes entry into Transmit Claim Token. The value of TBR shall be between 200 ms and 400 ms. A value of 200 ms is recommended.

3.4.2.3 Timer, beacon transmit (TBT)

Each station shall have a timer TBT. This timer is used to specify the length of time a station remains in Transmit Beacon before entering Beacon Test. The value of TBT shall be between 15.8 s and 26 s. A value of 18 s is recommended.

The recommended value for TBT has been changed to 18 s as the original value of 16 s could occasionally lead to a 4 Mbit/s station and a 16 Mbit/s coexisting on a two station ring.

3.4.2.4 Timer, claim token (TCT)

Each station shall have a timer TCT. This timer is reset either when entering Transmit Claim token or upon reception of the first claim token MAC frame. It is used to detect the failure of the Claim Token process and allow the Monitor FSM to enter Beacon Transmit state or the Join FSM to remove the station from the ring. The value of TCT shall be between 1.0 s and 1.2 s.

3.4.2.5 Timer, error report (TER)

Each station shall have a timer TER. This timer is used to report errors that have been detected by the station. When an error counter is incremented, TER, if not already running, is reset. The station queues a

Report Error PDU and resets all error counters to zero when TER expires. The default value of TER shall be between 2.0 s and 2.2 s. This value can be changed by the Change Parameters MAC frame or the Initialize Station MAC frame.

3.4.2.6 Timer, insert delay (TID)

Each station shall have a timer TID. This timer is used to delay the station’s recognition of its own beacon MAC frames until sufficient time has elapsed after Beacon Test to assure the PHY has put the station in the ring’s data path (no longer wrapped at the TCU). This is done to prevent premature resolution of beaconing (refer to clause 4 for details). The value of TID shall be between 5 s and 20 s.

3.4.2.7 Timer, join ring (TJR)

Each station shall have a timer TJR. This timer is used by the inserting station to detect the failure of the join ring process. The value of TJR shall be between 17.8 s and 18.0 s.

3.4.2.8 Timer, no token (TNT)

Each station shall have a timer TNT. This timer is used to recover from various token related error situations. The operation of timer TNT is described in clause 4. The value of TNT shall be between 2.6 s and 5.2 s. A value of 2.6 s is recommended.

3.4.2.9 Timer, queue PDU (TQP)

Each station shall have a pacing timer TQP. This timer is used to schedule the following events:

a) The enqueuing of a SMP PDU after reception of an AMP or a SMP MAC frame in which the A and C bits were equal to 0 when in the Repeat state (MS=RPT).
b) The transmission of a Claim Token MAC frame when in Transmit Claim Token state (MS=TCT).
c) The transmission of a Beacon MAC frame when in Transmit Beacon state (MS=TBN).

The value of TQP shall be between 10 ms and 20 ms. A value of 20 ms is recommended.

3.4.2.10 Timer, remove hold (TRH)

Each station shall have a timer TRH. This timer is used to delay the remove action for a minimum amount of time after insertion. This time allows the TCU to accomplish the station’s insert request before the station issues the remove request. This timer is specified for interoperability with TCUs built prior to this specification. The value of the TRH shall be between 5 and 7 s.

3.4.2.11 Timer, remove wait (TRW)

Each station shall have a timer TRW. This timer is used to control the time a station must remain in repeat to allow sufficient time for the TCU to remove the inserted station’s lobe from the ring (see 7.2.1.1). The value of TRW shall be a minimum of 200 ms.

3.4.2.12 Timer, request initialize (TRI)

Each station shall have a pacing timer TRI. This timer is used during join ring to detect that a response to the last Request Initialization MAC frame was not received and cause the Request Initialization MAC frame to be enqueued after the join ring FSM has completed neighbor notification. The value of TRI shall be between 2.5 s and 2.8 s.

3.4.2.13 Timer, return to repeat (TRR)

Each station shall have a timer TRR. This timer is used to ensure that the transmitting station returns to the Repeat state. The value of TRR shall be between 4.0 ms and 4.1 ms.

3.4.2.14 Timer, ring purge (TRP)

Each station shall have a timer TRP. This timer is used by the station to detect the failure of ring purge and exit to the Transmit Claim Token state. The value of TRP shall be between 1.0 s and 1.2 s.

3.4.2.15 Timer, signal loss (TSL)

Each station shall have a timer TSL. This timer is used to determine whether or not the PHY "signal_loss" condition is in a steady state. The value of TSL shall be between 200 ms and 250 ms.

3.4.2.16 Timer, standby monitor (TSM)

Each station shall have a timer TSM. This timer is used to assure that there is an active monitor on the ring or to detect token streaming. The value of TSM shall be between 14.8 s and 15.0 s.

3.4.2.17 Timer, valid transmission (TVX)

Each station shall have a timer TVX. This timer is used by the active monitor to detect the absence of frames or tokens. The value of this timer indirectly limits the maximum number of octets a station may transmit after capturing a token (token holding time) before it must release a token. The value of TVX shall be between 10.0 ms and 11.0 ms.

3.4.2.18 Timer, wire fault delay (TWFD)

Each station shall have a timer TWFD. This timer is used to delay the detection of wire fault after Join Ring completes (JS=JC, refer to clause 4) or the station inserts after Beacon Test. The value of TWFD shall be between 5.0 s and 10.0 s. The sum of TWFD + TWF shall be a minimum of 7.0 s.

3.4.2.19 Timer, wire fault (TWF)

Each station shall have a timer TWF. This timer is used to provide a sampling time for filtering wire fault. The value of TWF shall be between 0 s and 10 s. A value of 5 s is recommended.

3.5 Station policy flags

The station policy flags ("O"—suffix acronym) are set external to the MAC (see 6.1.3.1.2) and are not changed by the MAC state machines. Additional station operation flags are defined in clause 4. The station policy flags define the following policies:

a) Beacon handling method (FBHO)
b) Claim token contention method (FCCO)
c) Early token release support (FETO)
d) Medium rate (FMRO)
e) Token error isolation and reporting support (FTEO)
f) Token handling method (FTHO)
g) Reject remove frame support (FRRO)
h) Priority operation on multiple frame transmissions (FMFTO)
i) Error report timing method (FECO)

j) Good token detection method (FGTO)

3.5.1 Flag, beacon handling option (FBHO)

The flag FBHO is used to indicate how the station handles participation in beaconing prior to the station completing the join ring process. The station fully participates in beaconing as follows.

a) If the value of flag FBHO is 0, upon completion of the Join FSM Neighbor Notification state J5.
b) If the value of flag FBHO is 1, upon completion of the Join FSM Request Initialization state J6.

The FBHO value of 0 is recommended (improves beacon handling in the JS=RQI state).

3.5.2 Flag, claim contender option (FCCO)

The flag FCCO is used to indicate how the station behaves upon reception of a claim token MAC frame with a source address lower than its specific address. The station enters the Transmit Claim Token state when flag FCCO is set to 1. The station does not enter the Transmit Claim Token state when flag FCCO is set to 0. This policy flag is not taken into account if this station is the one that detects the error necessitating entry of the claim process.

3.5.3 Flag, early token release option (FETO)

The flag FETO is used to indicate whether or not the Early Token Release option for 16 Mbit/s operation is selected. If flag FETO is set to 1, the ETR option is active. If flag FETO is set to 0, the ETR option is inactive. For 4 Mbit/s operation, the value of the FETO flag has no significance.

3.5.4 Flag, error counting option (FECO)

A station accumulates errors for a period determined by the error report timer (TER) and reports all errors that occurred during that period. If FECO=0, the station resets TER when the first error is received and, when TER expires, sends an error report frame. If FECO=1, each time TER expires the station resets TER and, if any of the error counters are not zero, sends the error report frame.

3.5.5 Flag, medium rate option (FMRO)

The flag FMRO is used to indicate the operating speed of the ring as follows: If flag FMRO is set to 1, the station is operating at 16 Mbit/s. If flag FMRO is set to 0, the station is operating at 4 Mbit/s.

3.5.6 Flag, multiple frame transmission option (FMFTO)

The flag FMFTO is used to indicate how the station handles the reservation field during the transmission of multiple frames.

a) If flag FMFTO is set to 0, the station examines Pr and Rr each time it transmits the ending frame sequence. For those cases where the queue is not empty, the following occurs:
    1) If Pm<Rr or Pm<Pr, the frame being transmitted is the last frame to be transmitted.
    2) If Pm>=Pr and Pm>=Rr, the next frame from the queue is transmitted.
b) If flag FMFTO is set to 1, the station examines Pr (ignores Rr) each time it transmits the ending frame sequence. For those cases where the queue is not empty, the following occurs:
    1) If Pm<Pr, the frame being transmitted is the last frame to be transmitted.
    2) If Pm>=Pr, the next frame from the queue is transmitted.

The FMFTO value of 0 is recommended (allows high priority reservations to be honored earlier when the station is performing multiple frame transmit per token resulting in lower latency for higher priority traffic).

3.5.7 Flag, reject remove option (FRRO)

A station policy flag that controls how the MAC protocol responds to a Remove frame. If the value of FRRO is 0, then the station will remove from the ring. If the FRRO value is 1, then the station will send a response frame indicating that the function is disabled.

3.5.8 Flag, token error detection option (FTEO)

The flag FTEO is used to indicate how the station handles token errors as follows:

a) If flag FTEO is set to 0, isolation and reporting of code violations in a token is supported.
b) If flag FTEO is set to 1, isolation and reporting of code violations in a token is not supported.

The FTEO value of 0 is recommended (provides additional error checking of ring data).

3.5.9 Flag, token handling option (FTHO)

The flag FTHO is used to indicate how the active monitor station detects the absence of a token or frame.

a) If flag FTHO is set to 0, expiration of the timer TVX is used to detect the absence of tokens or frames.
b) If flag FTHO is set to 1, expiration of the timer TVX is used in conjunction with the FAT flag to detect the absence of tokens or frames. (FAT is defined in clause 4.)

3.5.10 Flag, transmit underrun behavior option (FTUBO)

The flag FTUBO is used to indicate the actions taken by the station when it is transmitting data (TS=DATA) and a STATION_ERR(tx_underrun) condition occurs. If flag FTUBO is set to 0, an abort delimiter is transmitted and the transmit machine returns to repeat (TS=RPT). If flag FTUBO is set to 1, an abort delimiter is transmitted and the transmit machine transitions to FILL (TS=FILL) allowing a token to be released without requiring the active monitor to purge the ring.

3.5.11 Flag, good token option (FGTO)

The flag FGTO is used to indicate how the station detects the presence of a token as follows:

a) If flag FGTO is 0, then good token detection consists of detecting either
    1) Avalid token with priority of 0, or
    2) A valid token with priority greater than 0 followed by a frame.
b) If flag FGTO is 1, then good token detection consists of detecting any valid token.

3.5.12 Flag, wire fault detection option (FWFDO)

The flag FWFDO is used to control at what point, during the join process, the wire fault detection process is enabled. If the flag FWFDO is set to 0, wire fault detection is started when the station enters Join Complete (JS=JC). If the flag FWFDO is set to 1, wire fault detection is started when the station enters Detect Monitor Presence (JS=DMP).

3.6 Error counters

An error counter (“E”—suffix acronym) is incremented when a particular error condition occurs and is set to zero (0) when the contents of the error counter is reported. These counters have a range of 0 to 255. When the count reaches 255 (X'FF'), the count freezes until its value is reported.

3.6.1 Counter, abort error (CABE)

The counter CABE is incremented when a station prematurely ends a transmission by transmitting an abort sequence.

3.6.2 Counter, ac error (CACE)

CACE is incremented when a station receives an AMP or SMP MAC frame in which A and C are both equal to 0, and then receives another SMP MAC frame with both A and C equal to 0 without first receiving an AMP frame.

3.6.3 Counter, burst error (CBE)

A burst5_error is indicated when the station detects the absence of transitions at the receiver input (see 5.4.2). During a single period of signal interruption, multiple burst errors are often indicated due to random noise; in addition, channel characteristics may cause a burst4 error repeated by a station to be detected as a burst5, at the next station. When this condition occurs, even though the station transforms the burst5 into a burst4, each time the burst4 propagates around the ring, it will be detected as a burst5, quickly causing the burst error counter to reach its maximum count. To aid in problem determination, the counter CBE is only required to be incremented once during each interval of signal disruption. The counter may be inhibited after a burst5_error has been indicated until an event is received that indicates the station is receiving a valid signal. At a minimum, the reception of a MAC frame addressed to the station shall enable the counting of CBE. Alternative events include the reception of a SD, the reception of an ED, or either. A station may count every burst5 error.

3.6.4 Counter, frame-copied error (CFCE)

The counter CFCE is incremented when a station recognizes a frame addressed to its specific address and detects that the A bits in the FS field are set to 1, indicating a possible duplicate address. Stations are not required to increment this counter for frames other than MAC frames or for frames that contain an RI field.

3.6.5 Counter, frequency error (CFE)

The counter CFE is incremented when a frequency error is indicated (see 5.7.2).

3.6.6 Counter, internal error (CIE)

The counter CIE is incremented when a station recognizes a recoverable internal error. Reports of this error may be used to identify a station in marginal operating condition.

3.6.7 Counter, line error (CLE)

The counter CLE is incremented when a frame is copied or repeated by a station, the E bit is zero, and the frame is a “FR_WITH_ERR” (see 4.3.2). Optionally, a station may also increment this counter when it repeats a token containing an error as specified in 4.3.1. See 3.5.8.

The first station detecting a line error sets E=1 in the ED of the frame or token to prevent other stations from also counting the error and to enable isolation of the source of the disturbance to a fault domain.

3.6.8 Counter, lost frame error (CLFE)

The counter CLFE is incremented when a station’s TRR timer expires during transmit strip operation. This count indicates how often frames transmitted by a station fail to return to it.

3.6.9 Counter, receive congestion error (CRCE)

The counter CRCE is incremented when a station sets the A bits to 1 in the FS field but does not copy the frame (C bits are not changed).

3.6.10 Counter, token error (CTE)

The counter CTE is incremented by an active monitor when it recognizes an error condition that requires it to purge the ring and restore a token. This occurs as the result of timer TVX expiring or when the active monitor receives a frame or token with the M bit set.

4. Token ring protocols

4.1 Overview

An informative overview of frame transmission, reception, and token ring operation is provided in 4.1 and 4.2. The normative specification of the operation is provided in 4.3. The handling of the undefined frame formats (3.2.3.1) is not described and is a subject for future standardization.

4.1.1 Frame transmission

Station access to the physical medium (the ring) is controlled by passing a token around the ring. When a PDU is requested to be transmitted it is normally queued awaiting the reception of a token.

As a token passes a station it may give that station an opportunity to transmit one or more frames. In order for a station to transmit a frame, the received token must have a priority less than or equal to that of the queued PDU. If the token priority is greater than the priority of the queued PDU, or a frame is repeated on the ring before a token is received, the station requests a token of appropriate priority by setting the R reservation bits in the AC field of the frame or token.

When a token of suitable priority is received, it is changed to a start-of-frame sequence by transmitting the token bit as a one and the M and R bits as zero. At this point the station stops repeating the incoming signal and begins transmitting the queued PDU by sending the FC, DA, SA, RI (if present), and INFO fields. During transmission, the FCS of the frame is calculated and transmitted between the INFO field and the ending delimiter. After transmission of the ending delimiter, a Frame Status field is transmitted with the A and C bits zero followed by the required interframe gap.

4.1.2 Token transmission

After transmission of the last or only frame has been completed, the station releases a token according to the token release station policy flag FETO (3.5.3).

a) *Early token release.* The station releases this token as soon as it has completed transmission of the inter-frame gap.

b) *Normal token release.* The station does not release the token until the station verifies that the last frame transmitted has returned to the station. The station does not release a token if the SA of the frame being stripped is not the same as the SA of the last frame it transmitted. If the SA has been verified at the time the station finishes transmitting the frame, then the token is transmitted at the end of the inter-frame gap; otherwise, the token is not transmitted until the SA is verified.

4.1.3 Stripping

The station removes (strips) all frames from the ring that it originates by transmitting fill until the ending delimiter of the last frame that the station originated is removed from the ring. The station performs the stripping process by counting the number of frames transmitted and not repeating the received symbols (stripping) until the same number of frames have been received.

Early token release and transmitting multiple frames per token require stations to compensate for understripping (failure of a station to strip all of the frames it transmitted) and overstripping (a station stripping the frame following the last frame it transmitted). A single bit error may destroy a starting delimiter or ending delimiter causing a frame to be lost (not counted) which results in overstripping. There are also error conditions that cause delimiters to be created. The most common source for generating erroneous delimiters is the switching action of certain types of TCUs. This TCU switching noise is characterized by multiple burst errors and ending delimiters. If the station stops stripping on the first of a series of erroneous ending delimiters, then not only will the station be understripping, but the following erroneous ending delimiters are detected by subsequent transmitting stations as indications of overstripping. Even though the understrip condition will be corrected by the active monitor (M bit inspection and setting), the ill effects of understripping are far more undesirable than those of overstripping; therefore, understripping should be avoided even at the cost of overstripping. Stripping is performed as follows:

a) *To account for lost frames*. Each time the station transmits a frame, it increments its frame counter and temporarily saves the SA of the frame (only the SA of the last frame transmitted is saved). Each frame received will decrement the frame counter. While the frame counter is greater than 1, the station strips all received frames. When the frame counter is 1 (or less), the station compares the SA of the last frame transmitted to the SA of the received frame. If the frame is not corrupted (SA composed of only Data_zero and Data_one symbols) and the SA's do not compare, the station stops stripping, allowing the ending delimiter to be repeated as an indication of overstripping to the next transmitting station. Otherwise, the station stops stripping after the ED is received, completely removing the frame from the ring.

b) *To account for overstripping by a previous station*. If after the station transmits its first ending delimiter and before receiving a starting delimiter, an ending delimiter is received, the station counts it as the end of its first frame and decrements its frame counter. If only a single frame was transmitted, (frame count=0), the station stops stripping. To prevent understripping, the exception to this rule is when a burst error is detected.

c) *To prevent understripping.* The station does not count frames (ending delimiters) received before the first ending delimiter is transmitted. After the first ending delimiter is received, the station only counts complete frames (i.e., one starting delimiter and one ending delimiter). When a burst error is detected, the station may count the current frame but should not stop stripping until after the next starting delimiter is detected. When the burst error occurs in the last or only frame, the station uses the value of the SA in the next frame (as described above) to determine if the station has lost a frame and thus is overstripping.

4.1.4 Frame reception

4.1.4.1 Frame reception process

While the station is inserted, the station inspects the incoming signal stream and checks it for frames that should be copied or acted upon. If the frame's DA field matches the station's individual address, relevant group address(s), an active functional address, or a broadcast address, the frame (FC, DA, SA, RI, INFO, and FS fields) is copied and subsequently indicated to the appropriate sublayer. While the station is receiving a frame (even if the frame is not being copied), the station checks the frames validity by 1) calculating the FCS and comparing it to the FCS of the received frame, and 2) observing Manchester code violations (presence of a Non-data_J or Non-data_K symbol) between the SD and ED fields. If either of these error conditions exist, the frame's address-recognized bits and frame-copied bits are not set by the station, but the station sets the error bit in the frame's ending delimiter.

4.1.4.2 Bridge interface interactions

To support MAC bridges (ISO/IEC 10038 : 1993), each valid received frame (FR) is indicated at the bridge service interface (M_UNITDATA.indication). If the bridge entity responds with the optional M_UNITDATA.response, then the station sets the A bits, and if the frame is successfully copied, the station sets the C bits. The bridge entity sends a frame by issuing an M_UNITDATA.request and the MAC will queue the frame for transmission.

4.1.5 Priority operation

The current ring service priority is indicated by the priority bits in the AC field of a frame or token on the ring. The priority bits (P) and the reservation bits (R) contained in the AC field work together in an attempt to match the service priority of the ring to the highest priority PDU that is queued for transmission. The value of the priority bits (P) and reservation bits (R) of the most recently received AC field are stored in registers as Pr and Rr, respectively.

The priority mechanism operates in such a way that *fairness* (equal access to the ring) is maintained for all stations within a priority level. This is accomplished by having the same station that raised the service priority level of the ring (*stacking station*) return the ring to the original service priority. To perform this function, two push-down stacks are defined whose top-most values are designated Sr and Sx. The following operations are defined on each stack:

— STACK: Push a value onto the top of the stack.
— POP: Remove the top-most value from the stack and discard. (Note that the top-most value of a stack is assumed to be visible for comparison with other parameters without popping.)
— RESTACK: Pop the top-most value and push a new value onto the stack.

The priority mechanism is explained as follows:

When a station has a priority (a value greater than zero) PDU (or PDUs) ready to transmit, it requests a priority token. This is done by changing the reservation bits (R) as the station repeats the AC field. If the priority level (Pm) of the PDU that is ready for transmission is greater than the R bits, the station increases the value of R field to the value Pm. If the value of the R bits is equal to or greater than Pm, the reservation bits (R) are repeated unchanged.

After a station has captured the token, it transmits PDUs that are at or above the present ring service priority level until it has completed transmission of those PDUs or until the transmission of another frame could not be completed within the allocated token holding time (see 3.4). The station may optionally terminate transmission of frames when the reservation field of the frame being stripped is higher than the pending transmit priority. All transmitted frames are sent at the present ring service priority value. The station will then generate a new token for transmission on the ring.

While the station is receiving the frames that it transmitted, the value of Rr indicates the highest reservation request by the other stations. If the station does not have additional PDUs to transmit that have a priority (Pm) greater than the present ring service priority, and does not have a reservation request (as contained in register Rr) which is greater than the present ring service priority (as contained in register Pr), the token is transmitted with its priority at the present ring service priority and the reservation bits (R) at the greater of Rr or Pm and no further action is taken.

However, if the station has a PDU ready for transmission or a reservation request (Rr), either of which is greater than the present ring service priority, the token is generated with its priority at the greater of Pm or Rr and its reservation bits (R) as 0. Since the station has raised the service priority level of the ring, the station becomes a stacking station and, as such, stores the value of the old ring service priority as Sr and the new ring service priority as Sx (STACK operation). These values will be used later to lower the service priority of the ring when there are no PDUs ready to transmit on the ring whose Pm is equal to or greater than the stacked Sx.

Note that since a station may have raised the service priority of the ring more than once before the service priority is returned to a lower priority (for example, from 1 to 3 and then 5 to 6), it may have multiple Sx and Sr values stored and, hence, be referred to as stacked.

Having become a stacking station, the station captures every token that it receives that has a priority (P) equal to its Sx in order to examine the R bits of the AC field for the purpose of lowering the service priority of the ring. If the station has no PDUs to transmit at the present ring service priority, the new token is transmitted with its P bits equal to either the value of the reservation bits (R), the value of the Sr, or the value of Pm, whichever is greater.

If the value of the new ring service priority is greater than Sr (P equal to Rr or Pm), then R bits are transmitted as 0, the old Sx value is replaced with a new value equal to the new P value, and the station continues its role as a stacking station (RESTACK operation).

However, if the Rr and Pm values are equal to or less than the value of the Sr, the new token is transmitted at a priority value of Sr and both Sx and Sr are removed (popped) from the stacks (POP operation). If no other values of Sx and Sr are stacked, the station discontinues its role as a stacking station. The R bits of the new token are transmitted as the value of Rr or Pm, whichever is greater.

Note that a stacking station that has captured the token may transmit PDUs instead of performing the priority operation as described above. Only those PDUs that have a priority equal to or greater than the ring service priority may be transmitted.

The frames that are transmitted to initialize the ring have a P field that is equal to 0. The receipt of a P field whose value is less than a stacked Sx will clear the stacks.

The complete specification of priority operation is contained in table 7.

4.1.6 Token ring fault detection, reporting, and recovery

Token ring error conditions are classified as either hard errors or soft errors.

a) Hard errors are defined as faults that prevent frames and/or tokens from circulating around the ring.
b) Soft errors are defined as faults that cause data corruption, but do not prevent frames and tokens from circulating around the ring.

Error detection, reporting, and recovery procedures require stations to identify and isolate fault conditions, and, in some cases, require the station causing the fault condition to remove from the ring.

Error conditions are isolated to fault domains (4.1.6.1) previously established by the Neighbor Notification process (4.1.6.2) and reported (4.1.6.3) by beacon MAC frames or Error Report MAC frames. Recovery from error conditions is managed by a strict ring recovery hierarchy (4.1.6.4) using the Ring Purge process, (4.1.6.5), the Claim Token process (4.1.6.6), and the Beacon process (4.1.6.7).

4.1.6.1 Fault domain

A fault domain establishes a boundary around a fault condition and identifies the location of the fault for the appropriate corrective action. The fault domain consists of the following as illustrated in figure 19:

a) The station *downstream* to the fault condition. This station reports the error (G).
b) The station *upstream* to the station reporting the error (F).
c) The components (e.g., ring medium, concentrators, etc.) between the reporting station (G) and its upstream station (F).

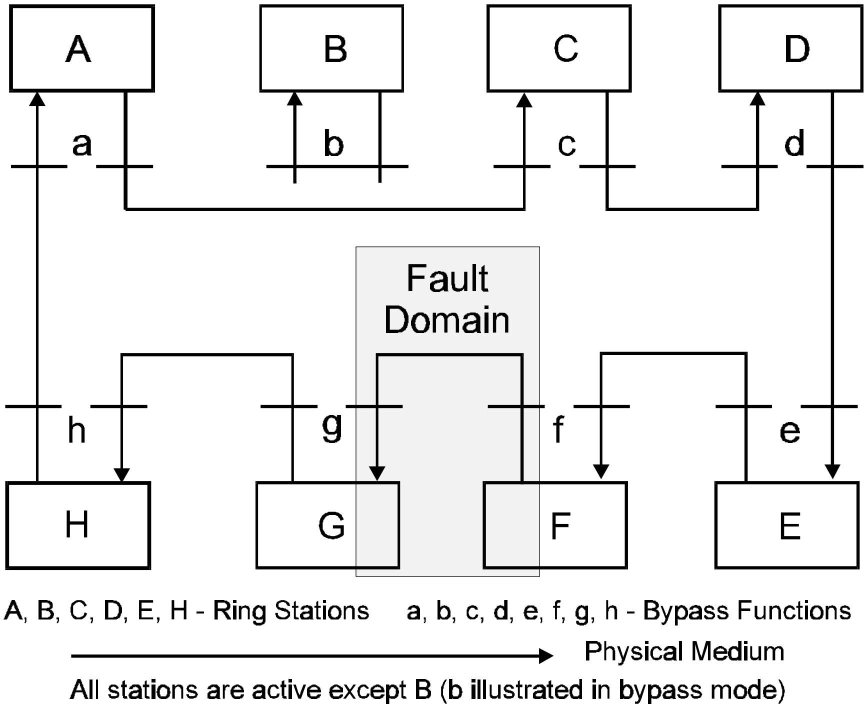

Figure 19—An example of a fault domain

Accurate problem determination is supported by requiring all stations to identify the fault domain when a fault condition is detected. This implies that at any instant in time, each station needs to know the identity of its upstream neighbor station. The Neighbor Notification process (4.1.6.2) provides each station the identity of its upstream neighbor.

4.1.6.2 Neighbor Notification process

The Neighbor Notification process depends on a capability provided by the A bits and the C bits of the FS field. The A and C bits are always transmitted as 0's. If a station recognizes the frame's destination address as a frame it should copy, the station sets the A bits to 1 and, if the station copies the frame, the station sets the C bits to 1.

When a frame is broadcast to all stations on a ring, the first station downstream from the broadcaster is the only station that receives the frame with the A and C bits as 0's. All other stations will see the A bits as 1's since this first station sets the A bits.

The Neighbor Notification process uses this feature by sending the AMP and SMP MAC frames to the broadcast address. If a station receives either of these frames with the A and C bits as 0's, then the frame's SA identifies this station's UNA and this SA is saved as the SUA. If the station is not the active monitor, it then sends a SMP frame to notify its downstream neighbor of its address. This process continues in a circular, daisy-chained fashion to let every station know the identity of its upstream neighbor.

The Neighbor Notification process is performed as follows:

a) The active monitor begins neighbor notification by broadcasting the AMP MAC frame. The station immediately downstream from it takes the following actions:

   1) If possible, copies the broadcast AMP MAC frame and, since the A and C bits are 0's, it stores the upstream station's identity as the SUA;
   2) Sets the A bits (and C bits if the frame was copied) of the AMP MAC frame to 1's;
   3) Resets its timer TQP;
   4) Upon expiration of TQP, the station transmits the SMP MAC frame.

b) One by one, the other stations receive a SMP frame with the A and C bits set to 0's, stores the SUA, resets its TQP timer, and continues the process by broadcasting a SMP frame when its TQP timer expires.

c) The active monitor completes the Neighbor Notification process upon reception of an AMP or a SMP MAC frame with its A and C bits set to 0's. If an AMP MAC frame is received with the A and C bits set to 0, then this is a single station ring.

(The AMP MAC frame is sent by the active monitor on a regular basis allowing all other stations (standby monitors) to detect the presence of the active monitor. Each station has a timer TSM that is reset each time an AMP MAC frame is received. If timer TSM expires, the station begins the Ring Recovery process by entering the Claim Token process.

4.1.6.3 Fault Reporting process

The token ring protocol uses two reporting schemes, one for soft errors and one for hard errors. The following explains the handling of hard errors and soft errors using the fault domain illustrated in figure 19.

Hard errors:

a) The token ring protocol uses the Beacon process in an attempt to recover from hard errors and restore normal token operation. If recovery is not possible, the beacon MAC frame identifies the fault domain for analysis by network management functions.

b) When a hard error is detected between stations F and G, the fault domain is identified and reported by station G transmitting beacon MAC frames. The beacon MAC frame contains the source address of the detecting station (G) and the address of its upstream station (F). This beacon MAC frame allows station (G) to declare the fault domain and alert other stations on the ring that normal token protocol is suspended until the hard error condition terminates or is removed.

Soft errors:

a) The token ring protocol uses either the Ring Purge process or normal token protocol to recover from soft errors and restore normal token operation when necessary.

b) When a soft error is detected between stations F and G, it is reported by station G, which transmits an Error Report MAC frame containing the source address of the detecting station (G), the address of its upstream station (F), and the count of the errors detected (see 4.1.7). The Error Report MAC frame is sent to the Ring Error Monitor for analysis by network management functions.

4.1.6.4 Ring recovery hierarchy overview

The token ring recovery concept is illustrated in figure 20.

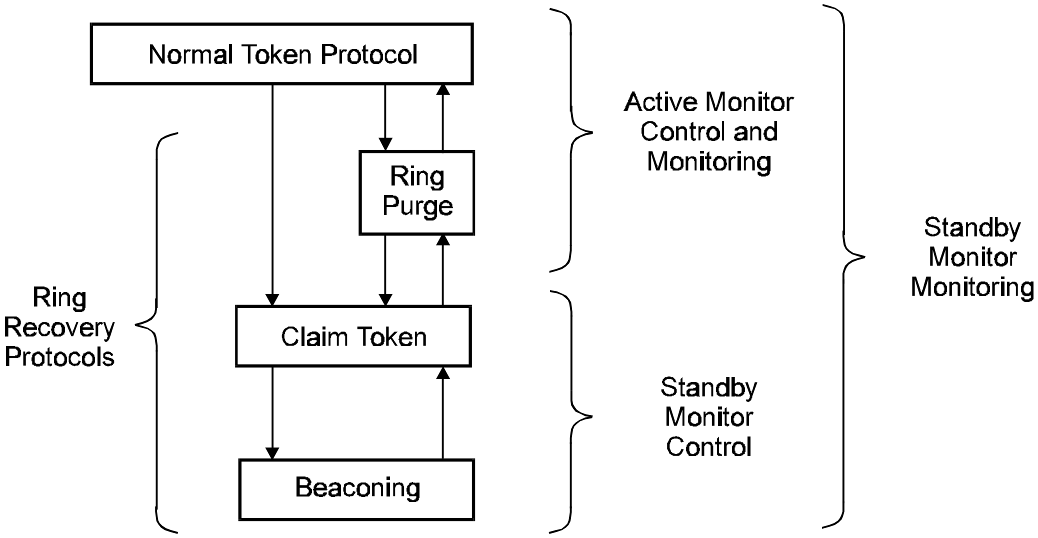

Figure 20—Token ring recovery overview

The token ring recovery concept is a simple one. Once the protocol detects that token passing has been interrupted, the Monitor FSM moves up and down the ring recovery hierarchy without exception as follows:

a) If a process is successful, the monitor FSM moves up the hierarchy.
b) If a process fails, the monitor FSM moves down the hierarchy.

Some examples follow:

a) If the active monitor detects the absence of a token, the Ring Purge process is executed. If the Ring Purge process is successful, the ring returns to normal token protocol operation. If the Ring Purge process fails, the Claim Token process is executed.

b) If a station detects the absence of an active monitor or that the active monitor is not performing its functions properly, the Claim Token process is executed.

c) If the Claim Token process is successful, the Ring Purge process is executed. If the Claim Token process fails, the Beacon process is executed.

4.1.6.5 Ring Purge process

The purpose of the Ring Purge process is to clean up the ring and release a token. The Ring Purge process is started when a station wins the Claim Token process (4.1.6.6) and becomes the active monitor or when the active monitor detects a failure in normal token protocol as follows:

a) There is no token on the ring (i.e., TVX expires without receiving a frame or token);
b) A station failed to strip its own frame from the ring (frame received with M bit set to 1);
c) A stacking station failed to lower the priority token that it released (token received with M bit set to 1).

The Ring Purge process is started when the active monitor enters the Transmit Ring Purge state and transmits Ring Purge MAC frames. When the station is not transmitting a Ring Purge MAC frame it is transmitting FILL. The Ring Purge process operates as follows:

a) Each time the active monitor transmits a Ring Purge MAC frame it resets timer TRR. If TRR expires before the station receives a valid ring purge frame, it transmits another ring purge frame.

b) If the SA of a received Ring Purge MAC frame does not match the station's MAC address, then the station disables its active monitor functions and returns to normal Repeat state.

c) When the active monitor receives its own purge MAC frame, it transitions to the normal Repeat state and transmits a token with the priority bits equal to the reservation bits of the received Ring Purge MAC frame and the reservation set to B'000'. If the priority of the token is greater than zero, the active monitor becomes a stacking station (4.1.5).

d) The active monitor uses the timer TRP to determine if the Ring Purge process fails to complete. Timer TRP is reset when the station first enters the Transmit Ring Purge state. If timer TRP expires while the station is still in Transmit Ring Purge, the station disables its active monitor functions and starts the Claim Token process.

4.1.6.6 Claim Token process

The Claim Token process determines which station becomes the active monitor when the Beacon process (4.1.6.7) is resolved, the Ring Purge process fails to complete, or a station detects that the active monitor functions are not being performed properly.

The Claim Token process is started by a station entering the Transmit Claim Token state and transmitting Claim Token MAC frames. All stations operate as follows:

a) Stations not in the Claim Token process that receive a Claim Token MAC frame either enter the Repeat Claim Token state or the Transmit Claim Token state as follows.

    1) If the station is the active monitor, it enters the Repeat Claim Token state and waits for the Claim Token process to complete. This is an indication that the station has failed to provide proper active monitor functions and should not participate in the Claim Token process nor be an active monitor.

    2) If the station is not the active monitor and is not a claim contender (FCCO=0), it enters the Repeat Claim Token state and waits for the Claim Token process to complete.

    3) If the station is not the active monitor and is a claim contender (FCCO=1), it compares the source address (SA) of the received claim token MAC frame to its own MAC address and takes the following actions.

        i) If the received SA is numerically higher than the station's own MAC address, then the station enters the Repeat Claim Token state and waits for the Claim Token process to complete.

        ii) If the received SA is the same as the station's own MAC address, then the station enters Repeat Claim Token state and waits for the Claim Token process to complete. This is an indication that there is another station on the ring with the same MAC address.

        iii) If the received SA is numerically lower than the station's own MAC address, then the station enters the Transmit Claim Token state, transmits Claim Token MAC frames, and waits for the reception of either its own Claim Token MAC frame or a Claim Token MAC frame with a higher address.

b) Stations in Repeat Claim Token state do not take any recovery action on received claim token MAC frames.

c) Stations in Repeat Claim Token state receiving a Ring Purge MAC frame move to normal token protocol by entering the normal Repeat state.

d) Stations in the Repeat Claim Token state or the Transmit Claim Token state receiving a beacon MAC frame move to the Beacon process by entering the Repeat Beacon state.

e) All stations in the Claim Token process use the timer TCT to determine if the Claim Token process fails to complete. Timer TCT is reset when the station first enters either the Repeat Claim Token or the Transmit Claim Token state. If timer TCT expires while the station is still in one of these states, the station starts the Beacon process by entering the Transmit Beacon state and transmitting beacon MAC frames.

f) Stations in Transmit Claim Token state transmit fill when they are not transmitting a claim token MAC frame and do not repeat frames. Claim Token MAC frames are transmitted without waiting for a token and the timer TQP is used to pace the transmission of the Claim Token MAC frames. When more than one station is in Transmit Claim Token state, a station does not receive its own Claim Token MAC frames.

g) When a station in Transmit Claim Token state receives a Claim Token MAC frame, it compares the source address (SA) of the received Claim Token MAC frame to its own MAC address and takes the following actions:

    1) If the received SA is numerically higher than the station’s own MAC address, then the station enters Repeat Claim Token state and waits for the Claim Token process to complete.

    2) If the received SA is numerically less than is station’s own MAC address, then the station remains in Transmit Claim Token state.

    3) If the received SA is the same as the station’s own MAC address and the UNA reported in the claim token MAC frame does not match the station’s SUA, then the station enters the Repeat Claim Token state and waits for the Claim Token process to complete. The frame’s UNA not matching the station’s SUA is an indication that there is another station on the ring with the same MAC address.

    4) If the received SA is the same as the station’s own MAC address, the UNA reported in the claim token MAC frame matches the station’s SUA, and the M bit is set to 1, then the station remains in the Transmit Claim Token state. The M bit being set to 1 indicates another station is still functioning as the active monitor.

    5) If the received SA is the same as the station’s own MAC address, the upstream neighbor’s address reported in the claim token MAC frame matches the station’s stored upstream neighbor’s address, and the M bit is not set, then the station wins the Claim Token process and starts the Ring Purge process.

A station may require multiple receptions of its own claim token MAC frames before winning the Claim Token process. Requiring multiple Claim Token MAC frames to be sent and received prevents a station from winning claim token prematurely and allows time for other stations to enter the Transmit Claim Token state and transmit its Claim Token MAC frames.

Following the rules stated above eventually all stations enter the Repeat Claim Token state except for the one in Transmit Claim Token state with the highest MAC address. The last station remaining in the Transmit Claim Token state receives its own claim token MAC frames and wins the Claim Token process. This station transitions to Transmit Ring Purge state and becomes the new active monitor. The other stations upon receiving the Ring Purge MAC frame, transition from Repeat Claim Token state to normal Repeat state.

A Claim Token MAC frame sent by one station is not repeated by another station in the Transmit Claim Token state. Therefore, only one station in the Transmit Claim Token state enters the Repeat Claim Token state per frame transmission. Claim Token MAC frames are typically paced at 20 ms intervals. To help speed up the Claim Token process and prevent stations in Repeat Claim Token from prematurely entering the Beacon Transmit state, an option allows a station in Transmit Claim Token state that receives a Claim Token MAC frame resulting in the station entering the Repeat Claim Token state, to transmit the received Claim Token MAC frame (using the SA of the original station).

4.1.6.7 Beacon process

The Beacon process is started by the failure of the Claim Token process (4.1.6.6). The Beacon process relies on the beacon MAC frame and its Beacon Type subvector. At the time the Claim Token process fails, the value of the Beacon Type subvector is determined by examination of the flag, signal loss (FSL) and the flag, claim token (FCT). The following is an explanation of each Beacon Type used in the Beacon process.

Beacon Type 1: Only stations implementing the Dual Ring protocol as specified in IEEE 802.5c [B5] generate beacon type 1 frames. Since Beacon Type 1 is the highest priority beacon type, stations not implementing the Dual Ring protocol will not interfere with the IEEE 802.5c protocol.

Beacon Type 2: The station is detecting a long-term loss of signal (flag FSL is 1). An example of a Beacon Type 2 condition is a break in the medium between the station and its upstream neighbor.

Beacon Type 3: The station has not received any Claim Token MAC frames (flag FCT is 0) but is not detecting loss of signal (flag FSL is 0). The Beacon Type 3 indicates bit streaming (data received does not correspond to expected protocol). An example of a Beacon Type 3 condition is when the upstream station is stuck transmitting fill.

Beacon Type 4: The station has received Claim Token MAC frames (flag FCT is 1) and is not detecting a loss of signal (flag FSL is 0). The Beacon Type 4 indicates a station streaming Claim Token MAC frames. An example of a Beacon Type 4 condition is when the upstream station is stuck transmitting Claim Token MAC frames.

All stations in the Beacon process operate as follows:

Beacon transmit: The station downstream from the fault enters the Transmit Beacon state by transmitting beacon MAC frames that contain the reason for beaconing (Beacon Type) and the address of its last known upstream neighbor (UNA). Stations in the Transmit Beacon state transmit fill when they are not transmitting a beacon MAC frame and do not repeat frames. When more than one station is in the Transmit Beacon state, a station does not receive its own beacon MAC frames. Beacon frames are transmitted without waiting for a token and the timer TQP is used to pace the transmission of the beacon frames.

a) If a station in the Transmit Beacon state receives its own beacon frame, then the Beacon process is resolved and the station enters the Claim Token process.
b) If a station in the Transmit Beacon state receives a beacon frame from another station with an equal or a lower Beacon Type then the station enters the Repeat Beacon state.

When a station first enters the Transmit Beacon state, it resets the TBT timer. When timer TBT expires the station enters Beacon Test state, removes itself from the ring, and performs a test to verify that it is not the source of the fault.

Beacon receive: A station not in Transmit Beacon state that receives a beacon frame either enters the Repeat Beacon state or removes itself from the ring. This procedure eventually causes all stations downstream from the station in Transmit Beacon (downstream to the fault) to enter the Repeat Beacon state.

The Beacon process requires a station that has recently joined the ring and has not yet become known to the stations on the ring (not completed the Neighbor Notification process), to remove from the ring when it either receives a beacon frame or detects a fault condition that would otherwise require it to start the beacon process. This preserves the fault domain established by the last Neighbor Notification cycle. The beacon frame contains the address of the last known upstream neighbor of the beaconing station (UNA). Each station that receives a beacon frame compares the UNA against its own address. Counter CBR is used to count the number of beacon frames which contain an UNA equal to the station's individual address. When the station has received 8 beacon frames in succession with UNA matching its own address (excluding frames with Beacon Type 1), it enters the Beacon Test state, removing itself from the ring, and performing a test to verify that it is not the source of the fault. Since beacon frames are typically paced at 20 ms intervals, this happens within 140–160 ms after receiving the first beacon frame directed at the station. Beacon Type 1 frames are used by dual ring stations during reconfiguration and are excluded from the Beacon Test process.

If a station in Repeat Beacon state receives a claim token MAC frame, the station moves to the Claim Token process entering either the Transmit Claim Token state or the Repeat Claim Token state.

Monitoring functions: The Beacon process contains two monitoring functions: 1) Detecting the absence of beacon MAC frames, and 2) optionally detecting a circulating beacon MAC frame.

a) All stations have a timer TBR that is reset each time a beacon frame is received. If TBR expires, it indicates that no beacon frames are being transmitted and causes the station to start the Claim Token process and enter the Transmit Claim Token state.

b) The downstream station from the beacon transmitter is responsible for monitoring for circulating beacon frames. The counter CBC is used to invoke this function. The counter CBC is set to its initial value when it first enters the Repeat Beacon state and anytime that it receives a beacon frame that was not transmitted by the station's upstream neighbor as determined by the comparison of the frame's SA with the station's stored upstream neighbor's address (SUA). Each time a frame is received that was transmitted by the station's last known upstream neighbor, counter CBC is decremented. When counter CBC reaches zero, the station will inspect and set the M bit in received frames. If the station downstream from the station transmitting the beacon frames receives a frame with the M bit set, the station starts the Claim Token process by entering the Transmit Claim Token state. This prevents a beacon frame transmitted by a station that has removed from continuously circulating around the ring.

4.1.7 Error reporting

When a station detects an error condition on the ring (such as an FCS error, a lost token, or a lost frame), it increments the appropriate error counter (see 3.6). The expiration of the error report timer (TER) causes the station to report the value of the error counters to the Ring Error Monitor (REM) using the Error Report MAC frame (which also identifies the fault domain—see 4.1.6.2) and the error counters are reset. Persistent errors can be detected, isolated, and, if required, necessary action can be taken by management.

4.1.8 Administration of ring parameters

Upon insertion into the ring, a station requests the value of the various ring operating parameters from the Ring Parameter Server (RPS) to assure compatible operation among stations on the ring. If a RPS is not on the ring, the station uses its default value for TER.

4.1.9 Configuration control

As part of the function of maintaining the configuration of the ring, stations notify the Configuration Report Server (CRS) when they detect a change in the address of their upstream neighbor (detected during the Neighbor Notification process). This indicates that either a station has been inserted into or removed from the ring. The CRS can alter the configuration of the ring by requesting stations to remove themselves from the ring. The CRS can also query ring stations for various status information.

4.1.10 Early token release (ETR)

The ETR option increases available bandwidth and improves the data transmission efficiency of the token ring protocol when the frame is shorter than the ring latency. ETR allows a transmitting station to release a token as soon as it completes frame transmission, whether or not the frame header has returned to the station. The priority used for tokens released prior to receiving the frame’s header will be derived from the priority and reservation of the most recently received frame or token (see priority operation).

It should be noted that the access delay for priority traffic may increase when an ETR system is heavily loaded with short frames. Short frames may impede the effect of priority reservations, since a short frame may be transmitted in its entirety and the next token released before the frame’s header returns to the originating station with the ring’s reservation.

Stations implementing ETR option are compatible and interoperable with stations that do not. An enhanced stripping algorithm is required to support early token release (see 4.1.3). This algorithm was not specified in previous versions of this standard, which only specified 4 Mbit/s operation; therefore, early token release is only specified for operation at 16 Mbit/s.

4.1.11 Repeat path

The MAC controls when the station repeats data received from the channel. The received data is repeated at the station’s output except for the following circumstances (see 5.1.2.2 and 5.1.2.4).

a) When modification of bits in the AC, ED, and FS fields of a repeated frame or token is necessary
b) During ring recovery
c) Transmission of a frame
d) Transmission of a token
e) Burst error correction (described in clause 5)

Modification of bits in the repeated frame or token involve

a) *Setting the M bit.* When stations originate frames or tokens, they transmit the M bit in the AC field as B'0'. The active monitor will set the M bit in the AC field to B'1' as specified in 4.3. All other stations repeat the M bit as received.

b) *Setting the R bits.* When stations originate frames, they transmit the R bits in the AC field as B'000'. When stations release tokens, the R bits are set as specified in 4.3. When the station has a frame to transmit, the station modifies the reservation bits in the AC field in a repeated frame or token as specified in 4.3 in order to request a token at the appropriate priority.

c) *Setting the E bit.* When stations originate frames or tokens, they transmit the E bit in the ED as B'0'. The station sets the E bit in the ED as specified in 4.3 when it detects an error in the frame or token.

d) *Setting the A bits.* When stations originate frames, they transmit both A bits in the FS field as B'0's. The station repeating the frame sets both A bits in the FS field to B'1's as specified in 4.3 when it detects that the frame is valid and addressed to the station.

e) *Setting the C bits.* When stations originate frames, they transmit both C bits in the FS field as B'0's. The station repeating the frame sets both C bits in the FS field to B'1's as specified in 4.3 when it detects that the frame is valid, the frame is addressed to the station, and the station successfully copies the frame.

During Ring Recovery and Normal Transmission of a frame, the flag FTI is used to control the repeat path [see 5.1.2.4 PS_CONTROL.request(Repeat_mode)] such that when FTI=1, the received data is not repeated and fill is transmitted:

a) During certain phases of ring recovery, the flag FTI is set to 1 to prevent the station from repeating the received data. During recovery, frames are transmitted without waiting for a token. The flag FTI is temporarily set to 0 to allow the frame to be transmitted, and then set back to 1 until either another frame is transmitted or the Monitor state returns to repeat at which time the FTI flag is set to 0 to enable the repeat path.

b) Capturing a token for transmission involves repeating the SD and P bits of the AC field of the token as received. Starting with the T bit the station stops repeating the token. The station changes the T bit to B'1', sets M to B'0', sets R to B'000', and appends the frame to the end of the AC field (starting with the FC field). After the station has transmitted the frame, flag FTI is set to 1 and the station continues to strip (not repeat) until the FTI flag is set to 0. The station will return to repeating when it

    1) Completely receives all of the frames it had transmitted,
    2) Detects that another station’s frame is being received, or
    3) Timer TRR expires indicating a transmitted frame was lost.

During transmission of a token (TX_TK), a station releases a token at its output independent of the value of the FTI flag. A token is released when

a) The station releases a token after capturing a token and transmitting its frames.

b) The active monitor creates a token after purging the ring.

c) A priority stacking station lowers the ring priority. The exact mechanism to lower the ring priority is not specified. Different mechanisms have been used to lower the priority of the ring. These mechanisms are as follows:

    1) The station changes the old token into an abort delimiter and (after a short delay) creates a new token.
    2) The station changes the old token into a start-of-frame sequence and then (after a short delay) creates a new token.
    3) The station removes the old token and (after a short delay) creates a new token.
    4) The station modifies the AC field as it repeats the token.

The Non-data_J and Non-data_K symbols are only transmitted in pairs and only with the Non-data_J symbol transmitted first. This limits the maximum number of signal elements consecutively transmitted at the same level to 3. When the station detects 4 consecutive elements at the same level, Burst4_error is indicated to the MAC protocol. If more than 4 consecutive elements are received at the same level, Burst5_error is indicated to the MAC protocol and transitions are introduced into the repeated data after the fourth element and the station continues generating transitions until a transition is received (as specified in clause 5). The first station downstream from a signal interruption will detect the Burst5 error condition. The other stations will receive and repeat the burst4. Note that certain physical characteristics may cause a Burst4 repeated by one station to be detected as a Burst5 at the next.

4.1.12 Assured delivery process

The MAC protocol supports assured delivery on Report Error MAC frames, Report SUA Change MAC frames, and Request Initialization MAC frames to guarantee that the MAC frame is received by its destination. The protocol uses the values of the E, A, and C bits as indicated below to decide if the frame is to be retransmitted.

When the station generates a Report Error frame to the Ring Error Monitor (REM) or a Report SUA Change frame to the Configuration Report Server (CRS), it relies on the value of the E, A, and C bits to indicate if the frame was successfully copied. When the frame is first queued for transmission, a counter (CER or CSC) is reset. When the frame returns to the station, the frame status is examined and the station takes the following actions:

a) If the frame fails to return to the station, the frame is re-queued.
b) If the frame status indicates that the frame was copied (C bits set) then the process terminates.
c) If the frame status indicates that the frame was good (E bit not set) and the server does not exist (A bits not set), then the process terminates.
d) If the counter reaches its terminal count (typically 4 tries), the process terminates.
e) If the status indicates that the server did not copy the frame (C bit not set) but either the server was present (A bit set) or the frame had an error (E bit set), then the frame is re-queued and the counter is incremented.

During the Join process, the station generates Request Initialization frames to the Ring Parameter Server (RPS). This process uses 2 counters (CRI & CRIN). The counter CRIN is used to determine that no RPS exists on the ring. When the RPS does exist on the ring, the timer TRI is used to pace the request frames and the counter CRI limits the number of request frames transmitted. When the first Request Initialization frame is queued, CRI, CRIN, and TRI are reset. When the frame returns to the station, the frame status is examined. The station takes the following actions:

a) If the frame fails to return to the station, then another request frame is queued.
b) If the frame had an error (E bit set), then another request frame is queued.
c) If the frame status indicates that the server does exist (A bits set) then the timer TRI is reset.
d) If the frame status indicates that the server does not exist (A bits not set), then counter CRIN is incremented, and if CRIN has not reached terminal count, another request frame is queued.
e) If the counter CRIN reaches terminal count (typically 4 tries) then the server does not exist and the process terminates.
f) If the station receives an Initialize Station MAC frame or a Change Parameters MAC frame, the process terminates.
g) If the timer TRI expires before an appropriate response frame is received, then CRI is incremented and, if CRI has not reached terminal count, another request frame is queued.
h) If the counter CRI reaches terminal count (typically 4 tries), then there is a failure to communicate and the Join process fails.

4.1.13 MAC service interfaces

The LLC, Bridge, and management interfaces are specified to provide the service required by other standards. These interfaces serve to specify operation and do not imply any particular implementation.

4.1.13.1 Service to LLC

On receipt of a MA_UNITDATA.request primitive from the local LLC entity, the MAC shall compose a frame using the parameters supplied to create the FC, DA, SA, RI (if present), INFO (user data) fields and calculates the FCS field. The RII bit shall be set to 1 if and only if routing information is present. The User_priority parameter establishes the requested access priority (Pm) of the frame and is also encoded in the YYY bits of the FC field. The value of the YYY bits shall be either the value of the User_priority parameter or may optionally be set to B'000'. It is recommended that the value of Pm be set to the value of the User_priority parameter. Annex I provides recommendations on the use of access priorities. The frame is then queued for transmission on a suitable token.

If a valid LLC frame is received with a DA that matches the station's individual address, matches one of the station's group addresses, matches an active functional address, or matches one of the broadcast addresses, then a MA_UNITDATA.indication is made to the local LLC entity. The value of the User_priority parameter shall be equal to the value of the YYY bits in the FC field.

4.1.13.2 Service to the bridge

On receipt of a M_UNITDATA.request primitive from the local bridge entity, the MAC shall compose a frame using the parameters supplied to create the FC, DA, SA, RI (if present), INFO (user data), and FCS fields. The RII bit shall be set to 1 if and only if routing information is present. The Access_priority parameter establishes the requested access priority (Pm) of the frame and User_priority parameter is encoded in the YYY bits of the FC field. If the FCS parameter is not specified, then the station calculates the FCS value. It is recommended that the value of Pm be set to the value of the Access_priority parameter. Annex I provides recommendations on the use of access priority. The frame is then queued for transmission on a suitable token.

When a valid frame is received (independent of any address match), a M_UNITDATA.indication is made to the bridge entity. If the frame type is LLC then the value for the User_priority parameter shall be the value of the YYY bits in the FC field, otherwise the value for the User_priority parameter is unspecified.

4.1.13.3 Service to management

On receipt of a MGT_UNITDATA.request primitive from the local management entity, the MAC shall compose a MAC frame using the parameters supplied to create the FC, DA, SA, RI (if present), and INFO (vector) fields, MAC information field, and calculated FCS field. The RII bit shall be set to 1 if and only if routing information is present. The requested access priority (Pm) of the frame shall be the value of the Access_priority parameter. The frame is then queued for transmission on a suitable token.

If a valid MAC frame with a non-zero destination class is received with a DA that matches the station's individual address, one of the group addresses, an active functional address, or one of the broadcast addresses, then a MGT_UNITDATA.indication is made to the MGT entity.

4.2 Specification definitions

The MAC protocol is specified in 4.3.4 as a table of events and actions. There are three basic state machines contained within the MAC protocol. They are the Join State Machine, the Monitor State Machine, and the Transmit State Machine. High level state machine diagrams are presented in figures 23, 24, and 25 as an overview of how the three state machines operate. A detailed set of state diagrams are presented in annex F.

This clause explains the notation used in the finite state machines (FSM), provides an overview of the Join Ring, Monitor, and Transmit FSMs, provides a list of abbreviations used, and defines the flags and counters used in the FSMs and Station Operation tables.

4.2.1 Finite state machine (FSM) notation

The notation used in the FSM diagrams is illustrated in figure 21.

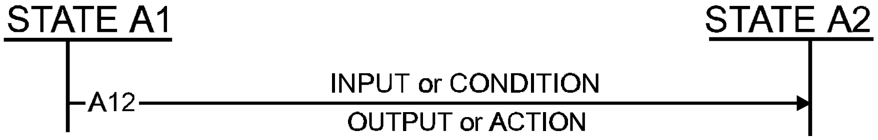

Figure 21—Sample FSM notation

States are shown as vertical lines. Transitions are shown as horizontal lines numbered with a letter representing the state machine and a two-digit number where the first digit represents the starting state and the second digit represents the ending state. For example, A12 indicates the transition from state 1 to state 2 of state machine A. Where there is more than one transition between any two states, a suffix is added to differentiate the transitions (e.g., A12a and A12b). The arrow indicates the direction of transition.

The inputs or conditions shown above the line are the requirements for making the transition. The output or action shown below the line occurs simultaneously with making the transition. The transition begins when the input occurs or the specified condition is met and is complete when the output or action has occurred. If the state transition is in progress, then no other FSM transition may be initiated.

If the conditions for transition out of a state are satisfied at the time the state is entered, no action is taken in that state and the exit transition is taken immediately.

The Station Operation table is an abstract model of the Token Ring Station protocol machine. For all possible combinations of events and station states it determines an unambiguous and consistent response by the station, including next state and primitives at the various service interfaces.

The Station Operation table and FSMs are evaluated as follows:

a) A snapshot of the value of all event terms is taken at an instant in time.
b) All event statements are evaluated as TRUE or FALSE, depending on the truth of each element in the event statements.
c) The actions for all TRUE events are performed.
d) Actions, except for transmit sequences and clock source changes are considered to occur simultaneously.
e) Transmit sequences are started and performed in the order specified.

   If table 7 indicates that a clock source (FTXC) change will occur while TS=DATA, then the clock source change shall be delayed until TS<>DATA. If table 7 indicates that a clock source change will occur at the same time as a transmit sequence starts, the station may delay the clock source change until after the transmit sequence is complete (TS<>DATA).

   The station should delay a transition into TS=DATA by at least 1.5 ms (see 5.7.1) after a clock source change.
f) Evaluation is repeated beginning with Step 1.

Timer=E events will evaluate as TRUE in the snapshot immediately following expiration of the timer, and are never TRUE in more than one snapshot.

The station must have a way of ordering simultaneous external events which would otherwise cause contradictory actions, for example, simultaneous frame arrival and timer expiration. Except for the case of simultaneous events, the actions associated with the set of TRUE event statements in the FSM and Station

Operation table are never contradictory. The method that a station uses to arbitrate simultaneous events is implementation-dependent and not constrained by this standard. The station may act on simultaneous events in any order prior to considering subsequent events.

Actions do not affect the truth of event statements until the next snapshot.

4.2.2 Join Ring FSM / Monitor FSM / Transmit FSM interaction

This clause and its figures explain the interaction of the Join Ring, Monitor, and Transmit FSMs. Functions performed by the FSMs are defined in 4.3. The Join Ring FSM is parent to the Monitor FSM which is parent to the Transmit FSM. Figure 22 illustrates the interaction of the Join Ring, Transmit, and Monitor FSMs.

The following FSM starting conditions (i.e., power on) are assumed:

a) Counters are set to 0.
b) Event Flags are set to 0.
c) Station Policy Flags are set according to station policy.
d) Timers are not running (stopped).

BYPASS. This state is entered as the result of station power-on-reset or as the result of detection of errors by the Join Ring or Monitor FSMs that cause the station to remove itself from the ring. The Transmit and Monitor FSMs are not specified while in the Join Ring bypass state (J0).

4.2.2.1 Join Ring FSM overview

The Join Ring FSM diagram is shown in figure 23. The Join Ring FSM always begins in BYPASS (J0). When the Connect.MAC request is detected, a Lobe Test is performed. During Lobe Test (JS=LT), the Monitor FSM is suspended and the station assumes it is the only entity on the lobe capable of putting data on the lobe. If the Lobe Test completes successfully, the Transmit and Monitor FSMs are activated in states TS=RPT and MS=RPT and the station requests insertion[7] into the ring. After requesting insertion, the Join Ring process will first determine if an active monitor exists on the ring. If no active monitor exists, the Monitor FSM is instructed to enter the Claim Token process. Once an active monitor exists, the DAT is performed to assure that the station’s individual address is unique. Until the address has been verified, the station is not allowed to participate in the Neighbor Notification process. Once the address has been verified, the station is required to participate.

The last phase of the Join sequence is to request the station’s ring parameters from the Ring Parameter Server (RPS). Once the join sequence is complete, the flag FJR is set. Depending on the FBHO policy flag, the FINS flag will be set to 1 either after the Neighbor Notification process has been completed, or after ring parameter initialization. Setting the FINS flag indicates that the station will participate in the token-ring recovery protocols and report ring errors.

The exact functions performed by the Join Ring FSM are specified in 4.3.4.

---

[7] The MAC protocol expects the TCU to insert the station into the ring when ring insertion is requested. The concentrator’s management may direct the TCU not to insert the station into the ring. The MAC protocol may be unable to detect that the concentrator’s management has denied station insertion into the ring. When this occurs, the MAC continues operation as a single station ring, yet higher layer protocols will not be able to establish peer-to-peer connections.

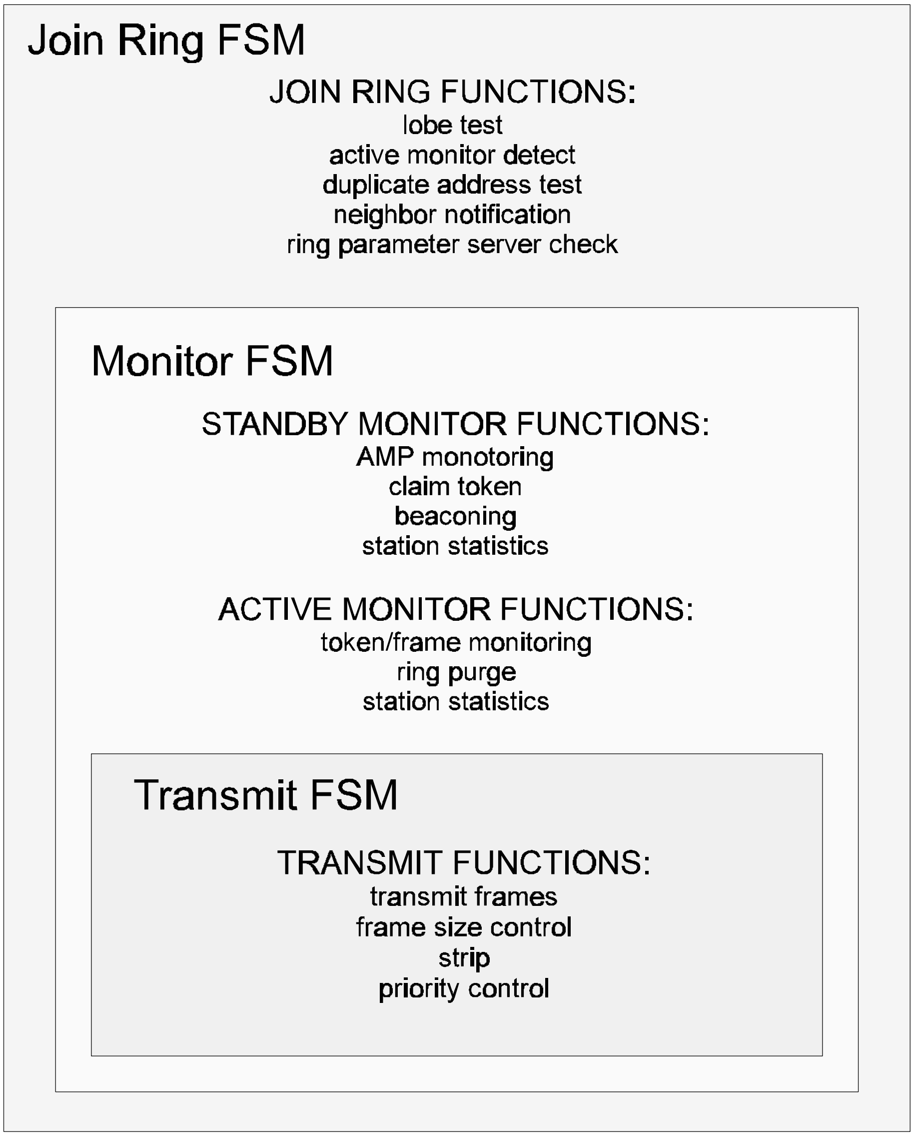

Figure 22—Join Ring, Transmit, and Monitor FSM interaction

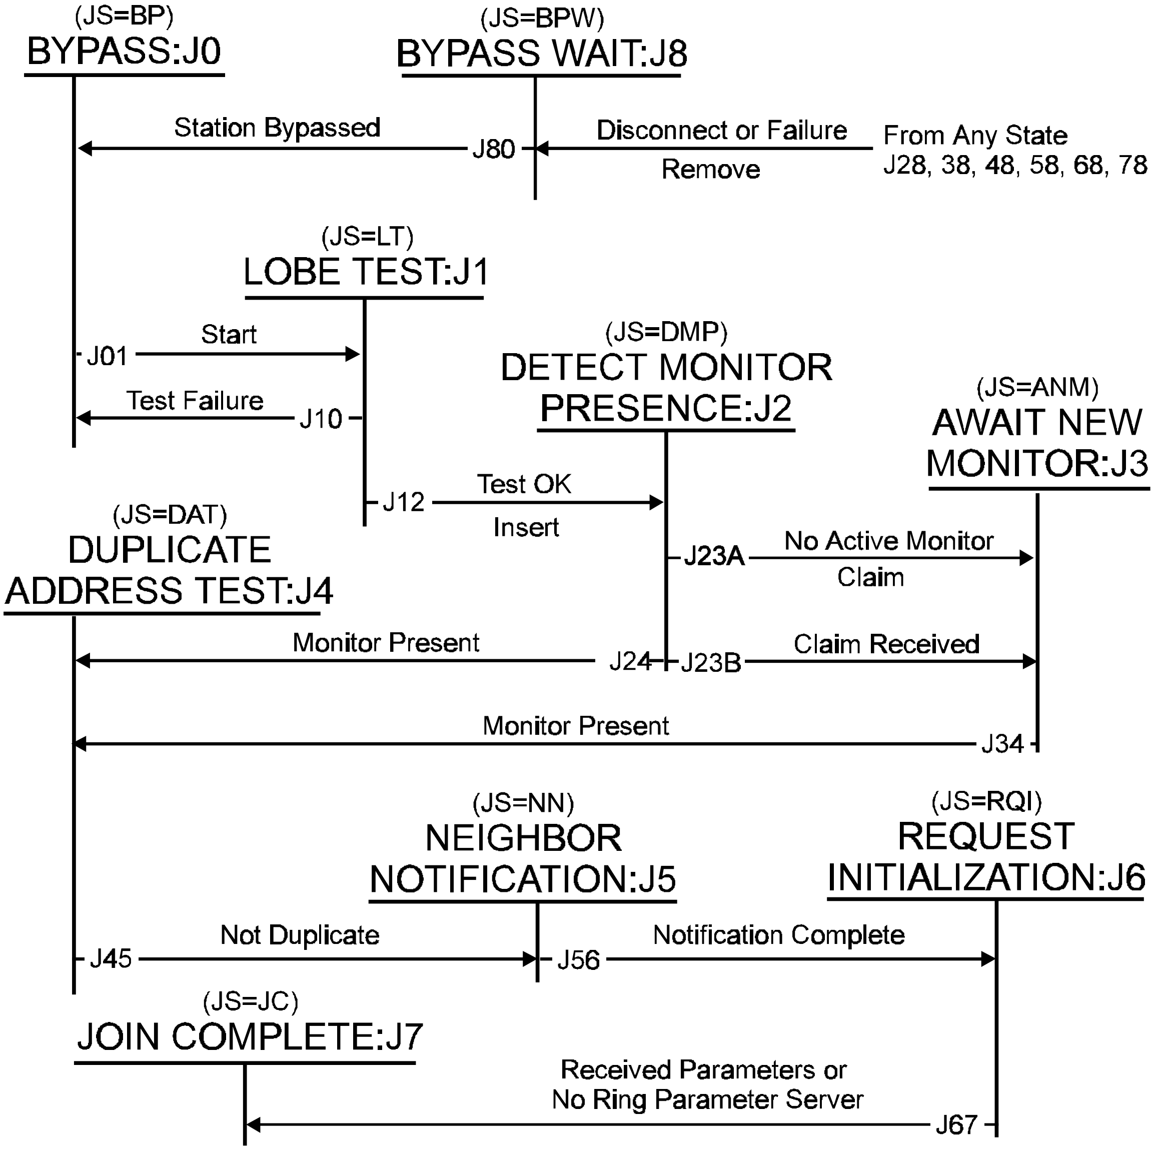

Figure 23—Join Ring FSM diagram

4.2.2.2 Transmit FSM overview

The Transmit FSM diagram is shown in figure 24. The transmit FSM is used to transmit frames following the token capture rules. The Transmit FSM captures a token, transmits the frame(s), strips the frame(s) from the ring, provides fill if necessary, and releases a correct token according to the rules of priority and early token release.

The purpose of recovery is to bring the ring back to a point at which a token can circulate the ring and thus allow stations to capture the token and transmit their frames. The reception of a token while in ring recovery is an unexpected event. A station is not required nor is it prevented from transmitting queued frames during recovery if a token is received.

The exact functions performed by the Transmit FSM are specified in the Station Operation tables (clause 4.3).

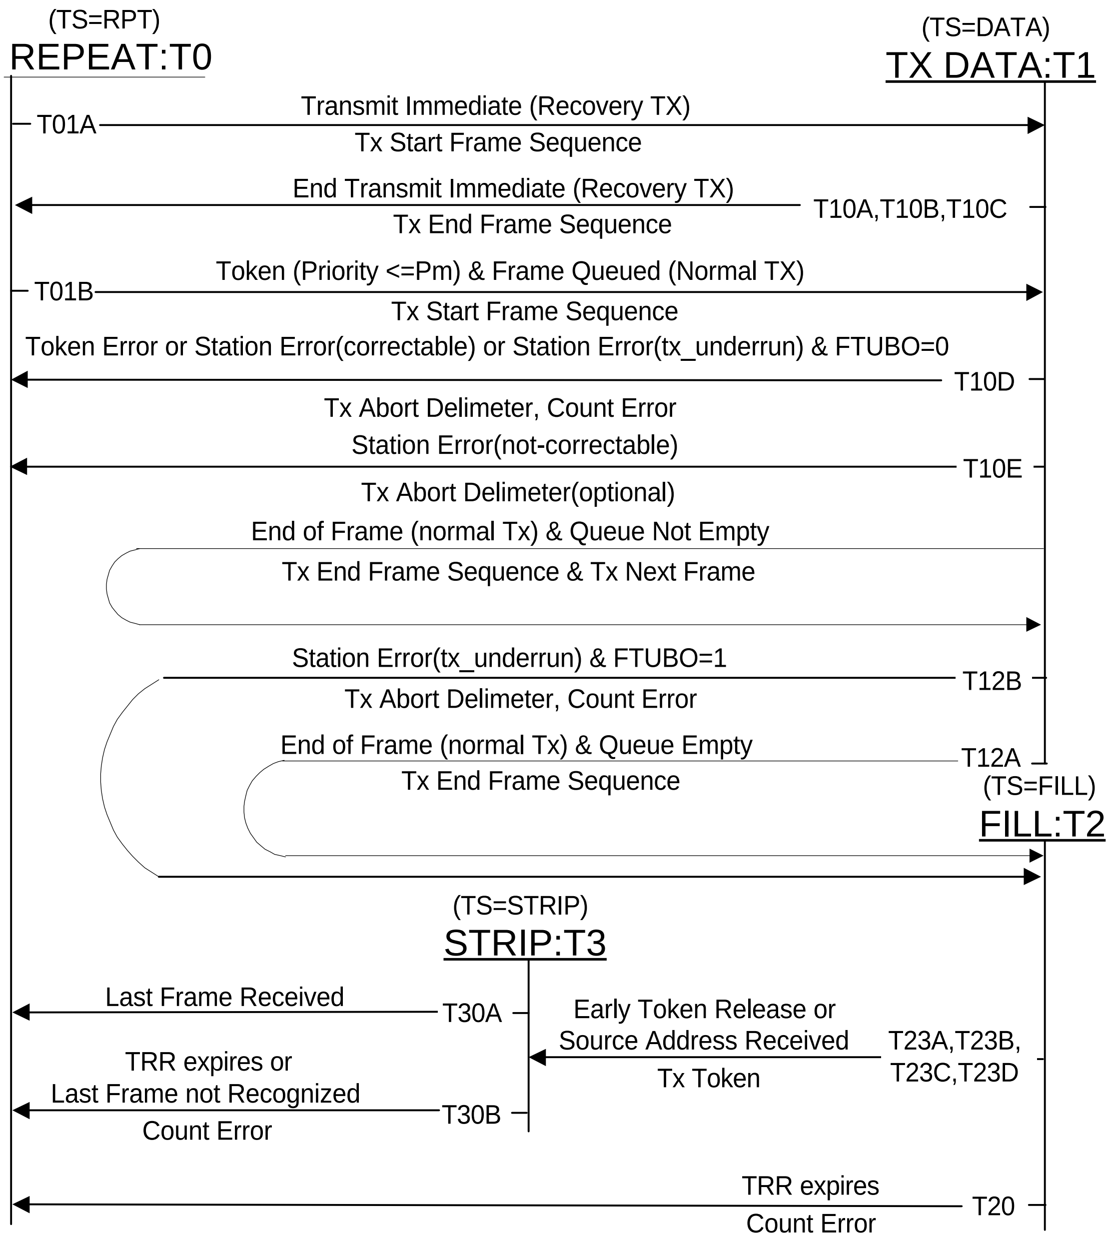

Figure 24—Transmit FSM diagram

4.2.2.4 Monitor FSM overview

The Monitor FSM is shown in figure 25. The Monitor FSM provides both the Standby Monitor and Active Monitor functions.

a) The Standby Monitor functions include monitoring of the active monitor, claiming to establish an active monitor, and beaconing to recover from ring faults.
b) The Active Monitor functions include token/frame monitoring to assure ring utilization and ring purge to recover from the loss of a valid token.

During normal token-ring operation, the Active Monitor functions are active in ONLY ONE station on the ring and the Standby Monitor functions are active in ALL other stations on the ring. However, during the claim token and beacon operations, ALL stations are using their Standby Monitor Functions.

The Monitor FSM shown in figure 25 has three states in which the station repeats received data: 1) Normal Repeat (MS=RPT), 2) Repeat Claim Token (MS=RCT), and 3) Repeat Beacon (MS=RBN).

a) In the normal repeat state, the station is monitoring the ring’s normal protocol operation.
b) In the repeat claim token or repeat beacon states, the station is not actively participating in ring recovery (that is, not transmitting a Ring Purge, Claim Token, or Beacon MAC frame). Instead the station is monitoring the ring for correct claim token and beacon operation.

When the Monitor FSM is not in one of these repeat states, the station is either transmitting one of the recovery frames or testing itself and does not repeat data. Frames received are examined to determine FSM’s actions. When testing itself, the station is not inserted in the ring.

The exact functions performed by the Monitor FSM are specified in table 7.

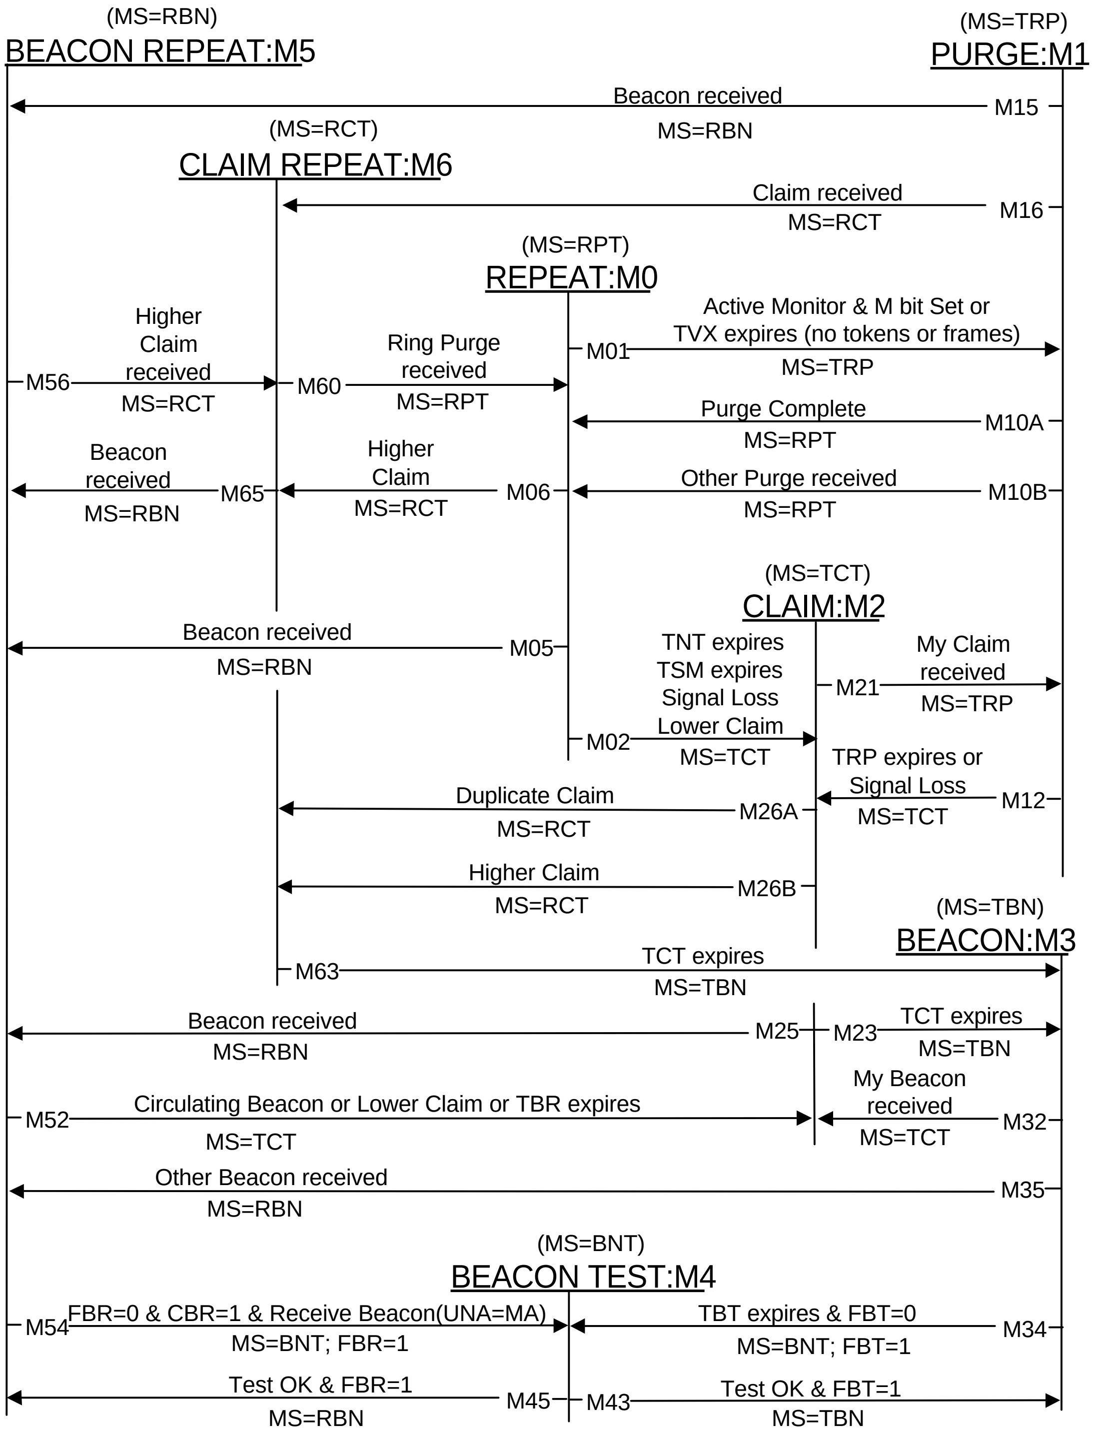

Figure 25—Monitor FSM diagram

4.2.3 Abbreviations and notations

The following list contains abbreviations and notations used in table 7 and in state machine descriptions.

Station policy flag notations

FBHO = Flag, beacon handling option
FCCO = Flag, claim contender option
FETO = Flag, early token release option
FECO = Flag, error counting option
FGTO = Flag, good token option
FMFTO = Flag, multi-frame transmit option
FMRO = Flag, medium rate option
FRRO = Flag, reject remove option
FTEO = Flag, token error option
FTHO = Flag, token handling option
FTUBO = Flag, transmit underrun behavior option
FWFDO = Flag, wire fault detection option

MAC protocol flag notations

FAM = Flag, active monitor
FANM = Flag, await new monitor
FAT = Flag, any token received
FBPF = Flag, bypass force
FBR = Flag, repeat beacon test
FBT = Flag, transmit beacon test
FCT = Flag, claim token
FDC = Flag, duplicate address test complete
FED = Flag, ending delimiter
FER = Flag, error report required
FID = Flag, insert delay
FINS = Flag, inserted
FJR = Flag, join ring process complete
FLF = Flag, lost frame
FMA = Flag, my address received
FMP = Flag, monitor present
FNC = Flag, neighbor notification complete
FNMA = Flag, not my address received
FNN = Flag, neighbor notification
FNW = Flag, neighbor notification waiting
FOP = Flag, operational
FPT = Flag, priority token
FRH = Flag, remove hold
FSL = Flag, signal loss
FSLD = Flag, signal loss detected
FSD = Flag, starting delimiter
FSMP = Flag, standby monitor present
FTI = Flag, transmit idles
FTW = Flag, test wait
FTXC = Flag, crystal transmit
FTXI = Flag, transmit immediate
FWF = Flag, wire fault
FWFA = Flag, wire fault active

Counter notations

CABE = Counter, abort error
CACE = Counter, AC error
CBC = Counter, beacon circulating
CBE = Counter, burst error
CBR = Counter, beacon repeat
CCR = Counter, claim receive
CCT = Counter, claim transmit
CDF = Counter, DAT failure
CDG = Counter, DAT good
CER = Counter, error report
CFCE = Counter, frame-copied error
CFE = Counter, frequency error
CFR = Counter, frames (frames to be stripped)
CIE = Counter, internal error
CLE = Counter, line error
CLFE = Counter, lost frame error
CRCE = Counter, receive congestion error
CRI = Counter, request initialization
CRIN = Counter, RQ_INIT not recognized
CNNR = Counter, neighbor notification requests
CSC = Counter, SUA change
CTE = Counter, token error
CTO = Counter, transmitted octets

Timer notations

TAM = Timer, active monitor
TBR = Timer, beacon repeat
TBT = Timer, transmit beacon test
TCT = Timer, claim token
TER = Timer, queue error report
TID = Timer, insert delay
TJR = Timer, join ring
TNT = Timer, no token
TQP = Timer, queue PDU
TRH = Timer, remove hold
TRI = Timer, ring initialization
TRP = Timer, transmit ring purge
TRR = Timer, return to repeat
TRW = Timer, remove wait
TSL = Timer, signal loss
TSM = Timer, standby monitor
TVX = Timer, valid transmission
TWF = Timer, wire fault
TWFD = Timer, wire fault delay

Join state notations (JS=)

ANM = await new monitor
BP = bypass
BPW = bypass wait
DAT = duplicate address test
DMP = detect monitor presence
JC = join complete
LT = lobe test

NN = neighbor notification
RQI = request initialization

Transmit state notations (TS=)

DATA = Transmit frame data
FILL = Transmit fill and release token
RPT = Repeat
STRIP = Strip transmitted frames and transmit fill

Monitor state notations (MS=)

BNT = beacon test
RBN = repeat beacon
RCT = repeat claim token
RPT = normal repeat
TBN = transmit beacon
TCT = transmit claim token
TRP = transmit ring purge

4.2.4 State machine elements

The state machines use the following counters, flags, and states to describe the operation of the station. These are logical elements used solely to describe the operation and do not specify an implementation. The value of the flags and counters are only meaningful internal to the state machine definition. Conformance will only be based on the station's ability to perform the protocol as specified by table 7.

4.2.4.1 Counters

Unless otherwise specified, all counters are set to 0 by the "Set_initial_conditions" action as the result of the Join Ring FSM transition detecting the Connect.MAC request.

A counter may be set to a value, counted up (incremented), or counted down (decremented) as a result of an action specified in the station operation table.

Counter, Beacon Circulating (CBC). The Monitor FSM uses the counter CBC to count the number of beacon MAC frames received from the station's upstream station before activating the circulating beacon frame detection (4.1.6.7).

Counter, Beacon Repeat (CBR). The Monitor FSM uses the counter CBR to control the number of beacon MAC frames received from its downstream neighbor before removing itself from the ring and entering the Beacon Test state.

Counter, Claim Receive (CCR). The Monitor FSM uses the counter CCR to count the number of Claim Token MAC frames transmitted by this station (MS=TCT) that have been successfully received and thus allow resolution of the Claim Token process.

Counter, Claim Transmit (CCT). The Monitor FSM uses the counter CCT to count the number of times the Claim Token MAC frame is transmitted by this station (MS=TCT).

Counter, DAT Failure (CDF). The Join Ring FSM uses the counter CDF to count the number of times the DAT MAC frame transmitted by this station is received with A=1. This counter is initialized to 0 by the Join Ring FSM.

Counter, DAT Good (CDG). The Join Ring FSM uses the counter CDG to count the number of times the DAT MAC frame transmitted by this station is received with A=0 and C=0. This counter is initialized to 0 by the Join Ring FSM.

Counter, Error Report (CER). The counter (CER) is used to control the number of times the Error Report MAC frame will be retransmitted by the assured delivery process.

Counter, Frame (CFR). The transmit FSM uses the counter CFR to indicate the number of frames originated by the station which, by station calculation, are still on the ring. Counter CFR is incremented when a SFS is transmitted. After the station has transmitted its first Ending Delimiter (FED=1), CFR is decremented when a Starting Delimiter has been detected (FSD=1) and the ED of a frame is stripped. This counter is defined for stations that transmit multiple frames per token.

Counter, Neighbor Notification Requests (CNNR). The neighbor notification counter counts the number of neighbor notifications since the last ring purge for the purpose of inhibiting counting neighbor notification errors due to normal recovery protocol.

Counter, Request Initialization (CRI). The Join Ring FSM uses the counter CRI to limit the number of RQ_INIT frames transmitted to the ring parameter server that were recognized by the RPS (A bits set).

Counter, Request Initialization Not recognized (CRIN). The join ring FSM uses the counter CRIN to limit the number of RQ_INIT frames transmitted to the ring parameter server that were not recognized by the RPS (A bits not set).

Counter, SUA Change (CSC). The counter (CSC) is used to control the number of times the SUA Change MAC frame will be retransmitted by the assured delivery process.

Counter, Transmitted Octets (CTO). The counter CTO is used by the Transmit FSM to limit the number of octets that can be transmitted after capturing the token and thus limit the token holding time of the station. Depending on the medium rate, the counter CTO is pre loaded with a value that indicates the maximum number of octets the station may transmit before releasing a token. Once a station has acquired the token, the CTO counter is decremented every 8 symbols (octet). A station may continue to transmit frames smaller than the remaining value of CTO. A frame with a size larger than the remaining value of CTO will be deferred until the next token capture. The maximum octet count values of CTO (based on a token holding time of 9.1 ms) at 4 Mbit/s is 4550 and at 16 Mbit/s is 18 200. A station may support any value for its maximum octet count (MAX_TX) that does not exceed these maximum octet count values. In addition, a station is not required to transmit more than one frame per token opportunity.

4.2.4.2 MAC protocol flags

Flags are used to remember occurrence of an event for later action by the FSMs and the Station Operation Table (table 7) and are not meant to imply any implementation requirements. In general, a flag is set to 1 when a condition occurs and set to 0 when the condition no longer exists or the appropriate action is taken. All MAC protocol flags (not station policy flags) are set to 0 by the “Set_initial_conditions” action of the Join Ring FSM.

NOTE—FTXC and FTI require specific actions when the flags are set.

The following flags, listed alphabetically, are defined:

Flag, Active Monitor (FAM). The flag FAM is set to 1 when the station wins the Claim Token process and becomes an active monitor. FAM is set to 0 when the station enters the standby monitor mode.

Flag, Await New Monitor (FANM). The flag FANM is only set to 1 when the station is joining the ring and has not detected an active monitor function on the ring. This allows the station to participate in the Claim Token process since it may be the only station on the ring.

Flag, Any Token (FAT). The flag FAT is used by a station in its role as Active Monitor (FAM=1). FAT is set to 1 when the station detects a frame or a token. FAT is tested and set to 0 upon expiration of TVX (timer valid transmission).

Flag, Bypass Force (FBPF). When a station in the Transmit Data state (TS=DATA), detects the "STATION_ERR(not_correctable)" condition, an abort sequence is optionally transmitted and the flag FBPF is set to 1. The join machine enters the Bypass Wait state (JS=BPW) upon detecting flag FBPF=1 and TS=RPT.

Flag, Beacon Repeat (FBR). The flag FBR is set to 1 when the station enters the beacon test state from the repeat beacon state to prevent multiple entries into the beacon test function and to indicate that the station should return to the repeat beacon state upon successful completion of beacon test.

Flag, Beacon Test (FBT). The flag FBT is set to 1 when the station enters the beacon test state from the transmit beacon state to prevent multiple entries into the beacon test state and to indicate the station should return to the transmit beacon state upon successful completion of beacon test.

Flag, Claim Token (FCT). The flag FCT is set to 1 upon reception of a claim token MAC frame for the purpose of determining the beacon type subvector (TX_BN_TYPE) of the beacon MAC frame when entering the transmit beacon state. FCT is set to 0 when entering the transmit claim token state.

Flag, Duplicate Address Test Complete (FDC). FDC is initially set to 0 the first time the station enters the ring as a result of the Join Ring FSM. The flag FDC is set to 1 upon completion of the duplicate address test (JS=DAT) and allows transition to the Neighbor Notification process (JS=NN). The flag FDC is also used by IEEE Std 802.5c-1991 [B5] to indicate whether or not DAT has completed for this MAC. See IEEE Std 802.5c-1991 for more information.

Flag, Ending Delimiter (FED). The FED flag is used to synchronize the frame counter (CFR) with the transmission of the first frame (ending delimiter). Flag FED is set to 0 upon entry into the transmit frame data state (in conjunction with the start of transmission of the first frame after token capture). The flag FED is set to 1 upon the transmission of a frame’s ending delimiter after token capture thus allowing CFR to be decremented by each received frame’s ending delimiter.

Flag, Error Report (FER). The flag FER is set to 1 when the first reportable error is detected and indicates that subsequent errors should not reset the error timer TER. Flag FER is set to 0 when the error timer expires and the Report Error MAC frame is transmitted.

Flag, Insert Delay (FID). Flag FID is used when returning from the beacon test state to assure sufficient time has elapsed to allow the PHY to insert the station in the ring’s data path (no longer wrapped at the TCU). Flag FID is set to 1 causing the station in the repeat beacon state (MS=RBN) or the transmit beacon state (MS=TBN) to wait for a) expiration of timer TID, b) reception of any frame, or c) the detection of filtered signal loss. Once one of these conditions is recognized, flag FID is set to 0 allowing the station to continue with normal resolution of the Beacon process.

Flag, Inserted (FINS). Flag FINS is initially set to 0 as a result of the Join Ring FSM. The flag FINS is set to 1 to indicate when the station may participate in beacon functions. A station may activate this capability either at the completion of neighbor notification (JS=NN) or at the completion of request initialization (JS=RQI) depending on the setting of the policy flag FBHO.

Flag, Join Ring (FJR). FJR is initially set to 0 as a result of the Join Ring FSM. The flag FJR is set to 1 upon completion of request initialization phase (JS=RQI).

Flag, Lost Frame (FLF). Flag FLF is used to count lost frame errors. The flag FLF is set to 1 at the time the station stops stripping if the transmitted frame had failed to return to the station. FLF is set to 0 when the CLFE error counter is incremented.

Flag, My Address (FMA). Flag FMA is used to indicate if the frame being stripped is actually the last frame transmitted. Flag FMA is set to 0 upon entry to the transmit frame data state (prior to transmission of the first frame after token capture), when a SD is received, and when an ED is received. When the frame counter indicates that the last transmitted frame is being stripped (CFR=1), the flag FMA is set to 1 upon receiving a frame with SA that is equal to this station’s last transmitted source address.

Flag, Monitor Present (FMP). Flag FMP is initially set to 0 as a result of the Join Ring FSM. The flag FMP is used while joining the ring and is set to 1 upon detecting the ring has an active monitor (receipt of a RP, AMP, or SMP frame).

Flag, Neighbor Notification Complete (FNC). Flag FNC is initially set to 0 as a result of the Join Ring FSM. The flag FNC is set to 1 upon completion of Join Ring state J5 (JS=NN) indicating neighbor notification has been completed (the station now knows its UNA).

Flag, Not My Address (FNMA). The flag FNMA is used to prevent overstripping caused by bit errors during multiple frame transmission and early token release. Flag FNMA is set to 0 upon entry to the transmit frame data state (prior to transmission of the first frame after token capture), when a SD is received, and when an ED is received. When the frame counter indicates that the last transmitted frame is being stripped (CFR=1), the flag FNMA is set to 1 if the SA of the received frame is not equal to the SA of the last frame transmitted. This flag is set to 1 only if the SA is composed of only 0 and 1 bits and the counter CFR is equal to 1. When this flag is set, it indicates the station has started to strip a frame sent from another station and the station returns to repeat allowing the other station to successfully count the ending delimiter. This flag is not required to be set to one if early token release is not implemented (FETO=0) and the station only transmits one frame per token opportunity.

Flag, Neighbor Notification (FNN). Flag FNN is only used by the active monitor. FNN is set to 0 when the active monitor starts the Neighbor Notification process. FNN is set to 1 when the active monitor receives an AMP or SMP frame with both the A and C bits set to 0, indicating the Neighbor Notification process has completed. When the active monitor starts a new neighbor notification cycle, the value of FNN indicates if the previous cycle failed to complete. If FNN is already set to 1 when a SMP MAC frame is received with both the A and C bits set to 0, it is considered an error.

Flag, Neighbor Notification Waiting (FNW). The station is inhibited from participating in neighbor notification until it has passed the DAT. Prior to DAT completion, Flag FNW is used to remember that the station has received a neighbor notification frame and that its neighbor is waiting. Flag FNW is initially set to 0 as a result of the Join Ring FSM. The flag FNW is set to 1 as the result of the reception of an appropriate AMP or a SMP MAC frame (A=0).

Flag, Operational (FOP). The flag FOP indicates when the station is operational. When FOP is set to 0 the station is not required to respond to signals at its receiver input. Flag FOP is initially set to 0 and is set to 1 at the time the station completes Lobe Test (JS=LT) and inserts into the ring. Flag FOP is set to 0 at the time the station removes from the ring (JS=BPW). FOP is also set to 0 when the station removes for testing (MS=BNT) and set back to 1 when the station inserts after successful completion of the beacon test process.

Flag, Priority Token (FPT). The flag FPT indicates the presence of a priority token (P>0). If a priority token is received, FPT is set to 1. If a frame or token with P=0 is received, FPT is set to 0.

Flag, Remove Hold (FRH). The flag FRH is set to 1 as the result of the Join Ring FSM. The flag FRH is set to 0 to indicate that the station has waited the required minimum time before removing from the ring.

Flag, Signal Loss (FSL). The flag FSL indicates the reception, or lack thereof, of valid signal from the ring [see 5.1.4.1, PM_STATUS.indication(Signal_detection)]. FSL is set to 1 when loss of signal has occurred for the period of TSL (filtered signal loss). FSL is set to 0 upon detection of valid signal.

Flag, Signal Loss Detected (FSLD). The flag FSLD is used to determine whether or not the PHY signal is a steady-state condition. FSLD is set to 0 upon receiving PM_STATUS.indication (Signal_detected=Signal_acquired), thus making the signal loss detection process inactive. FSLD is set to 1 upon receiving PM_STATUS.indication(Signal_detected=Signal_loss). When FSLD is set to 1, the signal loss detection process is active.

Flag, Standby Monitor Present (FSMP). Flag FSMP is only used by standby monitors. The flag FSMP is used to detect neighbor notification failures that would cause excessive neighbor notification frames. FSMP is set to 0 upon reception of an AMP MAC frame with the A and C bits not equal to 0 (indicating the start of a neighbor notification cycle). FSMP is set to 1 by the standby monitor upon receiving a SMP MAC frame with both the A and C bits set to 0, or an AMP MAC frame with both the A and C bits set to 0. If a subsequent SMP MAC frame with A and C bits set to 0 is received before the next neighbor notification cycle is started (FSMP=1), it is considered an error.

Flag, Starting Delimiter (FSD). The flag FSD is used when counting frames to indicate the presence of a frame. Flag FSD is set to 0 upon entry into the transmit frame data state (in conjunction with the start of transmission of the first frame after token capture) and upon the reception of an ending delimiter. The station may optionally set FSD to 0 when it detects a burst4 error to prevent a premature end to stripping. The flag FSD is set to 1 upon the reception of a starting delimiter. The station may optionally set FSD to 1 when the station transmits its first ED to account for the previous station overstripping the start of the first frame.

Flag, Test Wait (FTW). The flag, test wait (FTW) is used by the station when the repeat beacon function detects the need to execute the beacon test function. Using flag, FTW, the station ensures the concentrator has completed its insertion cycle before the station executes Remove_station. This flag is set to 0 when the station 1) sets its initial conditions and 2) after the station has determined that the remove hold function is complete. This flag is set to 1 when the repeat beacon function detects the need to enter the beacon test function.

Flag, Transmit Idles (FTI). The flag FTI is used to control the transmission of idles (fill). When flag FTI is set to 1, the MAC indicates PS_CONTROL.request(Repeat_mode=Fill) which causes the station to source fill rather than repeating the received data. When flag FTI is set to 0, the MAC indicates PS_CONTROL.request(Repeat_mode=Repeat), which causes the station to repeat the received data.

Flag, Transmit from Crystal (FTXC). The flag FTXC is used to select the station’s transmitter timing reference. FTXC is also used to indicate when the station’s fixed latency buffer is required to assure that the token can exist on the ring. FTXC is set to 1 as a result of entering and set to 0 upon exiting the FSM states requiring crystal transmit (transmit claim token, transmit beacon, or active monitor functions).

When the FTXC flag is set to 1, the MAC indicates PS_CONTROL.request(Crystal_transmit=Asserted) which causes the station (as specified in 5.8) to use its crystal clock as the transmit timing reference, provides an elastic buffer, and inserts a fixed latency buffer with a latency at least as large as a token in the data path.

When FTXC is set to 0, the PS_CONTROL.request(Crystal_transmit=Not_asserted) indicates that the station’s transmit timing reference is derived from the station’s clock recovery circuit and that the elastic and fixed latency buffers are no longer required and are removed from the data path.

Flag, Transmit Immediate (FTXI). The flag FTXI is used to indicate whether the station transmits frames with or without waiting for a token. When flag FTXI is set to 1, frame transmit occurs without waiting for a token. When flag FTXI is set to 0, frame transmit occurs after the appropriate token has been captured and indicates that a token needs to be released.

Flag, Wire Fault (FWF). Flag FWF is set to 1 to indicate that wire fault is present [see 5.1.4.1 PM_STATUS.indication (Wire_fault)]. If no wire fault is detected, FWF is set to 0.

Flag, Wire Fault Active (FWFA). Flag FWFA is set to 1 to activate wire fault detection and is set to 0 to deactivate wire fault detection.

4.2.4.3 States

There are a set of states for the Join Ring FSM, the Monitor FSM, and the Transmit FSM. An FSM can be in only one state at any instant in time.

4.2.4.3.1 Join states

The join state (JS=) notation is used to identify the current state of the Join Ring FSM. The join state values are Bypass, Lobe Test, Await New Monitor, Detect Monitor Presence, Duplicate Address Test, Neighbor Notification, Request Initialization, Join Complete, and Bypass Wait. During the Bypass and Bypass Wait states, normal operation is suspended and no assumptions can be made regarding the transmission or reception of data. Join states, listed by state value, are defined as follows:

Join State J3, Await New Monitor (JS=ANM) The join state JS=ANM is set when the join ring FSM enters the Await New Monitor state (J3). In this state the station waits for the Claim Token process to establish an active monitor.

Join State J0, Bypass (JS=BP) The join state JS=BP is set when the join ring FSM enters the bypass state (J0). In this state the station waits for the connect request and is not required to perform any other function.

Join State J8, Bypass Wait (JS=BPW) The join state JS=BPW is set when the join ring FSM enters the bypass wait state (J8). In this state the station must repeat received data and the station remains in this state until the TCU has had ample time to remove the station.

Join State J2, Detect Monitor Presence (JS=DMP) The join state JS=DMP is set when the join ring FSM enters the detect monitor presence state (J2). In this state the station waits for an indication that an active monitor exists on the ring.

Join State J4, Duplicate Address Test (JS=DAT) The join state JS=DAT is set when the join ring FSM enters the duplicate address test state (J4). In this state the station sends DAT MAC frames to verify that there is not another station on the ring that is using the same MAC address.

Join State J7, Join Complete (JS=JC) The join state JS=JC is set when the join ring FSM enters the join complete state (J7). The station has successfully joined the ring and is fully participating in all ring station functions.

Join State, Lobe Test J1, (JS=LT) The join state JS=LT is set when the join ring FSM enters the lobe test state (J1). Upon entering the lobe test state, the station shall transmit valid fill for a minimum of 20 ms. The station shall then perform a test to verify the operation of the lobe. The lobe test shall fail if the bit error rate does not meet the requirements defined in annex P. Annex P contains recommendations for implementing lobe media tests that meet the bit error rate requirements. While the station is in JS=LT, it shall transmit only valid frames, tokens, or fill and only test for errors in frames or tokens.

NOTE—The requirement for the station to transmit fill before performing the lobe test allows time for repeaters and the concentrator to synchronize with the transmitted data. A typical extended lobe might contain a copper to fiber repeater at the station and a fiber to copper repeater at the concentrator. Thus the path to the concentrator and back to the station would contain 5 receivers (not counting the station) each one requiring up to 1.5 ms to acquire synchronization. The 6 receivers plus the propagation delay consumes the entire 10 ms budget. Therefore it is recommended that the station transmit fill for a minimum of 20 ms.

Join State J5, Neighbor Notification (JS=NN) The join state JS=NN is set when the join ring FSM enters the neighbor notification state (J5). In this state the station waits until it has learned the address of its upstream neighbor.

Join State J6, Request Initialization (JS=RQI) The join state JS=RQI is set when the join ring FSM enters the request initialization state (J6). In this state the station transmits Request Initialization MAC frames to the RPS functional address. If the RPS is present, as determined by the settings of the A bits on these frames, the station expects to receive a valid Initialize Station or Change Parameters MAC frame in order to successfully complete the Join Ring process.

4.2.4.3.2 Monitor states

The monitor state (MS=) notation is used to identify the current state of the monitor FSM. The monitor state values are repeat, repeat beacon, repeat claim token, transmit beacon, transmit claim token, transmit ring purge, and beacon test. During the Beacon Test state, no assumptions can be made regarding the transmission or reception of data. Monitor states, listed by state value, are defined as follows:

Monitor State M4, Beacon Test (MS=BNT) The monitor state MS=BNT is set when the monitor FSM enters the beacon test state (M4). The station has removed itself from the ring and is performing a self test. During the Beacon Test state, the station assumes it is the only entity on the lobe capable of putting data on the lobe and no assumptions can be made regarding the station’s reception of data. Upon entering the lobe test state, the station waits for timer TRW to expire before starting its test sequence. Once the station has started the test sequence, the station shall transmit only valid frames, tokens, or fill and only test for errors in frames or tokens. At the start of the test sequence, the station shall transmit valid fill for a minimum of 10 ms. The station shall then perform a test to verify the operation of the lobe. Annex P contains the bit error rate requirements and recommendations for implementing lobe tests.

Monitor State M0, Repeat (MS=RPT) The monitor state MS=RPT is set when the monitor FSM enters the normal repeat state (M0). The ring is operational.

Monitor State M5, Repeat Beacon (MS=RBN) The monitor state MS=RBN is set when the monitor FSM enters the repeat beacon state (M5). The ring is in recovery, another station is transmitting beacon MAC frames.

Monitor State M6, Repeat Claim Token (MS=RCT) The monitor state MS=RCT is set when the monitor FSM enters the repeat claim token state (M6). The ring is in recovery, another station is transmitting claim token MAC frames.

Monitor State M3, Transmit Beacon (MS=TBN) The monitor state MS=TBN is set when the monitor FSM enters the transmit beacon state (M3). The ring is in recovery, this station is transmitting beacon MAC frames.

Monitor State M2, Transmit Claim Token (MS=TCT) The monitor state MS=TCT is set when the monitor FSM enters the transmit claim token state (M2). The ring is in recovery, and this station is transmitting claim token MAC frames (contending to be the active monitor).

Monitor State M1, Transmit Ring Purge (MS=TRP) The monitor state MS=TRP is set when the monitor FSM enters the Transmit Ring Purge state (M1). The ring is in recovery and this station is transmitting ring purge MAC frames to clean the ring.

4.2.4.3.3 Transmit states

The transmit state (TS=) notation is used to identify the current state of the transmit FSM. The transmit state values are repeat, transmit frame data, transmit fill, and transmit fill and strip. Transmit states, listed by state value, are defined as follows:

Transmit State T0, Repeat (TS=RPT) The transmit state TS=RPT is set when the transmit FSM enters the normal repeat state (T0) and is not transmitting (sourcing) frames.

Transmit State T1, Transmit Frame Data (TS=DATA) The transmit state TS=DATA is set when the transmit FSM is transmitting the data portion of the frame (T1).

Transmit State T2, Transmit Fill (TS=FILL) The transmit state TS=FILL is set when the transmit FSM has completed transmitting its frame(s), is stripping or waiting to strip its transmitted frame(s) from the ring and prepares to release a token (T2).

Transmit State T3, Transmit Fill and Strip (TS=STRIP) The transmit state TS=STRIP is set when the transmit FSM has released a token and is stripping the remainder of its frame(s) from the ring (T3).

4.3 Token ring MAC protocol specification

This clause specifies the procedures that are used in the MAC protocol.

The operation of the MAC is specified in this clause. In the case of a discrepancy between the FSM diagrams or table 7 and the supporting text, the FSM diagrams or table 7 take precedence. Table 7 takes precedence over the FSM diagrams.

The MAC receives from the PHY a serial stream of symbols (see 5.1.2.1 PS_UNITDATA.indication). Each symbol is one of the following:

a) Data_zero = zero symbol
b) Data_one = one symbol
c) Non-data_J = J symbol
d) Non-data_K = K symbol

(See 5.3 for a detailed description of these symbols.)

From the received symbols the MAC detects various types of input data, such as tokens, MAC frames, and LLC information frames.

In turn, the MAC stores values, sets flags, and performs certain internal actions as well as generating tokens, frames, fill, or modifying bits and delivering them to the PHY in the form of a serial stream of the Data_zero, Data_one, Non-data_J and Non-data_K symbols.

For the purpose of accumulating the FCS and storing the contents of a frame, Non-data_J and Non-data_K symbols that are not part of the SD or ED shall be interpreted as 1 and 0 values, respectively.

4.3.1 Token identification

Four varieties of token identification are used to describe station operation. They are Token, Good Token, Token Error, and Token with Error. These token varieties are indicated by combinations of the following properties:

Properties of a token

A—Starts with a valid SD.
B—Is three octets in length.
C—Is composed of only Data_zero and Data_one symbols between the SD and ED.
D—The T bit (token) is a Data_zero or optionally a Non-data_K.
E—Third octet is an ED and no ED_alignment_error (5.1.2.3).
F—No code violations in the P bits.

The four token varieties are defined below. This is not an inclusive list of all possible bit sequence formats; for example, other format sequences known in this standard are the frame and the abort sequence. Note that the value of the I, E, A, and C bits are not part of the following definitions. The notation -E means “not E.”

Good Token (TK_GOOD) defines the requirements for the entire token, while Token (TK) defines the requirements for a token that is captured to start transmission. In the latter case, only the first half of the token has been received which the transmit FSM changes into a start-of-frame sequence.

Good Token (TK_GOOD). A bit sequence that satisfies the following conditions based on the properties of a token listed above:

A & B & C & D & E

Token (TK). A bit sequence that satisfies the following condition, based on the properties of a token listed above:

A & D & F

Token Access Control (TK_AC). A bit sequence that satisfies the following condition, based on the properties of a token listed above:

A & D & optionally F

Token Error (TK_ERR). A bit sequence that satisfies the following condition, based on the properties of a token listed above:

A & D & -E

Token with Error (TK_WITH_ERR). A bit sequence that satisfies the following condition, based on the properties of a token listed above:

A & B & -C & D & E

4.3.2 Frame identification

Two definitions of frame validity are used to describe station operation. They are frame (FR) and frame with error (FR_WITH_ERR). The definition FR_AC is provided to define the valid start of a frame.

These frame validity definitions are indicated by combinations of the following properties:

Properties of a frame

A—Starts with a valid SD.
C—Is an integral number of octets in length and no ED_alignment_error (5.1.2.3).
D—Is composed of only Data_zero and Data_one symbols between the SD and ED.
E—Has the FF bits of the Frame Control field equal to 00 or 01.
F—Has a valid FCS.
G—Has a minimum of 18 octets between SD and ED.
H—Does not contain a valid SD or ED between the bounding SD and ED.
J—The end is delimited by a valid ED.
K—Has no code violations in P and T, and optionally M, and optionally R.
L—T bit is equal to a 1 indicating a frame.

The frame validity requirements are defined below. This is not an inclusive list of all possible bit sequence formats; for example, other format sequences known in this standard are the token and the abort sequence. Note that the value of the I, E, A, and C bits are not part of the following definitions.

Frame (FR). A valid frame is a bit sequence that satisfies the following condition, based on the properties of a frame listed previously:

	E & A & C & D & J & L & G & F	**(for MAC and LLC frames)**
or	**-E & A & C & D & J & L**	**(for undefined frame formats)**

Frame with Error (FR_WITH_ERR). A bit sequence that satisfies the following condition:

	E & A & H & J & L & (-C or -D or -F or -G)	**(for MAC and LLC frames)**
or	**-E & A & H & J & L & (-C or -D)**	**(for undefined frame formats)**

In addition to frame validity, the identification of the start of a frame is defined as:

Frame Access Control (FR_AC). A bit sequence that identifies the following condition:

A & K & L

Even though property K allows for variation, the station shall use a single definition for FR_AC for all instances of use.

4.3.3 Actions on frame errors

Frames received in error: Unless otherwise specified, a frame that is received with an error (FR_WITH_ERR) is not processed.

Transmission errors: When a station is transmitting, it shall inspect the received data to determine if its transmitted frame(s) was successfully propagated around the ring. Any frame error (as stated above) or the inability to recognize that the transmitted frame returned to the station is considered a transmission error. Frames that must be re-transmitted as a result of a transmission error are specified in tables 7.1, 7.3, and 7.4; otherwise, the MAC does not re-transmit the frame.

The capability to recognize that a transmitted frame has been lost (failed to return to the station) is provided by the TRR timer. When a station finishes transmitting its frame(s) it starts TRR and will normally wait in transmit state TS=FILL or TS=STRIP until it receives the SA of the last frame it transmitted. If TRR expires while the station is in TS=FILL or TS=STRIP, then the frame(s) is lost. The frame is also considered lost if, when the last frame returns, as determined by the frame counter (CFR), the SA does not match the stored SA of the last transmitted frame (LTA).

During the Ring Purge process, the TRR timer is also used to detect lost Ring Purge frames. If the frame returns to the station, then the station releases a token and transitions out of the MS=TRP state. If TRR expires while the station is in MS=TRP, then the Ring Purge frame was lost.

4.3.4 Station Operation tables

Tables 7.1 to 7.6 are used to precisely specify the operation of the token ring station. Each entry in the tables contains:

a) A state transition (S/T) indication that correlates to a transition on the appropriate FSM diagram in annex F.

b) A reference designator (REF) is assigned to the event. The sole purpose of the reference designator is to precisely identify the transition when necessary. The value of the reference designator is otherwise completely arbitrary and should not be interpreted to convey any other meaning.

c) Event/condition is a list of terms that can be equated to being true or false. If the event/condition equates true, then the specified action/output shall be performed. If the event/condition equates false, then the action/output shall be ignored. The highlighted term is the event that occurred or term that changed to cause the event/condition to be true.

d) Action/output is a list of terms that are equated to actions (internal to the MAC) or outputs (external to the MAC) that shall be performed as specified in 4.2.1 when the event/condition is true. If the event/condition equates false, then the action/output shall be ignored.

It is possible that a single occurrence of an event causes multiple event/condition terms to equate true. In this case, the actions/outputs of all table entries that equate true shall be performed.

The precise terms used for the events are specified in 4.3.5.1 using the definitions from previous clauses. The precise terms used for the actions/outputs are specified in 4.3.5.2 using the definitions from previous clauses.

Where possible, the first term of the event/condition is the dynamic event that occurred and the following terms are the static conditions. Tables 7.1 to 7.6 are ordered based on the alphanumeric value of the event. The position of the entry in the table has no significant meaning.

The starting state for the join state machine shall be JS=BP.

For the purpose of allowing flexibility among stations, parameters n1 through n5 are used in the Station Operation tables to represent certain counter thresholds. The minimum and maximum allowable threshold values are specified in table 6. The maximum values are the recommended values.

Table 6—Definition of Parameters n1 through n5

Parameters	Min	Max	Used with	Description
n1	1	3	CCR CCT	n1 is the number of the station’s own Claim Token MAC frames that must have been received before the station can become active monitor.
n2	1	2	CDG CDF	n2 is the number of DAT frames that must be received without address conflict before the station assumes its address is unique or the number of DAT frames received with conflict before the station assumes that it is a duplicate.
n3	1	4	CER CSC	n3 is the number of attempts to send an Error Report frame or a Report SUA Change frame when the destination server exists but fails to copy the frame.
n4	1	4	CRI CRIN	n4 is the number of Request Initialization frames sent to the RPS to determine that the RPS function is not present or if it is present, n4 is the number of Request Initialization frames sent waiting for the RPS to initialize the station.
n5	1	2	CNNR	n5 is the number of Neighbor Notification cycles that must be started after the ring is purged before the station reports AC errors.

Tables 7.1 to 7.6—MAC Station Operation tables

The following notation is used throughout tables 7.1, 7.2, 7.3, 7.4, 7.5, and 7.6:

(optional-i)—New implementations are encouraged to include this action.

(optional-x)—New implementations are encouraged to exclude this action.

(optional)—New implementations are allowed to include or exclude this action.

Table 7.1—MAC Join Station Operation table

S/T	REF	Event/condition	Action/output
J01	006	Connect.MAC & JS=BP	JS=LT; Set_initial_conditions; FTI=FTXC=x; TEST
J38	007	Disconnect.MAC & JS=ANM	JS=BPW
J48	008	Disconnect.MAC & JS=DAT	JS=BPW
J28	009	Disconnect.MAC & JS=DMP	JS=BPW
J78	010	Disconnect.MAC & JS=JC	JS=BPW; FINS=FJR=FDC=FNC=0
J58	011	Disconnect.MAC & JS=NN	JS=BPW; FDC=0
J68A	012	Disconnect.MAC & JS=RQI	JS=BPW; FINS=FDC=FNC=0
J38	255	FBPF=1 & TS=RPT & JS=ANM	JS=BPW
J48	256	FBPF=1 & TS=RPT & JS=DAT	JS=BPW
J78	257	FBPF=1 & TS=RPT & JS=JC	JS=BPW; FINS=FJR=FDC=FNC=0
J58	258	FBPF=1 & TS=RPT & JS=NN	JS=BPW; FDC=0
J68A	259	FBPF=1 & TS=RPT & JS=RQI	JS=BPW; FINS=FDC=FNC=0
	359	FDC=1 & FNC=0 & MS=RPT & FAM=1	[FNC=1; (optional-i)]
	027	FDC=1 & FNW=1 & FNC=0 & MS=RPT & FAM=0	FNC=1; QUE_SMP_PDU
	028	FDC=1 & FNW=1 & FNC=0 & MS=RPT & FAM=1	FNC=1
J24	058	FR_AMP & JS=DMP	JS=DAT; FMP=1; TJR=R; CDF=0; CDG=0; QUE_DAT_PDU
J38	067	FR_BN & JS=ANM	JS=BPW
J48	068	FR_BN & JS=DAT	JS=BPW
J28	069	FR_BN & JS=DMP	JS=BPW
J58	070	FR_BN & JS=NN	JS=BPW; FDC=0
J68B	071	FR_BN & JS=RQI & MS<>TBN & FINS=0	JS=BPW; FDC=FNC=0
J68B	081	FR_BN(SA<>MA & BN_TYPE<=TX_BN_TYPE) & MS=TBN & FINS=0 & JS=RQI	JS=BPW; FDC=FNC=0
J67	100	FR_CHG_PARM(DA=MA) & JS=RQI	JS=JC; FJR=1; If FWFDO=0 then (FWFA=FWF=0; TWFD=R)
J23B	106	FR_CT & JS=DMP	JS=ANM
	124	FR_DAT(SA=MA & A=0 & C=0) & CDG<(n2-1) & JS=DAT	CDG=(CDG+1); QUE_DAT_PDU
J45	125	FR_DAT(SA=MA & A=0 & C=0) & CDG=(n2-1) & JS=DAT	JS=NN; FDC=1; TJR=R
	126	FR_DAT(SA=MA & A|C<>0) & CDF<(n2-1) & JS=DAT	CDF=(CDF+1); QUE_DAT_PDU
J48	127	FR_DAT(SA=MA & A|C<>0) & CDF=(n2-1) & JS=DAT	JS=BPW
J67	131	FR_INIT(DA=MA) & JS=RQI	JS=JC; FJR=1; If FWFDO=0 then (FWFA=FWF=0; TWFD=R)
J38	156	FR_REMOVE(DA=Non_broadcast) & FRRO=0 & JS=ANM	JS=BPW
J48	157	FR_REMOVE(DA=Non_broadcast) & FRRO=0 & JS=DAT	JS=BPW
J28	158	FR_REMOVE(DA=Non_broadcast) & FRRO=0 & JS=DMP	JS=BPW

Table 7.1—MAC Join Station Operation table (*Continued*)

S/T	REF	Event/condition	Action/output
J78	159	FR_REMOVE(DA=Non_broadcast) & FRRO=0 & JS=JC	JS=BPW; FINS=FJR=FDC=FNC=0
J58	160	FR_REMOVE(DA=Non_broadcast) & FRRO=0 & JS=NN	JS=BPW; FDC=0
J68A	161	FR_REMOVE(DA=Non_broadcast) & FRRO=0 & JS=RQI	JS=BPW; FINS=FDC=FNC=0
J34	164	FR_RP & JS=ANM	JS=DAT; FMP=1; TJR=R; CDF=0; CDG=0; QUE_DAT_PDU
J24	165	FR_RP & JS=DMP	JS=DAT; FMP=1; TJR=R; CDF=0; CDG=0; QUE_DAT_PDU
	187	FR_RQ_INIT(SA=MA & A<>1) & CRIN<n4 & JS=RQI	TRI=R; CRIN=(CRIN+1); QUE_RQ_INIT
J67	188	FR_RQ_INIT(SA=MA & A<>1) & CRIN=n4 & JS=RQI	JS=JC; FJR=1; If FWFDO=0 then (FWFA=FWF=0; TWFD=R)
	189	FR_RQ_INIT(SA=MA & A=1) & CRI<n4 & JS=RQI	TRI=R
J24	191	FR_SMP & JS=DMP	JS=DAT; FMP=1; TJR=R; CDF=0; CDG=0; QUE_DAT_PDU
	206	FRH=0 & JS=BPW	MS=RPT; FAM=FOP=0; FRH=1; TRW=R; Remove_station; FA(monitor)=0
J48	210	FSL=1 & FAM=0 & JS=DAT	JS=BPW
J58	211	FSL=1 & FAM=0 & JS=NN	JS=BPW; FDC=0
J68B	209	FSL=1 & FINS=0 & FAM=0 & JS=RQI	JS=BPW; FDC=FNC=0
J38	224	INTERNAL_ERR(not_correctable) & JS=ANM	JS=BPW
J48	225	INTERNAL_ERR(not_correctable) & JS=DAT	JS=BPW
J28	226	INTERNAL_ERR(not_correctable) & JS=DMP	JS=BPW
J78	227	INTERNAL_ERR(not_correctable) & JS=JC	JS=BPW; FINS=FJR=FDC=FNC=0
J58	228	INTERNAL_ERR(not_correctable) & JS=NN	JS=BPW; FDC=0
J68A	229	INTERNAL_ERR(not_correctable) & JS=RQI	JS=BPW; FINS=FDC=FNC=0
	230	JS=BPW & FAM=1	[FAM=0; FA(monitor)=0 (optional)]
	231	JS=BPW & FOP=1	[FOP=0 (optional)]
	232	JS=BPW & MS<>RPT	[MS=RPT (optional)]
J28	238	PM_STATUS.indication(Medium_rate_error) & JS=DMP	JS=BPW
J68B	263	TBT=E & MS=TBN & FINS=0 & JS=RQI	JS=BPW; FDC=FNC=0
J38	265	TCT=E & JS=ANM	JS=BPW
J48	267	TCT=E & MS=RCT & JS=DAT	JS=BPW
J58	268	TCT=E & MS=RCT & JS=NN	JS=BPW; FDC=0
J48	269	TCT=E & MS=TCT & JS=DAT	JS=BPW
J58	270	TCT=E & MS=TCT & JS=NN	JS=BPW; FDC=0
J10	276	TEST_FAILURE & JS=LT	JS=BP
J70	277	TEST_FAILURE & MS=BNT & JS=JC	JS=BP
J60	278	TEST_FAILURE & MS=BNT & JS=RQI	JS=BP
J12	281	TEST_OK & JS=LT	JS=DMP; MS=RPT; TS=RPT; FTI=FTXC=0; FOP=FRH=1; TJR=R; TRH=R; SUA=0; INSERT If FWFDO=1 then FWFA=FWF=0; TWFD=R

Table 7.1—MAC JoinStation Operation table (*Continued*)

S/T	REF	Event/condition	Action/output
J48	284	TJR=E & JS=DAT	JS=BPW
J23A	287	TJR=E & JS=DMP	JS=ANM; FANM=1
J58	285	TJR=E & JS=NN	JS=BPW; FDC=0
J68A	286	TJR=E & JS=RQI	JS=BPW; FINS=FDC=FNC=0
	315	TRH=E	FRH=0
	316	TRI=E & JS=RQI & CRI<n4	TRI=R; CRI=(CRI+1); QUE_RQ_INIT_PDU
J68A	317	TRI=E & JS=RQI & CRI=n4	JS=BPW; FINS=FDC=FNC=0
J68B	321	TRP=E & MS=TRP & FINS=0 & JS=RQI	JS=BPW; FDC=FNC=0
J38	318	TRP=E & MS=TRP & JS=ANM	JS=BPW
J48	319	TRP=E & MS=TRP & JS=DAT	JS=BPW
J58	320	TRP=E & MS=TRP & JS=NN	JS=BPW; FDC=0
J80	325	TRW=E & JS=BPW	JS=BP; FTI=x
J38	368	TWF=E & FWFA=1 & FWF=1 & MS<>BNT & JS=ANM	JS=BPW
J48	369	TWF=E & FWFA=1 & FWF=1 & MS<>BNT & JS=DAT	JS=BPW
J28	370	TWF=E & FWFA=1 & FWF=1 & MS<>BNT & JS=DMP	JS=BPW
J78	345	TWF=E & FWFA=1 & FWF=1 & MS<>BNT & JS=JC	JS=BPW; FINS=FJR=FDC=FNC=0
J58	371	TWF=E & FWFA=1 & FWF=1 & MS<>BNT & JS=NN	JS-BPW; FDC=FNC=0
J68A	372	TWF=E & FWFA=1 & FWF=1 & MS<>BNT & JS=RQI	JS=BPW; FINS=FDC=FNC=0
	347	TX_ERR(DAT) & JS=DAT	QUE_DAT_PDU
	350	TX_ERR(RQ_INIT) & JS=RQI & FOP=1	QUE_RQ_INIT_PDU

Table 7.2—MAC Transmit Station Operation table

S/T	REF	Event/condition	Action/output
	013	EOD & TS=DATA & FED=0 & FTXI=0	FED=1; [FSD=1 (optional-i)]
T12A	014	EOD & TS=DATA & FTXI=0 & QUE_EMPTY	TS=FILL; FTI=1; TRR=R; TX_FCS; TX_EFS(I=0)
T12A	015	EOD & TS=DATA & FTXI=0 & QUE_NOT_EMPTY & (CTO-FR_LTH<0)	TS=FILL; FTI=1; TRR=R; TX_FCS; TX_EFS(I=0)
T12A	016	EOD & TS=DATA & FTXI=0 & QUE_NOT_EMPTY & (CTO-FR_LTH>=0) & Pm<Pr	TS=FILL; FTI=1; TRR=R; TX_FCS; TX_EFS(I=0)
T12A	017	EOD & TS=DATA & FTXI=0 & QUE_NOT_EMPTY & (CTO-FR_LTH>=0) & Pm>=Pr & FMFTO=0 & Pm<Rr	TS=FILL; FTI=1; TRR=R; TX_FCS; TX_EFS(I=0)
	018	EOD & TS=DATA & FTXI=0 & QUE_NOT_EMPTY & (CTO-FR_LTH>=0) & Pm>=Pr & FMFTO=0 & Pm>=Rr	CFR=(CFR+1); CTO=(CTO-FR_LTH); TX_FCS; TX_EFS(I=x); TX_SFS(P=Pr;R=0); LTA=TX_SA
	019	EOD & TS=DATA & FTXI=0 & QUE_NOT_EMPTY & (CTO-FR_LTH>=0) & Pm>=Pr & FMFTO=1	CFR=(CFR+1); CTO=(CTO-FR_LTH); TX_FCS; TX_EFS(I=x); TX_SFS(P=Pr;R=0); LTA=TX_SA
T10A	020	EOD & TS=DATA & FTXI=1 & MS=RBN	TS=RPT; FTXI=FTXC=0; TX_FCS; TX_EFS(I=0)
T10A	021	EOD & TS=DATA & FTXI=1 & MS=RCT	TS=RPT; FTXI=FTXC=0; TX_FCS; TX_EFS(I=0)
T10A	022	EOD & TS=DATA & FTXI=1 & MS=RPT	TS=RPT; FTXI=FTXC=0; TX_FCS; TX_EFS(I=0)
T10B	023	EOD & TS=DATA & FTXI=1 & MS=TBN	TS=RPT; FTXI=0; FTI=1; TX_FCS; TX_EFS(I=0)
T10B	024	EOD & TS=DATA & FTXI=1 & MS=TCT	TS=RPT; FTXI=0; FTI=1; TX_FCS; TX_EFS(I=0)
T10C	025	EOD & TS=DATA & FTXI=1 & MS=TRP	TS=RPT; FTXI=0; FTI=1; TX_FCS; TX_EFS(I=0); TRR=R
T30B	036	FNMA=1 & CFR=1 & TS=STRIP	TS=RPT; FTI=0; FLF=1
	040	FR(DA=Any_recognized_address) & FTI=0 & TS=RPT	SET A=1
	043	FR(P<Sx)	[CLEAR_STACKS (optional-i)]
	054	FR_AC(R<Pm) & PDU_QUEUED & FTI=0 & TS=RPT & FOP=1	SET R=Pm
	103	FR_COPIED(DA=Any_recognized_address) & FTI=0 & TS=RPT	SET C=1
	203	FR_WITH_ERR & FTI=0 & TS=RPT	SET E=1
	044	RCV(SA<>LTA) & CFR=1	FNMA=1
	045	RCV(SA=LTA) & CFR=1	FMA=1
	246	RCV_AC	STORE(Pr;Rr)
	247	RCV_ED	FMA=FNMA=FSD=0
	250	RCV_ED & FSD=1 & FED=1 & CFR>1	CFR=(CFR-1)
T30B	249	RCV_ED & FSD=1 & FMA=0 & CFR=1 & TS=STRIP	TS=RPT; FTI=0; FLF=1
T30A	248	RCV_ED & FSD=1 & FMA=1 & CFR=1 & TS=STRIP	TS=RPT; FTI=0
	251	RCV_SD	FMA=FNMA=0; FSD=1
T10D	254	STATION_ERR(correctable) & TS=DATA	TS=RPT; TX_AB
T10E	363	STATION_ERR(not-correctable) & TS=DATA	TS=RPT; [TX_AB (optional)]; FBPF=1
T10D	364	STATION_ERR(tx_underrun) & FTUBO=0 & TS=DATA	TS=RPT; TX_AB
T12B	365	STATION_ERR(tx_underrun) & FTUBO=1 & TS=DATA	TS=FILL; FED=FTI=1; TRR=R; TX_AB
T01B	293	TK(P<=Pm) & PDU_QUEUED & FTI=0 & FOP=1 & TS=RPT	TS=DATA; FED=FMA=FSD=FTXI=0; CFR=1; CTO=(MAX_TX-FR_LTH); TX_SFS(P=Pr;R=0); LTA=TX_SA
	294	TK(P<>Sx) & PDU_QUEUED(P>Pm>R) & FTI=0 & TS=RPT & FOP=1	SET R=Pm
	300	TK(P>0 & M=0) & FAM=1 & FTI=0 & TS=RPT	SET M=1
	295	TK_AC(P<Sx)	CLEAR_STACKS
	296	TK_AC(P=Sx) & Sr<Px & PDU_QUEUED(Pm<Sx) & FTI=0 & FOP=1 & TS=RPT	TX_TK(P=Px;M=0;R=0); RESTACK(Sx=Px)

Table 7.2—MAC Transmit Station Operation table (*Continued*)

S/T	REF	Event/condition	Action/output
	297	TK_AC(P=Sx) & Sr<Px & QUE_EMPTY & FTI=0 & FOP=1 & TS=RPT	TX_TK(P=Px;M=0;R=0); RESTACK(Sx=Px)
	298	TK_AC(P=Sx) & Sr>=Px & PDU_QUEUED(Pm<Sx) & FTI=0 & FOP=1 & TS=RPT	TX_TK(P=Sr;M=0;R=Px); POP(Sx;Sr)
	299	TK_AC(P=Sx) & Sr>=Px & QUE_EMPTY & FTI=0 & FOP=1 & TS=RPT	TX_TK(P=Sr;M=0;R=Px); POP(Sx;Sr)
T10D	301	TK_ERR & FTXI=0 & TS=DATA	TS=RPT; TX_AB
	307	TK_WITH_ERR & FTEO=0 & FTI=0 & TS=RPT	SET E=1
T20	323	TRR=E & TS=FILL	TS=RPT; FTI=0; FLF=1
T30B	324	TRR=E & TS=STRIP	TS=RPT; FTI=0; FLF=1
T23A	327	TS=FILL & FMA=0 & FETO=1 & FMRO=1& Pr=Sx & Sr<Px	TS=STRIP; TX_TK(P=Px;M=0;R=0); RESTACK(Sx=Px)
T23B	328	TS=FILL & FMA=0 & FETO=1 & FMRO=1& Pr=Sx & Sr>=Px	TS=STRIP; TX_TK(P=Sr;M=0;R=Px); POP(Sx;Sr)
T23C	329	TS=FILL & FMA=0 & FETO=1 & FMRO=1& Pr>Sx & Pr<Px	TS=STRIP; TX_TK(P=Px;M=0;R=0); STACK(Sx=Px;Sr=Pr)
T23D	330	TS=FILL & FMA=0 & FETO=1 & FMRO=1& Pr>Sx & Pr>=Px	TS=STRIP; TX_TK(P=Pr;M=0;R=Px)
T23A	331	TS=FILL & FMA=1 & Pr=Sx & Sr<Px	TS=STRIP; TX_TK(P=Px;M=0;R=0); RESTACK(Sx=Px)
T23B	332	TS=FILL & FMA=1 & Pr=Sx & Sr>=Px	TS=STRIP; TX_TK(P=Sr;M=0;R=Px); POP(Sx;Sr)
T23C	333	TS=FILL & FMA=1 & Pr>Sx & Pr<Px	TS=STRIP; TX_TK(P=Px;M=0;R=0); STACK(Sx=Px; Sr=Pr)
T23D	334	TS=FILL & FMA=1 & Pr>Sx & Pr>=Px	TS=STRIP; TX_TK(P=Pr;M=0;R=Px)
T01A	352	TXI_REQ & TS=RPT	TS=DATA; FTI=0; FTXI=1; TX_SFS(P=0;R=0)

Table 7.3—MAC Monitor Station Operation table

S/T	REF	Event/condition	Action/output
M54	003	CBR=0 & MS=RBN & FBR=0	MS=BNT; FBR=FTW=1
M02A	026	FANM=1 & MS=RPT & JS=ANM	MS=TCT; FANM=FCT=0; FTXC=FTI=1; TCT=R; TQP=R; CCT=1; CCR=1; TXI(CT _PDU)
M52A	038	FR(DA=Any_recognized_address) & FID=1 & MS=RBN	MS=TCT; FCT=FID=0; FTXC=FTI=1; TCT=R; TQP=R; CCT=1; CCR=1; TXI(CT_PDU)
	039	FR(DA=Any_recognized_address) & FID=1 & MS=TBN	FID=0; FTI=1; TQP=R; TXI(BN_PDU)
	046	FR_AC & FAM=0 & FPT=1 & FGTO=0	TNT=R; FPT=0
	048	FR_AC(M=0) & FAM=1 & FTHO=0	TVX=R
	049	FR_AC(M=0) & FAM=1 & FTHO=1	FAT=1
	047	FR_AC(M=0) & FAM=1 & FTI=0 & TS=RPT	SET M=1
	050	FR_AC(M=0) & MS=RBN & CBC=0 & FTI=0 & TS=RPT	[SET M=1 (optional-i)]
M01	051	FR_AC(M=1) & FAM=1 & MS=RPT	MS=TRP; FTI=1; TRP=R; TXI(RP_PDU)
	055	FR_AMP & CNNR<>0	CNNR=(CNNR-1)
	057	FR_AMP & FAM=0 & FINS=1	TSM=R
	056	FR_AMP & FAM=1	[TSM=R (optional-i)]
	059	FR_AMP(A=0 & C=0) & FAM=0 & FOP=1	FSMP=1; TQP=R
	060	FR_AMP(A|C<>0) & FAM=0	FSMP=0
	061	FR_AMP(SA<>MA) & FAM=1	QUE_ACT_ERR_PDU(EC=2); FAM=0; FA(monitor)=0
	062	FR_AMP(SA<>MA) & FAM=1 & FINS=1	TNT=R
	063	FR_AMP(SA<>MA) & FAM=1 & TS=RPT	FTXC=FTI=0
	064	FR_AMP(SA<>SUA & A=0 & C=0)& FOP=1	CSC=1; SUA=SA ; QUE_SUA_CHG_PDU
	065	FR_AMP(SA=MA & A=0 & C=0) & FAM=1	FNN=FNW=1; LMP=NULL
	066	FR_AMP(SA=MA & A|C<>0) & FAM=1	LMP=SA
	072	FR_BN & MS<>TBN & MS<>BNT & FINS=1	TBR=R
M65	073	FR_BN & MS=RCT & FINS=1	MS=RBN
	074	FR_BN & MS=RPT & FAM=1 & TS=RPT	FTXC=0
M05	075	FR_BN & MS=RPT & FINS=1	MS=RBN; FAM=0; FA(monitor)=0
M25	076	FR_BN & MS=TCT & FINS=1	MS=RBN
	077	FR_BN & MS=TCT & FINS=1 & TS=RPT	FTXC=FTI=0
M15	078	FR_BN & MS=TRP & FINS=1	MS=RBN; FAM=0; FA(monitor)=0
	079	FR_BN & MS=TRP & FINS=1 & TS=RPT	FTXC=FTI=0
M32	080	FR_BN(M=0 & SA=MA) & MS=TBN & FID=0	MS=TCT; FCT=0; TCT=R; TQP=R; CCT=1; CCR=1; TXI(CT_PDU)
M35	082	FR_BN(SA<>MA & BN_TYPE<=TX_BN_TYPE) & MS=TBN & FID=0 & FINS=1	MS=RBN; TBR=R
	083	FR_BN(SA<>MA & BN_TYPE<=TX_BN_TYPE) & MS=TBN & FID=0 & FINS=1 & TS=RPT	FTXC=FTI=0
	094	FR_BN(SA<>MA & UNA=MA & BN_TYPE<>1) & FBR=0 & MS<>RBN & MS<>BNT	CBR=7
	084	FR_BN(SA<>SUA)	CBC=2
	087	FR_BN(SA=SUA & UNA<>MA) & CBC=0 & MS=RBN	CBC=1

Table 7.3—MAC Monitor Station Operation table (*Continued*)

S/T	REF	Event/condition	Action/output
	090	FR_BN(SA=SUA & UNA<>MA) & FBR=0 & CBC>0 & MS=RBN	CBC=(CBC-1)
	086	FR_BN(SA=SUA & UNA<>MA) & MS<>RBN & FBR=0 & MS<>BNT	CBC=1
	085	FR_BN(SA=SUA & UNA=MA) & FBR=0	CBC=2
	089	FR_BN(SA=SUA) & FBR=1 & CBC=0 & MS=RBN	CBC=1
	091	FR_BN(SA=SUA) & FBR=1 & CBC>0 & MS=RBN	CBC=(CBC-1)
	088	FR_BN(SA=SUA) & FBR=1 & MS<>RBN & MS<>BNT	CBC=1
	092	FR_BN(UNA<>MA)	CBR=8
	095	FR_BN(UNA=MA & BN_TYPE<>1) & FBR=0 & MS=RBN	CBR=(CBR-1)
	093	FR_BN(UNA=MA & BN_TYPE=1)	CBR=8
M52C	096	FR_BN_CIR & MS=RBN & CBC<>2 & FINS=1	[MS=TCT; FCT=0; FTXC=FTI=1; TCT=R; TQP=R; CCT=1; CCR=1; TXI(CT_PDU) (optional-i)]
	104	FR_CT & FAM=1	FAM=0; QUE_ACT_ERR_PDU(EC=1); FA(monitor)=0
	105	FR_CT & FAM=1 & TS=RPT	FTXC=FTI=0
	107	FR_CT & MS<>TBN & FCT=0	FCT=1
M16	108	FR_CT & MS=TRP	MS=RCT; TCT=R
M26A	109	FR_CT(M=0 & SA=MA & UNA<>SUA) & MS=TCT	MS=RCT; TCT=R; QUE_ACT_ERR_PDU(EC=3)
	110	FR_CT(M=0 & SA=MA & UNA<>SUA) & MS=TCT & TS=RPT	FTXC=FTI=0
M21	111	FR_CT(M=0 & SA=MA & UNA=SUA) & CCT>=n1 & CCR>=n1 & MS=TCT	MS=TRP; FAM=1; TRP=R; TXI(RP_PDU); QUE_NEW_MON_PDU; FA(monitor)=1
	112	FR_CT(M=0 & SA=MA & UNA=SUA) & MS=TCT	CCR=(CCR+1)
M56	113	FR_CT(SA<MA) & MS=RBN & FCCO=0 & FID=0	MS=RCT; TCT=R
M52B	114	FR_CT(SA<MA) & MS=RBN & FCCO=1 & FINS=1	MS=TCT; FCT=0; FTXC=FTI=1; TCT=R; TQP=R; CCT=1; CCR=1; TXI(CT_PDU)
M06	115	FR_CT(SA<MA) & MS=RPT & FAM=1	MS=RCT; TCT=R
M06	116	FR_CT(SA<MA) & MS=RPT & FCCO=0 & FOP=1	MS=RCT; TCT=R
M02B	117	FR_CT(SA<MA) & MS=RPT & FCCO=1 & FAM=0 & FINS=1	MS=TCT; FCT=0; FTXC=FTI=1; TCT=R; TQP=R; CCT=1; CCR=1; TXI(CT_PDU)
M06	118	FR_CT(SA<MA) & MS=RPT & FINS=0 & FOP=1	MS=RCT; TCT=R
M56	119	FR_CT(SA>MA) & MS=RBN & FID=0	MS=RCT; TCT=R
M06	120	FR_CT(SA>MA) & MS=RPT & FOP=1	MS=RCT; TCT=R
M26B	121	FR_CT(SA>MA) & MS=TCT	MS=RCT; TCT=R
	122	FR_CT(SA>MA) & MS=TCT	[TXI(RCV_CT_PDU) (optional-i)]
	123	FR_CT(SA>MA) & MS=TCT & TS=RPT	FTXC=FTI=0
	163	FR_RP	CNNR=n5
M60	169	FR_RP & MS=RCT	MS=RPT
	166	FR_RP & MS=RCT & FINS=1	FBR=FBT=0
	170	FR_RP & MS=RCT & FINS=1	TNT=R; TSM=R

Table 7.3—MAC Monitor Station Operation table (*Continued*)

S/T	REF	Event/condition	Action/output
	171	FR_RP & MS=RPT & FAM=1	FAM=0; QUE_ACT_ERR_PDU(EC=2); FA(monitor)=0
	173	FR_RP & MS=RPT & FAM=1 & FINS=1	TNT=R; [TSM=R (optional-i)]
	172	FR_RP & MS=RPT & FAM=1 & TS=RPT	FTXC=FTI=0
	167	FR_RP & MS=RPT & FINS=1	FBR=FBT=0
	168	FR_RP & MS=TRP & FINS=1	FBR=FBT=0
	174	FR_RP(SA<>MA) & FAM=0 & FINS=1	[TNT=R (optional-x)]
M10B	175	FR_RP(SA<>MA) & MS=TRP	MS=RPT; FAM=0; QUE_ACT_ERR_PDU(EC=2); FA(monitor)=0
	176	FR_RP(SA<>MA) & MS=TRP & TS=RPT	FTXC=FTI=0
M10A	177	FR_RP(SA=MA & R=0) & MS=TRP	MS=RPT; FAT=FTI=FNN=0; FMP=1; TVX=R; TAM=R; TX_TK(P=0; M=0;R=0); CLEAR_STACKS; QUE_AMP_PDU; LMP=(MA|NULL); [TSM=R (optional-i)]
M10C	178	FR_RP(SA=MA & R>0) & MS=TRP	MS=RPT; FAT=FTI=FNN=0; FMP=1; TVX=R; TAM=R; TX_TK(P=Rr; M=0;R=0); RESET_STACKS(Sx=Rr;Sr=0); QUE_AMP_PDU; LMP=(MA|NULL); [TSM=R (optional-i)]
	192	FR_SMP(A=0 & C=0) & FAM=1	FNN=FNW=1; LMP=NULL
	193	FR_SMP(A=0 & C=0) & FSMP=0 & FAM=0 & FOP=1	FSMP=1; TQP=R
	198	FR_SMP(A|C<>0) & FAM=1	LMP=SA
	199	FR_SMP(SA<>SUA & A=0 & C=0) & FAM=0 & FOP=1	CSC=1; SUA=SA; QUE_SUA_CHG_PDU
	200	FR_SMP(SA<>SUA & A=0 & C=0) & FAM=1 & FNN=0 & FOP=1	CSC=1; SUA=SA; QUE_SUA_CHG_PDU
	201	FR_SUA_CHG(SA=MA & E=0 & A=1 & C=0) & CSC<n3 & FOP=1	CSC=(CSC+1); QUE_SUA_CHG_PDU
	202	FR_SUA_CHG(SA=MA & E=1 & C=0) & CSC<n3 & FOP=1	CSC=(CSC+1); QUE_SUA_CHG_PDU
	207	FSL=0 & MS=TBN & TX_BN_TYPE=2	TX_BN_TYPE=3
	208	FSL=1 & FID=1 & MS=TBN	FID=0; FTI=1; TQP=R; TX_BN_TYPE=2; TXI(BN_PDU)
M52A	212	FSL=1 & FINS=1 & MS=RBN	MS=TCT; FCT=FID=0; FTXC=FTI=1; TCT=R; TQP=R; CCT=1; CCR=1; TXI(CT_PDU)
M62	213	FSL=1 & FINS=1 & MS=RCT	MS=TCT; FCT=0; FTXC=FTI=1; TCT=R; TQP=R; CCT=1; CCR=1; TXI(CT_PDU)
M02C	214	FSL=1 & FINS=1 & MS=RPT	MS=TCT; FAM=FCT=0; FTXC=FTI=1; TCT=R; TQP=R; CCT=1; CCR=1; TXI(CT_PDU); FA(monitor)=0
M12A	215	FSL=1 & FINS=1 & MS=TRP	MS=TCT; FAM=FCT=0; TCT=R; TQP=R; CCT=1; CCR=1; TXI(CT_PDU); FA(monitor)=0
	216	FSL=1 & MS=TBN	TX_BN_TYPE=2
	219	FTW=1 & FRH=0 & MS=BNT	FOP=FTW=0; TRW=R; Remove_station
	260	TAM=E & FNN=0 & FAM=1 & MS=RPT	TAM=R; QUE_NN_INCMP_PDU; QUE_AMP_PDU; LMP=(MA|NULL)
	261	TAM=E & FNN=1 & FAM=1 & MS=RPT	FNN=0; TAM=R; QUE_AMP_PDU; LMP=(MA|NULL)
M52B	262	TBR=E & MS=RBN & FINS=1	MS=TCT; FCT=0; FTXC=FTI=1; TCT=R; TQP=R; CCT=1; CCR=1; TXI(CT_PDU)
M34	264	TBT=E & MS=TBN & FINS=1 & FBT=0	MS=BNT; FOP=0; FBT=1; TRW=R; Remove_station

Table 7.3—MAC Monitor Station Operation table (*Continued*)

S/T	REF	Event/condition	Action/output
M63	266	TCT=E & MS=RCT & FNC=1	MS=TBN; FTXC=FTI=1; TBT=R; TQP=R; TX_BN_TYPE=4; TXI(BN_PDU)
M23B	271	TCT=E & MS=TCT & FNC=1 & FCT=0 & FSL=0	MS=TBN; TBT=R; TQP=R; TX_BN_TYPE=3; TXI(BN_PDU)
M23A	273	TCT=E & MS=TCT & FNC=1 & FCT=1 & FSL=0	MS=TBN; TBT=R; TQP=R; TX_BN_TYPE=4; TXI(BN_PDU)
M23C	272	TCT=E & MS=TCT & FNC=1 & FSL=1	MS=TBN; TBT=R; TQP=R; TX_BN_TYPE=2; TXI(BN_PDU)
M45	279	TEST_OK & FBR=1 & MS=BNT	MS=RBN; FTI=FTXC=FWFA=FWF=0; FBT=FID=FOP=FRH=1; TWFD=R; TRH=R; TID=R; INSERT
M43	280	TEST_OK & FBT=1 & MS=BNT	MS=TBN; FTI=FWFA=FWF=0; FBR=FID=FOP=FRH=FTXC=1; TID=R; TRH=R; TWFD=R; INSERT
M52A	282	TID=E & MS=RBN & FID=1	MS=TCT; FCT=FID=0; FTXC=FTI=1; TCT=R; TQP=R; CCT=1; CCR=1; TXI(CT_PDU)
	283	TID=E & MS=TBN & FID=1	FID=0; FTI=1; TQP=R; TXI(BN_PDU)
M01	290	TK(M=1) & FAM=1 & MS=RPT	MS=TRP; FTI=1; TRP=R; TXI(RP_PDU)
	304	TK_GOOD & FAM=0 & FGTO=1	TNT=R
	305	TK_GOOD(P=0) & FAM=0 & FGTO=0	TNT=R; FPT=0
	306	TK_GOOD(P>0) & FAM=0 & FGTO=0	FPT=1
M02B	310	TNT=E & FAM=0 & FINS=1 & MS=RPT	MS=TCT; FCT=0; FTXC=FTI=1; TCT=R; TQP=R; CCT=1; CCR=1; TXI(CT_PDU)
	311	TQP=E & MS=RPT & FDC=0 & FAM=0	FNW=1
	312	TQP=E & MS=RPT & FDC=1 & FAM=0	FNC=1; QUE_SMP_PDU
	313	TQP=E & MS=TBN & FID=0	TQP=R; TXI(BN_PDU)
	314	TQP=E & MS=TCT	TQP=R; CCT=(CCT+1); TXI(CT_PDU)
M12B	322	TRP=E & MS=TRP & FINS=1	MS=TCT; FAM=FCT=0; TCT=R; TQP=R; CCT=1; CCR=1; TXI(CT_PDU); QUE_ACT_ERR_PDU(EC=1); FA(monitor)=0
	326	TRW=E & MS=BNT	FTXC=1; FTI=x; TEST
	335	TSL=E & FSLD=1	FSL=1
M02C	336	TSM=E & MS=RPT & FINS=0 & FAM=1	MS=TCT; FAM=FCT=0; FTXC=FTI=1; TCT=R; TQP=R; CCT=1; CCR=1; TXI(CT_PDU); FA(monitor)=0
M02C	337	TSM=E & MS=RPT & FINS=1	MS=TCT; FAM=FCT=0; FTXC=FTI=1; TCT=R; TQP=R; CCT=1; CCR=1; TXI(CT_PDU); FA(monitor)=0
M01	338	TVX=E & FTHO=0 & FAM=1 & MS=RPT	MS=TRP; FTI=1; TRP=R; TXI(RP_PDU)
M01	341	TVX=E & FTHO=1 & FAT=0 & FAM=1 & MS=RPT	MS=TRP; FTI=1; TRP=R; TXI(RP_PDU)
	344	TVX=E & FTHO=1 & FAT=1 & FAM=1	FAT=0; TVX=R
	346	TWFD=E & FWFA=0	FWFA=1
	348	TX_ERR(RP) & MS=TRP	TXI(RP_PDU)
	351	TX_ERR(SUA_CHG) & CSC<n3 & FOP=1	QUE_SUA_CHG_PDU; [CSC=CSC+1 (optional)]

Table 7.4—MAC Error Handling Station Operation table

S/T	REF	Event/condition	Action/output
	001	Burst5_error_event & MS=RPT & FINS=1 & FER=0	TER=R; FER=1; CBE=(CBE+1)
	002	Burst5_error_event & MS=RPT & FINS=1 & FER=1 & CBE<255	CBE=(CBE+1)
	004	CER=n3 & FECO=0	FER=FLF=0; CER=0; SET ERR_CNTR to 0;
	005	CER=n3 & FECO=1	FLF=0; CER=0; SET ERR_CNTR to 0;
	030	FJR=1 & FBHO=1 & FINS=0 & FECO=1	TER=R; FER=1
	031	FLF=1 & MS=RPT & FINS=1 & FER=0	FLF=0; FER=1; TER=R; CLFE=(CLFE+1)
	032	FLF=1 & MS=RPT & FINS=1 & FER=1 & CLFE<255	FLF=0; CLFE=(CLFE+1)
	034	FNC=1 & FBHO=0 & FINS=0 & FECO=1	FER=1; TER=R
	052	FR_AC(M=1) & FAM=1 & MS=RPT & FINS=1 & FER=0	FER=1; TER=R; CTE=(CTE+1)
	053	FR_AC(M=1) & FAM=1 & MS=RPT & FINS=1 & FER=1 & CTE<255	CTE=(CTE+1)
	361	FR_LLC(DA=MA & A=1) & MS=RPT & FINS=1 & FER=0	[FER=1; TER=R; CFCE=(CFCE+1) (optional-x)]
	362	FR_LLC(DA=MA & A=1) & MS=RPT & FINS=1 & FER=1 & CFCE<255	[CFCE=(CFCE+1) (optional-x)]
	041	FR_MAC(DA=MA & A=1) & MS=RPT & FINS=1 & FER=0	FER=1; TER=R; CFCE=(CFCE+1)
	042	FR_MAC(DA=MA & A=1) & MS=RPT & FINS=1 & FER=1 & CFCE<255	CFCE=(CFCE+1)
	152	FR_NOT_COPIED & MS=RPT & FINS=1 & FER=0	FER=1; TER=R; CRCE=(CRCE+1)
	153	FR_NOT_COPIED & MS=RPT & FINS=1 & FER=1 & CRCE<255	CRCE=(CRCE+1)
	179	FR_RPRT_ERR(SA=MA & A=1 & C<>0) & CER<n3 & FECO=0 & FOP=1	FER=FLF=0; CER=0; SET ERR_CNTR to 0
	180	FR_RPRT_ERR(SA=MA & A=1 & C<>0) & CER<n3 & FECO=1 & FOP=1	FLF=0; CER=0; SET ERR_CNTR to 0
	181	FR_RPRT_ERR(SA=MA & E=0 & A=0) & CER<n3 & FECO=0 & FOP=1	FER=FLF=0; CER=0; SET ERR_CNTR to 0
	182	FR_RPRT_ERR(SA=MA & E=0 & A=0) & CER<n3 & FECO=1 & FOP=1	FLF=0; CER=0; SET ERR_CNTR to 0
	183	FR_RPRT_ERR(SA=MA & E=0 & A=1 & C=0) & CER<n3 & FOP=1	CER=(CER+1); QUE_RPRT_ERR_PDU
	184	FR_RPRT_ERR(SA=MA & E=1 & C=0) & CER<n3 & FOP=1	CER=(CER+1); QUE_RPRT_ERR_PDU
	194	FR_SMP(A=0 & C=0) & FSMP=1 & MS=RPT & FINS=1 & FAM=0 & CNNR=0 & FER=0	FER=1; TER=R; CACE=(CACE+1)
	195	FR_SMP(A=0 & C=0) & FSMP=1 & MS=RPT & FINS=1 & FAM=0 & CNNR=0 & FER=1 & CACE<255	CACE=(CACE+1)
	196	FR_SMP(A=0 & C=0) & MS=RPT & FINS=1 & FAM=1 & FNN=1 & CNNR=0 & FER=0	TER=R; FER=1; CACE=(CACE+1)
	197	FR_SMP(A=0 & C=0) & MS=RPT & FINS=1 & FAM=1 & FNN=1 & CNNR=0 & FER=1 & CACE<255	CACE=(CACE+1)
	204	FR_WITH_ERR(E=0) & MS=RPT & FINS=1 & FER=0	FER=1; TER=R; CLE=(CLE+1)
	205	FR_WITH_ERR(E=0) & MS=RPT & FINS=1 & FER=1 & CLE<255	CLE=(CLE+1)
	222	INTERNAL_ERR(correctable) & MS=RPT & FINS=1 & FER=0	FER=1; TER=R; CIE=(CIE+1)

Table 7.4—MAC Error Handling Station Operation table (*Continued*)

S/T	REF	Event/condition	Action/output
	223	INTERNAL_ERR(correctable) & MS=RPT & FINS=1 & FER=1 & CIE<255	CIE=(CIE+1)
	252	STATION_ERR(correctable) & TS=DATA & FINS=1 & MS=RPT & FER=0	FER=1; TER=R; CABE=(CABE+1)
	253	STATION_ERR(correctable) & TS=DATA & FINS=1 & MS=RPT & FER=1 & CABE<255	CABE=(CABE+1)
	366	STATION_ERR(tx_underrun) & TS=DATA & FINS=1 & MS=RPT & FER=0	FER=1; TER=R; CABE=(CABE+1)
	367	STATION_ERR(tx_underrun) & TS=DATA & FINS=1 & MS=RPT & FER=1 & CABE<255	CABE=(CABE+1)
	274	TER=E & FINS=1 & ERR_CNTR<>0	CER=1; QUE_RPRT_ERR_PDU
	275	TER=E & FINS=1 & FECO=1	TER=R
	291	TK(M=1) & FAM=1 & MS=RPT & FINS=1 & FER=0	FER=1; TER=R; CTE=(CTE+1)
	292	TK(M=1) & FAM=1 & MS=RPT & FINS=1 & FER=1 & CTE<255	CTE=(CTE+1)
	302	TK_ERR & TS=DATA & FINS=1 & FER=0	FER=1; TER=R; CABE=(CABE+1)
	303	TK_ERR & TS=DATA & FINS=1 & FER=1 & CABE<255	CABE=(CABE+1)
	308	TK_WITH_ERR(E=0) & FTEO=0 & MS=RPT & FINS=1 & FER=0	FER=1; TER=R; CLE=(CLE+1)
	309	TK_WITH_ERR(E=0) & FTEO=0 & MS=RPT & FINS=1 & FER=1 & CLE<255	CLE=(CLE+1)
	339	TVX=E & FTHO=0 & FAM=1 & MS=RPT & FINS=1 & FER=0	FER=1; TER=R; CTE=(CTE+1)
	340	TVX=E & FTHO=0 & FAM=1 & MS=RPT & FINS=1 & FER=1 & CTE<255	CTE=(CTE+1)
	342	TVX=E & FTHO=1 & FAT=0 & FAM=1 & MS=RPT & FINS=1 & FER=0	FER=1; TER=R; CTE=(CTE+1)
	343	TVX=E & FTHO=1 & FAT=0 & FAM=1 & MS=RPT & FINS=1 & FER=1 & CTE<255	CTE=(CTE+1)
	349	TX_ERR(RPRT_ERR) & CER<n3 & FOP=1	QUE_RPRT_ERR_PDU; [CER=CER+1 (optional)]

Table 7.5—MAC Interface Signals Station Operation table

S/T	REF	Event/condition	Action/outout
	037	FR(RI_NOT_PRESENT) & FJR=1 & MS<>BNT & TS=RPT	M_UNITDATA.indication
	360	FR(RI_PRESENT) & FJR=1 & MS<>BNT	M_UNITDATA.indication
	134	FR_LLC(DA=Any_recognized_address) & FJR=1 & MS<>BNT	MA_UNITDATA.indication
	135	FR_MAC(DA=Any_recognized_address & DC<>0) & FJR=1 & MS<>BNT	MGT_UNITDATA.indication
	217	FTI=0	PS_CONTROL.request(Repeat_mode=Repeat)
	218	FTI=1	PS_CONTROL.request(Repeat_mode=Fill)
	220	FTXC=0	PS_CONTROL.request(Crystal_transmit=Not_asserted)
	221	FTXC=1	PS_CONTROL.request(Crystal_transmit=Asserted)
	233	M_UNITDATA.request & FJR=1 & FR_LTH<=MAX_TX	QUE_PDU
	235	M_UNITDATA.response & FR_COPIED & FTI=0 & TS=RPT	SET C=1
	234	M_UNITDATA.response & FTI=0 & TS=RPT	SET A=1
	236	MA_UNITDATA.request & FJR=1 & FR_LTH<=MAX_TX	QUE_PDU
	237	MGT_UNITDATA.request (SC<>RS) & FJR=1 & FR_LTH<=MAX_TX	QUE_PDU
	239	PM_STATUS.indication(Signal_detection= Signal_acquired) & FSLD=1	FSL=FSLD=0
	240	PM_STATUS.indication(Signal_detection= Signal_loss) & FSLD=0	FSLD=1; TSL=R
	241	PM_STATUS.indication(Wire_fault=Detected) & FWFA=1 & FWF=0	FWF=1; TWF=R
	242	PM_STATUS.indication(Wire_fault=Not_detected) & FWF=1	FWF=0
	243	PS_STATUS.indication(Burst4_error)	[FSD=0 (optional-i)]
	244	PS_STATUS.indication(Frequency_error) & MS=RPT & FINS=1 & FER=0	FER=1; TER=R; CFE=(CFE+1)
	245	PS_STATUS.indication(Frequency_error) & MS=RPT & FINS=1 & FER=1 & CFE<255	CFE=(CFE+1)
	288	TK(M=0) & FAM=1 & FTHO=0	TVX=R; PS_EVENT.response(Token_received)
	289	TK(M=0) & FAM=1 & FTHO=1	FAT=1; PS_EVENT.response(Token_received)

Table 7.6—MAC Miscellaneous Frames Handling Signals Station Operation table

S/T	REF	Event/condition	Action/outout
	097	FR_CHG_PARM & FOP=1	SET APPR_PARMS
	098	FR_CHG_PARM(DA=broadcast & A|C<>1 & CORR_NOT_PRESENT) & FOP=1	QUE_RSP_PDU(DC=RCV_SC; SC=RS; CORR=UNK_VALUE (optional-x); RSP_TYPE=0001)
	099	FR_CHG_PARM(DA=broadcast & A|C<>1 & CORR_PRESENT) & FOP=1	QUE_RSP_PDU(DC=RCV_SC; SC=RS; CORR=RCV_CORR; RSP_TYPE=0001)
	101	FR_CHG_PARM(DA=Non_broadcast & CORR_NOT_PRESENT) & FOP=1	QUE_RSP_PDU(DC=RCV_SC; SC=RS; CORR=UNK_VALUE (optional-x); RSP_TYPE=0001)
	102	FR_CHG_PARM(DA=Non_broadcast & CORR_PRESENT) & FOP=1	QUE_RSP_PDU(DC=RCV_SC; SC=RS; CORR=RCV_CORR; RSP_TYPE=0001)
	128	FR_INIT & FOP=1	SET APPR_PARMS
	129	FR_INIT(DA=broadcast & A|C<>1 & CORR_NOT_PRESENT) & FOP=1	QUE_RSP_PDU(DC=RPS; SC=RS; CORR=UNK_VALUE (optional-x); RSP_TYPE=0001)
	130	FR_INIT(DA=broadcast & A|C<>1 & CORR_PRESENT) & FOP=1	QUE_RSP_PDU(DC=RPS; SC=RS;CORR=RCV_CORR; RSP_TYPE=0001)
	132	FR_INIT(DA=Non_broadcast & CORR_NOT_PRESENT) & FOP=1	QUE_RSP_PDU(DC=RPS; SC=RS; CORR=UNK_VALUE (optional-x); RSP_TYPE=0001)
	133	FR_INIT(DA=Non_broadcast & CORR_PRESENT) & FOP=1	QUE_RSP_PDU(DC=RPS; SC=RS;CORR=RCV_CORR; RSP_TYPE=0001)
	136	FR_MAC_INV(ERR_COND=LONG_MAC & SC<>RS & CORR_NOT_PRESENT) & FOP=1	QUE_RSP_PDU(DC=RCV_SC; SC=RS; CORR=UNK_VALUE (optional-x); RSP_TYPE=8009)
	137	FR_MAC_INV(ERR_COND=LONG_MAC & SC<>RS & CORR_PRESENT) & FOP=1	QUE_RSP_PDU(DC=RCV_SC; SC=RS; CORR=RCV_CORR; RSP_TYPE=8009)
	138	FR_MAC_INV(ERR_COND=SC_INVALID & SC<>RS & CORR_NOT_PRESENT) & FOP=1	QUE_RSP_PDU(DC=RCV_SC; SC=RS; CORR=UNK_VALUE (optional-x); RSP_TYPE=8004)
	139	FR_MAC_INV(ERR_COND=SC_INVALID & SC<>RS & CORR_PRESENT) & FOP=1	QUE_RSP_PDU(DC=RCV_SC; SC=RS; CORR=RCV_CORR; RSP_TYPE=8004)
	140	FR_MAC_INV(ERR_COND=SHORT_MAC & SC_NOT_PRESENT) & FOP=1	[QUE_RSP_PDU(DC<>RS; SC=RS; CORR=UNK_VALUE (optional-x); RSP_TYPE=8001) (optional-x)]
	141	FR_MAC_INV(ERR_COND=SHORT_MAC & SC_PRESENT & SC<>RS) & FOP=1	QUE_RSP_PDU(DC=RCV_SC; SC=RS; CORR=UNK_VALUE (optional-x); RSP_TYPE=8001)
	142	FR_MAC_INV(ERR_COND=SV_LTH_ERR & SC<>RS & CORR_NOT_PRESENT) & FOP=1	QUE_RSP_PDU(DC=RCV_SC; SC=RS; CORR=UNK_VALUE (optional-x); RSP_TYPE=8005)
	143	FR_MAC_INV(ERR_COND=SV_LTH_ERR & SC<>RS & CORR_PRESENT) & FOP=1	QUE_RSP_PDU(DC=RCV_SC; SC=RS; CORR=RCV_CORR; RSP_TYPE=8005)
	144	FR_MAC_INV(ERR_COND=SV_MISSING & SC<>RS & CORR_NOT_PRESENT) & FOP=1	QUE_RSP_PDU(DC=RCV_SC; SC=RS; CORR=UNK_VALUE (optional-x); RSP_TYPE=8007)
	145	FR_MAC_INV(ERR_COND=SV_MISSING & SC<>RS & CORR_PRESENT) & FOP=1	QUE_RSP_PDU(DC=RCV_SC; SC=RS; CORR=RCV_CORR; RSP_TYPE=8007)
	146	FR_MAC_INV(ERR_COND=SV_UNK & SC<>RS & CORR_NOT_PRESENT) & FOP=1	QUE_RSP_PDU(DC=RCV_SC; SC=RS; CORR=UNK_VALUE (optional-x); RSP_TYPE=8008)
	147	FR_MAC_INV(ERR_COND=SV_UNK & SC<>RS & CORR_PRESENT) & FOP=1	QUE_RSP_PDU(DC=RCV_SC; SC=RS; CORR=RCV_CORR; RSP_TYPE=8008)
	148	FR_MAC_INV(ERR_COND=VI_LTH_ERR & SC<>RS & CORR_NOT_PRESENT) & FOP=1	QUE_RSP_PDU(DC=RCV_SC; SC=RS; CORR=UNK_VALUE (optional-x); RSP_TYPE=8002)
	149	FR_MAC_INV(ERR_COND=VI_LTH_ERR & SC<>RS & CORR_PRESENT) & FOP=1	QUE_RSP_PDU(DC=RCV_SC; SC=RS; CORR=RCV_CORR; RSP_TYPE=8002)
	150	FR_MAC_INV(ERR_COND=VI_UNK & SC<>RS & CORR_NOT_PRESENT) & FOP=1	QUE_RSP_PDU(DC=RCV_SC; SC=RS; CORR=UNK_VALUE (optional-x); RSP_TYPE=8003)
	151	FR_MAC_INV(ERR_COND=VI_UNK & SC<>RS & CORR_PRESENT) & FOP=1	QUE_RSP_PDU(DC=RCV_SC; SC=RS; CORR=RCV_CORR; RSP_TYPE=8003)
	154	FR_REMOVE(DA=Broadcast & A<>1) & FOP=1	[QUE_RSP_PDU(DC=RCV_SC; SC=RS; RSP_TYPE=800A) (optional-i)]
	155	FR_REMOVE(DA=Broadcast & A=1) & FOP=1	[QUE_RSP_PDU(DC=RCV_SC; SC=RS; RSP_TYPE=800A) (optional-x)]

Table 7.6—MAC Miscellaneous Frames Handling Signals Station Operation table (*Continued*)

S/T	REF	Event/condition	Action/outout
	162	FR_REMOVE(DA=Non_broadcast) & FRRO=1 & FOP=1	QUE_RSP_PDU(DC=RCV_SC; SC=RS; RSP_TYPE=800A)
	185	FR_RQ_ADDR & FOP=1	QUE_RPRT_ADDR_PDU
	186	FR_RQ_ATTACH & FOP=1	QUE_RPRT_ATTACH_PDU
	190	FR_RQ_STATE & FOP=1	QUE_RPRT_STATE_PDU

4.3.5 Precise specification of terms

4.3.5.1 Precise specification of events/conditions

The following definitions are applied to the terms used for events in the FSMs and Station Operation table.

Unless otherwise specified, the following terms and operations are defined:

{term1} = {term2}. Term 1 is equal to term 2.
{term1} < **{term2}.** Term 1 is less than term 2.
{term1} <= **{term2}.** Term 1 is less than or equal to term 2.
{term1} > **{term2}.** Term 1 is greater than term 2.
{term1} >= **{term2}.** Term 1 is greater than or equal to term 2.
{term1} <> **{term2}.** Term 1 is not equal to term 2.
{flag}=0. The specified flag is set to zero (false).
{flag}=1. The specified flag is set to one (true).
{timer}=E. The specified timer has expired.
& means “and”.
| means “or”.

A=0. Both the A bits in the received frame’s FS field (bits 0 and 4) were 0.
A=1. Both the A bits in the received frame’s FS field (bits 0 and 4) were 1.
A<>1. Either or both A bits in the received frame’s FS field (bits 0 and 4) were 0.
A|C<>0. At least one of the frame status bits in the received frame’s FS field (bits 0,1,4, or 5) is 1.
A|C<>1. At least one of the frame status bits in the received frame’s FS field (bits 0,1,4, or 5) is 0.
BN_TYPE. The value of the beacon type subvector received.
Burst5_error_event. A conditional PM_STATUS.indication(Burst5_error) has occurred. The conditions under which a Burst5_error is excluded is not uniquely specified by this standard (see 3.6.3). At a minimum, the station shall include the first Burst5_error following a valid MAC frame copied by the station if the Burst5_error occurs within a frame. The station may include every Burst5_error.
C=0. Both the C bits in the received frame’s FS field (bits 1 and 5) were set to 0.
C=1. Both the C bits in the received frame’s FS field (bits 1 and 5) were set to 1.
C<>0. Either or both C bits in the received frame’s FS field (bits 1 and 5) were set to 1.
Connect.MAC. The MAC receives the command from management to start the connection process to join the ring.
CORR_NOT_PRESENT. The received frame did not contain a correlator subvector.

CORR_PRESENT. The received frame did contain a correlator subvector.

CTO-FR_LTH<0. The value of the Transmit Octet counter is less than the next frame length to be transmitted.

CTO-FR_LTH>=0. The value of the Transmit Octet counter is greater than, or equal to, the next frame length to be transmitted.

DA=any_recognized_address. The DA of the received frame matches any of the station's addresses as follows:

a) Is the station's individual address (DA=MA), or
b) Is one of the station's group addresses, or
c) Is one of the station's functional addresses, or
d) Is one of the broadcast addresses as defined in 3.2.4.1.

DA=MA. The DA of the received frame is equal to the individual address of the station. If the station's individual address is a universally administered address, then all 48 bits must match. If the station's individual address is a locally administrated address, then either a hierarchical address match or a 48-bit address match is allowed.

DA=Non_broadcast. The received frame was not sent to a broadcast address, but otherwise addressed to the station.

Disconnect.MAC. The request from local management to remove the station from the ring.

E=0. The error bit in the received ED field is zero.

E=1. The error bit in the received ED field is one.

EOD. End of data: The last octet of the Information field has been transmitted.

ERR_CNTR<>0. Any error counter not zero.

ERR_COND=LONG_MAC. MAC frame too long—INFO field larger than maximum allowed VL value.

ERR_COND=SC_INVALID. Invalid source class.

ERR_COND=SHORT_MAC. MAC frame not long enough to contain VL, VC, and VI fields.

ERR_COND=SV_LTH_ERR. Subvector length error.

ERR_COND=SV_MISSING. Missing required subvector.

ERR_COND=SV_UNK. Unknown subvector SVI value.

ERR_COND=VI_LTH_ERR. Vector length error. VL is not equal to the sum of all the SVLs plus the length of VL,VC, and VI fields, or VL does not agree with the length of the frame.

ERR_COND=VI_UNK. Unrecognized vector ID value.

FR. A frame has been received that meets the criteria specified in 4.3.2.

FR(criteria). A frame has been received which meets the specified criteria and the criteria specified in 4.3.2.

FR_AC(criteria). A frame's access control field has been received which meets the specified criteria and the criteria specified in 4.3.2.

FR_AMP(criteria). A verified Active Monitor Present frame (3.3.5.2) is received which meets the specified criteria.

FR_BN(criteria). A verified Beacon frame (3.3.5.2) is received which meets the specified criteria.

FR_CHG_PARM(criteria). A verified Change Parameters MAC frame (3.3.5.2) is received which meets the specified criteria.

FR_BN_CIR. A frame's access control field has been detected with the M bit set to 1 indicating a circulating frame. The method used to detect this condition is outside the scope of this standard, but reception of a valid beacon frame (3.3.5.2) with the M bit set to 1 shall satisfy this condition.

FR_COPIED(criteria). The MAC successfully copied the received frame which meets the specified criteria.

FR_CT(criteria). A verified Claim Token MAC frame (3.3.5.2) is received which meets the specified criteria.

FR_DAT(criteria). A verified DAT MAC frame (3.3.5.2) is received which meets the specified criteria.

FR_INIT(criteria). A verified Initialize Station MAC frame (3.3.5.2) is received which meets the specified criteria.

FR_LLC(criteria). An LLC frame is received which meets the specified criteria and the criteria specified in 4.3.2.

FR_LTH. The length of the frame to be transmitted. The value for the frame length includes all of the frame format fields beginning with the starting delimiter (SD) and including the interframe gap (IFG).

FR_LTH<=MAX_TX. The length of the frame to be transmitted is less than or equal to the maximum allowed frame length.

FR_MAC(criteria). A valid MAC frame is received which meets the specified criteria and the criteria specified in 4.3.2.

FR_MAC_INV(reason). A valid (4.3.2) MAC frame is received which fails verification (3.3.5.2) for the reason specified.

FR_NOT_COPIED. The station sets the A bits (ref 040 or 234) but does not copy the frame.

FR_REMOVE(criteria). A verified Remove MAC frame (3.3.5.2) is received which meets the specified criteria.

FR_RP(criteria). A verified Ring Purge MAC frame (3.3.5.2) is received which meets the specified criteria.

FR_RPRT_ERR(criteria). A Report Error MAC frame (3.3.5.1) is received which was transmitted by the station without error (4.3.3) and that meets the specified criteria.

FR_RQ_ADDR. A verified Request Address MAC frame (3.3.5.2) is received.

FR_RQ_ATTACH. A verified Request Attachment MAC frame (3.3.5.2) is received.

FR_RQ_INIT(criteria). A Request Initialization MAC frame (3.3.5.1) is received which was transmitted by the station without error (4.3.3) and meets the specified criteria, is received.

FR_RQ_STATE. A verified Request Station State MAC frame (3.3.5.2) is received.

FR_SMP(criteria). A verified SMP MAC frame (3.3.5.2), which meets the specified criteria, is received.

FR_SUA_CHG(criteria). A Report SUA Change MAC frame (3.3.5.1), which was transmitted by the station without error (4.3.3) and meets the specified criteria, is received.

FR_WITH_ERR. A frame is received with errors (see 4.3.2).

FR_WITH_ERR(criteria). A Frame with Error (see 4.3.2) is received which meets the specified criteria.

INTERNAL_ERR. Any internal error occurred that prevented the station from following the established protocol (i.e., parity error, etc.).

JS=state. The Join State is in the specified state.

LTA. The stored SA field from the last transmitted frame.

M=0. The Monitor bit in the AC field is received as zero.

M=1. The Monitor bit in the AC field is received as one.

M_UNITDATA.request—The bridge interface requests a frame be transmitted.

M_UNITDATA.response—The optional indication from the bridge interface in response to an M_UNITDATA.indication that requests setting of the A and C bits.

MA_UNITDATA.request—The LLC interface requests a frame be transmitted.

MGT_UNITDATA.request—The management interface requests a frame be transmitted.

MAX_TX. The maximum number of octets that may be transmitted (including fill) after capturing the token as specified in 4.2.4.1, Counter, Transmitted Octets (CTO).

MS=state. The Monitor State is in the specified state.

P. Value of P bits in AC field.

P>Pm>R. A token or frame is received and the priority of the PDU queued for transmission is greater than the received reservation (R) but less that the current service priority (P).

PM_STATUS.indication(Signal_detection=Signal_acquired). The PHY indicates valid receiver signal (see 5.1.4.1).

PM_STATUS.indication(Signal_detection=Signal_loss). The PHY indicates loss of valid receiver signal (see 5.1.4.1).

PM_STATUS.indication(Wire_fault=Detected). The PHY indicates a wiring fault (see 5.1.4.1).

PM_STATUS.indication(Wire_fault=Not_detected). The PHY indicates no wiring fault (see 5.1.4.1).

PDU_QUEUED. A frame is queued for transmission.

PDU_QUEUED(criteria). A frame is queued for transmission which meets the specified criteria. Note that a queued PDU is transmitted when a token is received with a priority less than or equal to the priority of the queued PDU (Pm). Frames that do not wait for a token are not queued but are indicated by TXI_REQ.

PS_STATUS.indication(Frequency_error). The station indicates the frequency of the received data is out of tolerance (see 5.1.2.3).

PS_STATUS.indication(Medium_rate_error). The station indicates the frequency of the received data is not the selected rate (see 5.1.2.3).

PS_STATUS.indication(Burst4_error). The station indicates the received data contains a Burst4_error (see 5.1.2.3).

PS_STATUS.indication(Burst5_error). The station indicates the received data contains a Burst5_error (see 5.1.2.3).

Pm. The priority of the queued PDU. If no PDU is queued, Pm is assumed to have a value of zero.

Px. A priority value representing the higher value of either (a) the received reservation or (b) the priority of the frame queued for transmission.
If PDU_QUEUED and Pm>Rr then Px=Pm
Else Px=Rr.

QUE_EMPTY. No frames are queued for transmission.

QUE_NOT_EMPTY. Another PDU is queued for transmission.

R. Value of R bits in AC field.

RCV_AC. The Access Control field is received.

RCV_ED. An ending delimiter (PS_STATUS.indication(Ending_delimiter) see 5.1.2.3) is received.

RCV_SD. A starting delimiter (PS_STATUS.indication(Starting_delimiter) see 5.1.2.3) is received.

RCV(SA<>LTA). The source address of the received frame is not the same as the source address of the last frame transmitted by this station. This condition requires that no code violations are present in the SA field and optionally any of the preceding fields.

RCV(SA=LTA). The source address of the received frame is the same as the source address of the last frame transmitted by this station. The station is not required to, but may check for code violations in the SA or any preceding field to determine the validity of this event.

RI_PRESENT. The RII bit (3.2.4.2) of the received frame is one indicating the Routing Information Field (RIF) is present.

RI_NOT_PRESENT. The RII bit (3.2.4.2) of the received frame is zero indicating the Routing Information Field (RIF) is not present.

SA<>MA. The source address (SA) of the received frame is not equal to the individual address of the station.

SA<>SUA. The source address (SA) of the received frame is not the same as the address stored as the upstream neighbor’s address (SUA).

SA<MA. The value of the source address (SA) of the received frame is numerically less than the individual address of the station. For purpose of comparison, first bit received is most significant.

SA>MA. The value of the source address (SA) of the received frame is numerically greater than the individual address of the station. For purpose of comparison, first bit received is most significant.

SA=MA. The source address (SA) of the received frame is equal to the individual address of the station.

SA=SUA. The source address (SA) of the received frame is the same as the address stored as the upstream neighbor’s address (SUA).

SC_NOT_PRESENT. The MAC frame is too short to contain the source class.

SC_PRESENT. The MAC frame does contain the source class.

SC<>RS. The source class is not 0 (Ring Station).

SC=CRS. The source class is 4 (Configuration Report Server)

SC=RPS. The source class is 5 (Ring Parameter Server).

STATION_ERR. Any internal condition which prevents the successful completion of the PDU transmit operation.

TEST_FAILURE. The station failed its self test.

TEST_OK. The station passed its self test.

TK. A token is received which meets the criteria specified in 4.3.1.

TK(criteria). A token is received which meets the specified criteria and the criteria specified in 4.3.1.

TK_AC(criteria). A token is received which meets the specified criteria and the criteria specified in 4.3.1.

TK_GOOD. A good token is received which meets the criteria specified in 4.3.1.

TK_GOOD(criteria). A good token is received which meets the specified criteria and the criteria specified in 4.3.1.

TK_ERR. The token is not valid (see 4.3.1).

TK_WITH_ERR. A token is received which contains errors (see 4.3.1).

TK_WITH_ERR(criteria). A token is received which contains errors (see 4.3.1) that meets the specified criteria.

TS=state. The transmit state is in the specified state.

TX_BN_TYPE. The value of the beacon type subvector being transmitted.

TX_ERR(DAT). During the transmission of the DAT MAC frame, a transmission error is encountered (see 4.3.3).

TX_ERR(RP). During the transmission of the Ring Purge MAC frame, a transmission error is encountered (see 4.3.3).

TX_ERR(RPRT_ERR). During the transmission of the RPRT_ERR MAC frame, a transmission error is encountered (see 4.3.3).

TX_ERR(RQ_INIT). During the transmission of the Request Initialization MAC frame, a transmission error is encountered (see 4.3.3).

TX_ERR(SUA_CHG). During the transmission of the Report SUA Change MAC frame, a transmission error is encountered (see 4.3.3).

TXI_REQ. A frame is requested to be transmitted without waiting for a token. This request is generated by the TXI(frame_type) action (e.g., TXI(BN_PDU).

UNA. The upstream neighbor’s address (UNA) subvector in the received frame.

UNA<>MA. The reported upstream neighbor’s address (UNA) in the received frame is not equal to the station’s individual address.

UNA<>SUA. The value of the reported upstream neighbor’s address (UNA) in the received frame is not the same as the stored upstream address (SUA).

UNA=MA. The reported upstream neighbor’s address (UNA) in the received frame is equal to the station’s individual address.

UNA=SUA. The value of the received reported upstream neighbor’s address (UNA) is the same as the stored upstream address (SUA).

4.3.5.2 Precise specification of actions

The following definitions are applied to the terms used for actions in the FSMs and table 7. Actions are separated by a semicolon (;).

Unless otherwise specified, the following operations are defined:

variable = value. Set the variable to the specified value.
{counter}={counter}+1. Increment the specified counter.
{counter}={counter}-1. Decrement the specified counter.
{counter}=value. Set the specified counter to the specified value.
{flag}=0. Set the value of the specified flag to zero (false).
{flag}=1. Set the value of the specified flag to one (true).
{timer}=R. The specified timer will be set to its initial value and started.

A=0. Both A bits in the FS field shall be transmitted as zero.

A=1. Both A bits in the FS field shall be set to one as the frame is repeated.

C=0. Both C bits in the FS field shall be transmitted as zero.

C=1. Both C bits in the FS field shall be set to one as the frame is repeated.

CLEAR_STACKS. The station shall clear all stacked values for Sx and Sr (i.e., the list is empty and the effective value of Sx and Sr is -1). The station shall no longer be responsible for lowering the ring priority.

CTO=CTO-FR_LTH. The value of the Transmit Octet counter is decreased by the number of octets (frame length) of the frame to be transmitted. The value for the frame length includes all of the frame format fields beginning with the starting delimiter (SD) and including the interframe gap (IFG). This correlates to the counter CTO being decremented every 8 bits after the capture of the token.

CTO=MAX_TX-FR_LTH. The value of the Transmit Octet counter is set to the value of the maximum frame length in octets less the number of octets (frame length) of the frame to be transmitted. The value for the frame length includes all of the frame format fields beginning with the starting delimiter (SD) and including the interframe gap (IFG).

CORR=RCV_CORR. The value of the correlator subvector will be the same value as the received correlator subvector.

CORR=UNK_VALUE (optional-x). The frame received did not contain a correlator subvector (3.3.4), thus the value of the correlator subvector to be transmitted is unspecified and the subvector may be omitted. The standard recommends new implementations not transmit the correlator subvector when no correlator subvector was received.

DC=CRS. The value of the destination class is 4 (Configuration Report Server).

DC=RCV_SC. The destination class field DC shall contain the value of the source class field (SC) of the received frame.

DC=RPS. The value of the destination class is 5 (Ring Parameter Server).

DC<>RS. The destination class field DC shall not be 0. Note that the source class field (SC) of the received frame was not present and thus the destination class of the response frame is not defined but shall not be the ring station class.

E=0. The error (E) bit in the ending delimiter (ED) field shall be transmitted as zero.

E=1. The error (E) bit in the ending delimiter (ED) field shall be set to one as the frame is repeated.

FA(monitor)=0. Disable the functional address corresponding to the active monitor function.

FA(monitor)=1. Enable the functional address corresponding to the active monitor function.

FTI=x. The value of FTI is not specified.

INSERT. Request the PHY to physically connect the station into the ring [5.1.4.2 PM_CONTROL.request(Insert_station)].

JS=state. The join state is changed to the specified state.

LTA=TX_SA. Capture the source address (SA) of the last transmitted frame as variable LTA.

LMP=(MA|NULL). The address for the Last Monitor Present frame (X'0A' subvector 3.3.4) is set to either the station’s address or the null address.

LMP=NULL. The address for the Last Monitor Present frame (X'0A subvector 3.3.4) is the null address.

LMP=SA. The SA is saved for reporting the address of the Last Monitor Present frame (X'0A' subvector 3.3.4).

M=0. The station shall transmit the monitor bit (M) in the AC field as a zero.

M=1. The station shall set the monitor bit (M) in the received AC field to one as the AC field is repeated.

M_UNITDATA.indication—The frame is indicated to the bridge interface.

MA_UNITDATA.indication—The frame is indicated to the LLC interface.

MGT_UNITDATA.indication—The frame is indicated to the management interface.

MS=state. The monitor state is changed to the specified state.

P. The value of the P bits in the AC field.

Pm. The priority of the PDU being queued.

Pr. The value of the P bits in the last received AC field.

PM_CONTROL.request(Insert_station). The MAC requests the station be inserted into the ring (see 5.1.4.2).

PM_CONTROL.request(Remove_station). The MAC requests the station be removed from the ring (see 5.1.4.2).

POP(Sx,Sr). Remove the last values of Sx and Sr from the list of stacked priorities on the Sx and Sr stacks.

PS_CONTROL.request(Crystal_transmit=Asserted). The MAC requests Crystal_transmit (see 5.1.2.4).

PS_CONTROL.request(Crystal_transmit=Not_asserted). The MAC removes the Crystal_transmit request (see 5.1.2.4).

PS_CONTROL.request(Repeat_mode=Fill). The MAC requests the station sources fill (see 5.1.2.4).

PS_CONTROL.request(Repeat_mode=Repeat). The MAC requests the station repeat (see 5.1.2.4).

PS_EVENT.response(Token_received). The MAC indicates token received (see 5.1.2.5).

QUE_ACT_ERR_PDU(EC=value). Queue a Report Active Monitor Error MAC PDU as defined in 3.3.5.1 for transmission with the specified error code (EC).

QUE_AMP_PDU. Queue an Active Monitor Present (AMP) MAC PDU as defined in 3.3.5.1 for transmission.

QUE_DAT_PDU. Queue a Duplicate Address Test (DAT) MAC PDU as defined in 3.3.5.1 for transmission.

QUE_PDU. Queue the PDU for transmission.

QUE_NEW_MON_PDU. Queue a Report New Active Monitor MAC PDU as defined in 3.3.5.1 for transmission.

QUE_NN_INCMP_PDU. Queue a Report Neighbor Notification Incomplete MAC PDU as defined in 3.3.5.1 for transmission.

QUE_RPRT_ADDR_PDU. Queue a Report Station Addresses MAC PDU as defined in 3.3.5.1 for transmission.

QUE_RPRT_ATTCH_PDU. Queue a Report Station Attachment Report Attachment MAC PDU as defined in 3.3.5.1 for transmission.

QUE_RPRT_ERR_PDU. Queue a Report Error MAC PDU as defined in 3.3.5.1 for transmission.

QUE_RPRT_STATE_PDU. Queue a Report Station State MAC PDU as defined in 3.3.5.1 for transmission.

QUE_RQ_INIT_PDU. Queue a Request Initialization MAC PDU as defined in 3.3.5.1 for transmission.

QUE_RSP_PDU. Queue a Response MAC PDU as defined in 3.3.5.1 for transmission. This PDU is not required to be queued if the received frame initiating the response was sent to a broadcast address and the address-recognized and frame-copied bits in the Frame Status field were all ones.

QUE_SMP_PDU. Queue a Standby Monitor Present MAC PDU as defined in 3.3.5.1 for transmission.

QUE_SUA_CHG_PDU. Queue a SUA Change MAC PDU as defined in 3.3.5.1 for transmission.

R. The value of the R bits in the AC field.

Rr. The value of the R bits in the last received AC field.

R=Pm. Set the reservation bits (R) in the AC field to the value of the queued PDU.

RESET_STACKS(Sx=P; Sr=0). The station clears all stacked values for Sx and Sr and then adds the new Sx and Sr values.

RESTACK(Sx=Px). Replace the last value of stack Sx with Px (the value of P bits of the token to be transmitted).

Remove_station. Request the PHY to physically disconnect the station from the ring [5.1.4.2 PM_CONTROL.request(Remove_station)].

RSP_TYPE=value. The Response Code subvector shall have the hexadecimal value specified.

SC=RS. The Source Class field (SC) shall contain the value zero (Ring Station).

SET <field> = <value>. When the MAC repeats the field, it will set the field to the specified value (see M=1, R=Pm, E=1, A=1, and C=1).

SET APPR_PARMS. The station shall set the station’s parameters to the values indicated in the received frame.

SET ERR_CNTR to 0. Set the values for all of the error counters reported in the Report Error MAC frame to zero.

Set_initial_conditions. The station shall set all MAC flags to zero, set all MAC counters to zero, and stop all timers. The Monitor FSM and Transmit FSM are not specified. The PS_CONTROL.request(Medium_rate) and PM_CONTROL.request(Medium_rate) shall indicate to the PHY the value of FMRO.

STACK(Sx=Px; Sr=Pr). Add the new Sx and Sr values to the list of stacked priorities on the Sx and Sr stacks.

STORE(Pr; Rr). Save the value of P and R in the received AC field as Pr and Rr, respectively.

SUA=0. Store the null address as the station's SUA.

SUA=SA. Store the value of the SA from the received frame as the station's SUA.

TEST. The station shall perform a test of its transmit functions, its receive functions, and the medium between the station and the TCU. It is recommended that the data path includes the elastic buffer and the fixed latency buffer (5.8). A station shall fail the test if the sustained bit error rate does not meet the criteria specified in annex P. A station shall only transmit valid frames, tokens, and fill during the test and shall only count errors in frames and tokens.

TS=state. The transmit state is changed to the specified state.

TX_AB. The station shall transmit an abort delimiter.

TX_BN_TYPE. The value of the beacon type subvector to be transmitted.

TX_EFS(I=0). The station shall transmit an EFS composed of ED, FS, and IFG fields. The E,I,A, and C bits shall be zero.

TX_EFS(I=x). The station shall transmit an EFS composed of ED, FS, and IFG fields. The I bit may be zero or one, and the E,A, and C bits shall be zero.

TX_FCS. The station shall transmit FCS for the frame as defined in 3.2.7.

TX_SFS(P=value; M=0; R=value). The station shall transmit the SFS with the priority and reservation values as specified. The token busy bit (T) shall be one and the monitor bit (M) shall be zero.

TX_TK(P=value; M=0; R=value). The station shall transmit a token with the priority and reservation fields as specified. The monitor (M) bit and the token busy (T) bit shall be transmitted as zero.

TXI(BN_PDU). The station shall transmit a Beacon MAC frame with the AC field values of P=000, T=1, M=0, R=000. The frame shall contain all of the required subvectors. The transmission of the frame shall occur at the earliest opportunity (after completion of any transmission in progress) and not wait for a token. This action generates the TXI_REQ event.

TXI(CT_PDU). The station shall transmit a Claim Token MAC frame with the AC field values of P=000, T=1, M=0, R=000. The frame shall contain all of the required subvectors. The transmission of the frame shall occur at the earliest opportunity (after completion of any transmission in progress) and not wait for a token. This action generates the TXI_REQ event.

TXI(RCV_CT_PDU). The station shall transmit the received Claim Token MAC frame as received with the AC field values of P=000, T=1, M=0, R=000. The transmission of the frame shall occur at the earliest opportunity (after completion of any transmission in progress) and not wait for a token. This action generates the TXI_REQ event.

TXI(RP_PDU). The station shall transmit a Ring Purge MAC frame with the AC field values of P=000, T=1, M=0, R=000. The frame shall contain all of the required subvectors. The transmission of the frame shall occur at the earliest opportunity (after completion of any transmission in progress) and not wait for a token. This action generates the TXI_REQ event.

5. Station specific components

5.1 General

The station's physical layer (PHY) is divided into two sublayers: physical signaling components (PSC) and physical medium components (PMC). This clause describes the operation of the station's PSC. The components described in this clause provide the physical coupling between the Medium Access Control (MAC) protocol and the physical media components (PMC).

This clause defines the signal processing and system level physical layer (PHY) signaling specifications (independent of the medium type) that are specific to a ring station. These station specific component specifications include symbol timing, symbol encoding/decoding, latency, and burst error detection/correction. The specifications in this clause are generic for all media types. Frame formats are defined in clause 3 and the MAC protocol is specified in clause 4. The specifics for the PMC, which include the receiver, clock recovery, ring access control, and transmitter, are covered in clause 7.

This clause addresses the PSC that couple the PMC to the MAC and, for the purpose of specifying operation, defines the service boundaries to the PMC and to the MAC. The definition of these components are solely for the purpose of specifying operation and do not imply any particular implementation. The operation specified in this clause requires the MAC protocol and shall not be considered appropriate operation for devices that do not contain a MAC. Devices that do not support the MAC protocol (e.g., repeaters) are subject for future standardization. This clause requires that any station that modifies the data stream (e.g., perform burst error correction) shall conform to the MAC protocol, except that a device not implementing the MAC protocol may modify a data stream containing a burst6 so long as the resulting data stream has only that burst modified, and the resultant is a burst5.

5.1.1 Overview

Figure 26 shows an example of a station's data path and illustrates the service interface to the MAC and PMC.

The term used for the element of data transfer between the MAC and the differential Manchester encoder/decoder is "symbol." A symbol is characterized by one of the values of Data_zero, Data_one, Non-data_J, Non-data_K. The term used for the element of data transfer between components in the PHY is "signal element." A signal element is characterized solely by its polarity (Logic_0, Logic_1) and thus two (2) signal elements are required to represent a single symbol.

The signaling method used to send data to the ring is differential Manchester, which contains both data and timing reference information. This signaling method permits a station to recover timing information from the received data signals. Timing information recovery is specified in clause 7. When the station is not transmitting and not in recovery, it is normally repeating the received data. The repeat path inspects the received data for invalid signal patterns. By the nature of differential Manchester coding, with the pairing of Non-data_J and Non-data_K symbols, and in the absence of errors, there should never be more than 3 signal elements received in succession with the same polarity. The reception of 4 consecutive signal elements with the same polarity (burst4) is a coding error that indicates the potential loss of data and timing information. When more than 4 signal elements are received with the same polarity (burst5), the station introduces transitions into the repeated data such that only 4 signal elements are repeated with the same polarity. The station continues to introduce transitions until a transition is received. One common source of burst errors is caused by a TCU when a station is inserted or removed from the ring. The source of the burst5 is isolated by changing the burst5 into a burst4. Thus only the first station after the source of the burst condition reports it. The burst4 propagates around the ring as an indication that data has been corrupted. Note that certain physical characteristics may cause a burst4 repeated by one station to be detected as a burst5 at the next station and thus becomes a potential source for reported burst errors. The MAC protocol is notified of burst errors both for the purpose of isolation/reporting the error and also reducing over-stripping (miss-counting of transmitted frames) when two or more frames are concatenated into a single frame when a burst error destroys the ending delimiter of one frame and the starting delimiter of the next.

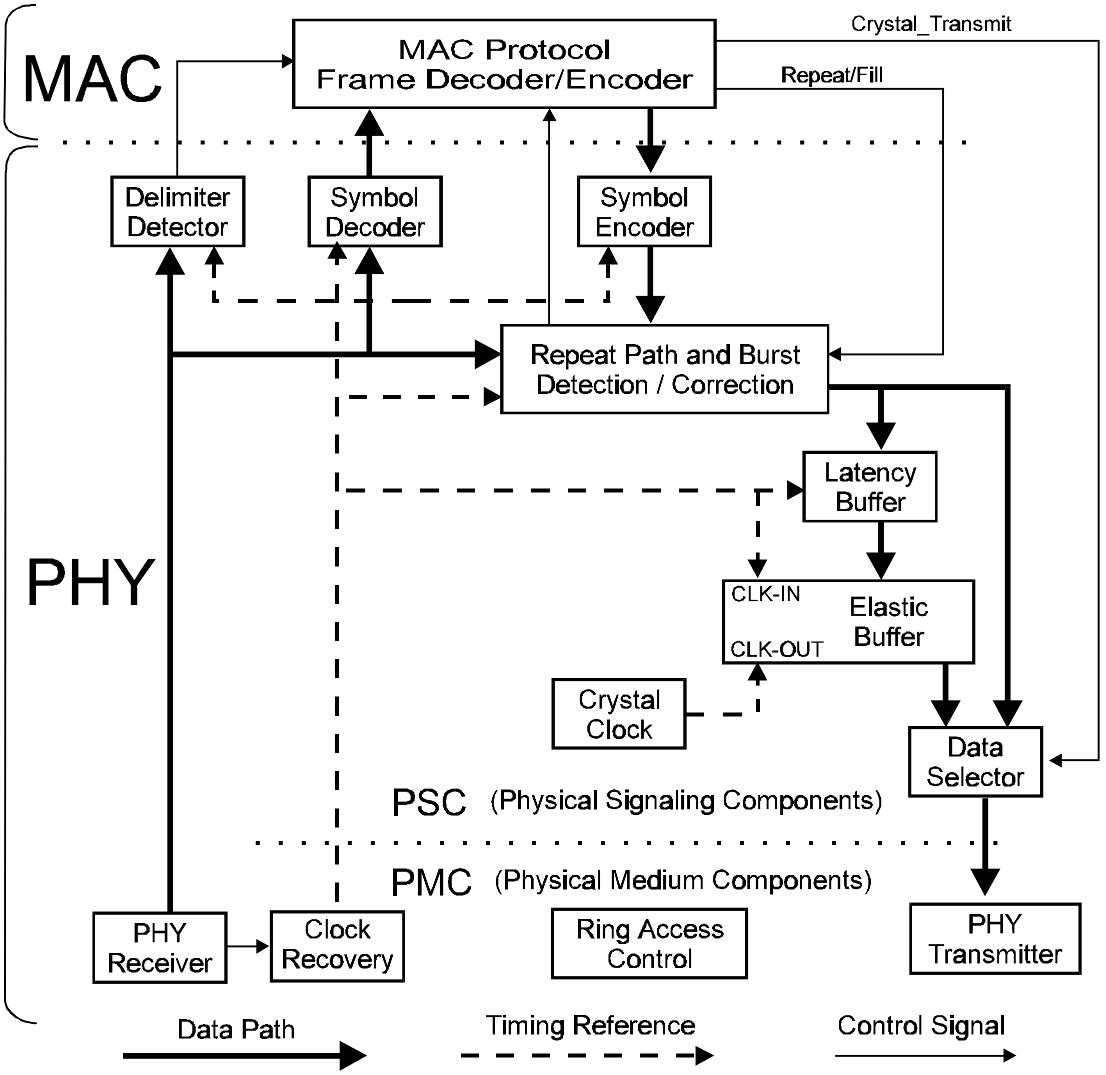

Figure 26—Station data path example

The received data is also sent to the Symbol Decoder where it is inspected for Starting Delimiter (SD) and Ending Delimiter (ED) patterns. When a delimiter is detected, the delimiter event is indicated to the MAC and the decoding of the received data stream is synchronized such that for every 2 signal elements, a symbol (Data_zero, Data_one, Non-data_J, Non-data_K) is indicated to the MAC protocol.

The MAC protocol controls the repeat path such that received data is normally repeated at the station's output except when the MAC protocol interrupts the repeat path forcing either fill or data (frames and tokens) to be transmitted. The phrase "repeated data" means that a transition on the input results in a transition on the output independent of the polarity of either the input or output signals The symbol encoder allows the MAC protocol to modify individual symbols as they are being repeated and source symbols creating frames and tokens.

The MAC protocol requests repeat, fill, or symbols (Data_zero, Data_one, Non-data_J, Non-data_K) to the PSC. The symbol encoder converts each symbol to differential Manchester code to be transmitted. Differential Manchester code is characterized by the relationship of the signal element value to the value of the previous signal element. The symbol encoder operates by creating or removing the transitions between signal elements.

The fixed latency buffer and elastic buffer are inserted in the data path when requested by the MAC protocol to provide 1) the latency required to assure that an entire token can exist within the total ring latency and 2) an elastic buffer to allow the station's output to be timed from a fixed frequency reference (crystal clock). One station provides the ring timing reference by timing its output from a fixed frequency reference while the other stations derive their output timing from their received data (recovered clock). The elastic buffer allows the station supplying the ring timing reference (crystal transmit) to compensate for dynamic variations in the cumulative ring latency.

5.1.2 MAC interface service specification

The following service primitives specify the required information that is passed between the MAC and the PSC. The PSC/MAC service specification is solely for the purpose of explaining token ring operation and does not imply any particular implementation.

5.1.2.1 PS_UNITDATA.indication

This primitive defines the transfer of data from the PSC to the MAC.

PS_UNITDATA.indication[Rcv_symbol (5.6)]

The Rcv_symbol specified is one of the following:

Data_zero
Data_one
Non-data_J
Non-data_K

When generated. The PSC generates a PS_UNITDATA.indication for each symbol received.

Effect of receipt. The MAC uses this information to process frames and tokens.

5.1.2.2 PS_UNITDATA.request

This primitive defines the transfer of data from the MAC to the PSC.

PS_UNITDATA.request [Tx_symbol (5.5)]

The Tx_symbol specified is one of the following:

Data_zero
Data_one
Non-data_J
Non-data_K

When generated. The MAC generates a PS_UNITDATA.request for each symbol of a token or frame originated by the station and for each symbol to be modified in the repeated data.

Effect of receipt. The PSC uses this information to generate the transmitted signal.

5.1.2.3 PS_STATUS.indication

This primitive is used by the PSC to inform the MAC of errors and significant status changes.

PS_STATUS.indication [Frequency_error (5.7.2),
Medium_rate_error (5.2),
Starting_delimiter (5.6),
Ending_delimiter (5.6),
ED_alignment_error (5.6),
Burst4_error (5.4.2),
Burst5_error (5.4.2)]

When generated. Upon detection of any of the conditions, the PSC generates a PS_STATUS.indication.

Effect of receipt. These signals are processed by the MAC protocol.

5.1.2.4 PS_CONTROL.request

This primitive is used by the MAC to request certain actions of the PSC.

PS_CONTROL.request [Crystal_transmit (5.8),
Medium_rate (5.2),
Repeat_mode (5.4.1)]

Crystal_transmit is specified as one of the following:

Asserted
Not asserted

Medium_rate is specified as one of the following:

4 Mbit/s
16 Mbit/s

Repeat_mode specified is one of the following:

Repeat
Fill

When generated. The MAC generates a PS_CONTROL.request for each action request.

Effect of receipt. The PSC performs the appropriate action.

5.1.2.5 PS_EVENT.response

This primitive is used by the MAC to indicate certain events to the PSC.

PS_EVENT.response [Token_received (5.8.3.1)]

When generated. The MAC generates a PS_EVENT.response for each event detected.

Effect of receipt. The PSC takes appropriate action (i.e., initializes the elastic buffer).

5.1.3 PMC interface service specification

The following service primitives specify the required information that is passed between the PSC and the PMC. The PSC/PMC service specification is solely for the purpose of explaining token ring operation and does not imply any particular implementation.

5.1.3.1 PM_UNITDATA.indication

This primitive defines the transfer of data from the PMC to the PSC.

PM_UNITDATA.indication [Rcv_signal_element (5.6)]

The Rcv_signal_element specified is one of the following:

Logic_0
Logic_1

When generated. The PHY generates a PM_UNITDATA.indication for each signal element received.

Effect of receipt. The PSC processes this information for providing data to the MAC, detecting frame delimiters, and detecting signaling errors.

5.1.3.2 PM_UNITDATA.request

This primitive defines the transfer of data from the PSC to the PMC.

PM_UNITDATA.request [Tx_signal_element (5.5)]

The Tx_signal_element specified is one of the following:

Logic_0
Logic_1

When generated. The PSC generates a PM_UNITDATA.request every signal element interval.

Effect of receipt. The PMC uses this information to generate the transmitted signal.

5.1.4 MAC/PMC interface service specification

The following service primitives specify the information that is passed between the MAC and the PMC. The MAC/PMC service specification is solely for the purpose of explaining token ring operation and does not imply any particular implementation.

5.1.4.1 PM_STATUS.indication

This primitive is used by the PMC to inform the MAC of errors and significant status changes.

PM_STATUS.indication [Signal_detection (5.7.1),
Wire_fault (5.9)]

Signal detection is specified as one of the following:

Signal_acquired

Signal_loss

Wire_fault is specified as one of the following:

Detected
Not_detected

When generated. Upon detection of any of the conditions, the PMC generates a PM_STATUS.indication.

Effect of receipt. These signals are processed by the MAC protocol.

5.1.4.2 PM_CONTROL.request

This primitive is used by the MAC to request certain actions of the PSC.

PM_CONTROL.request [Insert_station (5.9),
Remove_station (5.9),
Medium_rate (5.2)]

Medium_rate is specified as one of the following:

4 Mbit/s
16 Mbit/s

When generated. The MAC generates a PM_CONTROL.request for each action request.

Effect of receipt. The PHY performs the appropriate action.

5.2 Data signaling rate

A station shall be capable of transmitting and receiving symbols at one or both of the following nominal bit rates:

a) 4 Mbit/s
b) 16 Mbit/s

One symbol (two signal elements) shall be output per bit period.

The MAC request Medium_rate is used to select the data rate for stations supporting both data rates. The signal Medium_rate_error is optionally provided to indicate to the MAC that the station has detected the ring is operating at a different rate than the rate requested by Medium_rate. The mechanism to determine the ring media rate is not specified nor required by this standard. When the Medium_rate_error function is implemented in a station, certain ring conditions, such as circulating burst4 data patterns, may create an incorrect frequency determination that falsely set this error flag. The MAC protocol specifies the actions taken (i.e., if Medium_rate_error is detected during the join process, the station removes from the ring).

The source of a station’s transmitter timing is selected by the MAC request Crystal_transmit. All stations shall transmit signal elements timed from one of the following two sources based on the Crystal_transmit request from the MAC:

a) A stable timing reference (hereafter referred to as crystal clock) with a frequency tolerance of ±0.01% when the MAC asserts Crystal_transmit.

b) A clock recovery circuit that extracts timing information inherent in the transitions from the received data signals when the MAC is not asserting Crystal_transmit. The clock recovery circuit tracks the frequency and phase of the received data stream as specified in 7.2.3.

Under normal ring conditions, the token ring protocol places one station, called the active monitor, in the state where it derives its transmit timing from the local crystal clock. The remainder of the stations transmit symbols using the timing information recovered from the received signal. This protocol results in all the other stations transmitting data at the same long-term time-averaged rate as the received signal. However, the instantaneous rate of the data signal at each station does not remain constant due to jitter from various sources. Because the received data signal timing information is used to reclock the station's transmitted data, jitter accumulates in the ring starting at the active monitor. The jitter accumulation continues around the ring until it reaches the active monitor. Jitter does not propagate through the active monitor since its data transmission is timed from its crystal clock. An elastic buffer is used in the active monitor station to absorb the jitter accumulation and avoid adding or removing signal elements from the repeated data stream.

5.3 Symbol coding (differential Manchester 1, 0, J, K)

Symbols shall be output to the medium in the form of differential Manchester code which is characterized by the transmission of two signal elements per symbol as shown in Figure 27. Symbols are output at the data bit rate of 4 or 16 Mbit/s as described in 5.2.

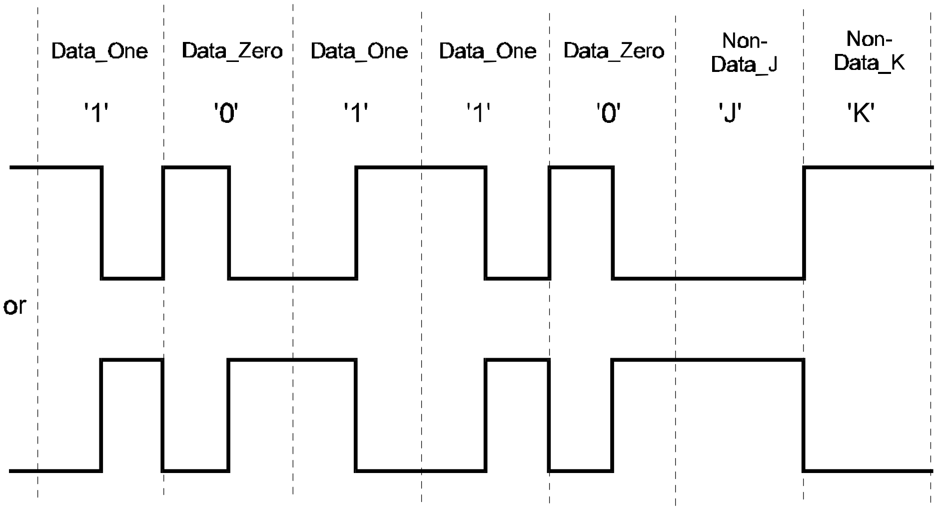

Figure 27—Example of symbol encoding

In the case of two symbols, Data_one and Data_zero, a signal element of one polarity is output for one half the duration of the symbol being transmitted, followed by the contiguous output of a signal element of the opposite polarity for the remainder of the symbol duration. This coding scheme provides the following distinct advantages:

a) The resulting signal has no DC component and can readily be inductively or capacitively coupled.
b) The forced mid-symbol transition conveys inherent timing information on the channel.
c) The signals are independent of channel polarity reversals.

The value of the four differential Manchester symbols depends on the polarity of the trailing signal element of the preceding symbol. The symbol values have the following characteristics.

a) Data_zero symbol: The polarity of the leading signal element is the opposite polarity to the trailing signal element of the preceding symbol. Consequently, a transition occurs at the symbol boundary as well as at mid-symbol.

b) Data_one symbol: The polarity of the leading signal element is the same polarity as the trailing signal element of the preceding symbol. Consequently, only one transition occurs (at mid-symbol).

c) Non-data_J symbol: Both signal elements are the same polarity and they are the same polarity as the trailing signal element of the preceding symbol.

d) Non-data_K symbol: Both signal elements are the same polarity and they are the opposite polarity to the trailing signal element of the preceding symbol.

Although the Non-data_J and Non-data_K symbols are not true differential Manchester symbols (no mid-symbol transition), they are referred to as differential Manchester symbols. The Non-data_J and Non-data_K symbols are used by the starting and ending delimiters (specified in 3.2.1 and 3.2.8) to define the beginning and end of tokens and frames. Since the transmission of single Non-data_J or Non-data_K symbols introduces a DC component on the ring, they are normally transmitted as a pair of J-K symbols to minimize the accumulation of a DC signal component (the combination of a Non-data_J followed immediately by a Non-data_K has no DC component).

5.4 Repeat path

5.4.1 Repeat/fill

When the value for Repeat_mode indicated by the MAC is “repeat,” the station shall repeat the received data such that each received transition (reversal of signal element polarity from one signal element to the next) results in a corresponding transition in the transmitted data.

When the value for Repeat_mode indicated by the MAC is “fill,” that station shall source any combination of Data_zero and Data_one symbols (Data_zero symbols preferred) and not repeat the received data.

5.4.2 Burst error detection, correction, and indication

The PSC indicates a Burst4_error to the MAC when it detects four signal elements with the same polarity in succession. The PSC indicates a Burst5_error to the MAC when it detects five or more signal elements with the same polarity in succession, and, if the station is repeating, the station shall introduce a change of polarity (that is, a transition) for every transmitted signal element starting at the end of the fourth signal element in the sequence until a transition is received from the ring, at which time the station returns to repeating. This process shall assure that no more than 4 signal elements with the same polarity are ever repeated by the station. Note that when a Burst5_error condition is detected, the PSC indicates a Burst4_error and one signal element later indicates a Burst5_error.

5.5 Symbol encoder

The symbol encoder allows the MAC to modify individual symbols in the repeated frames and tokens as well as generate frames and tokens.

The symbol encoder shall encode for transmission, the symbols presented to it by the MAC (Tx_symbol). The values for the symbols exchanged between MAC and symbol encoder are shown as follows. (Specific implementations are not constrained in the method of making this information available.)

Data_zero
Data_one
Non-data_J
Non-data_K

5.6 Symbol decoder

Received symbols (Rcv_symbol) shall be decoded using an algorithm that is the inverse of the one described for symbol encoding. The station shall monitor the received data for the signal element pattern corresponding to a starting delimiter and ending delimiter as shown in figure 28.

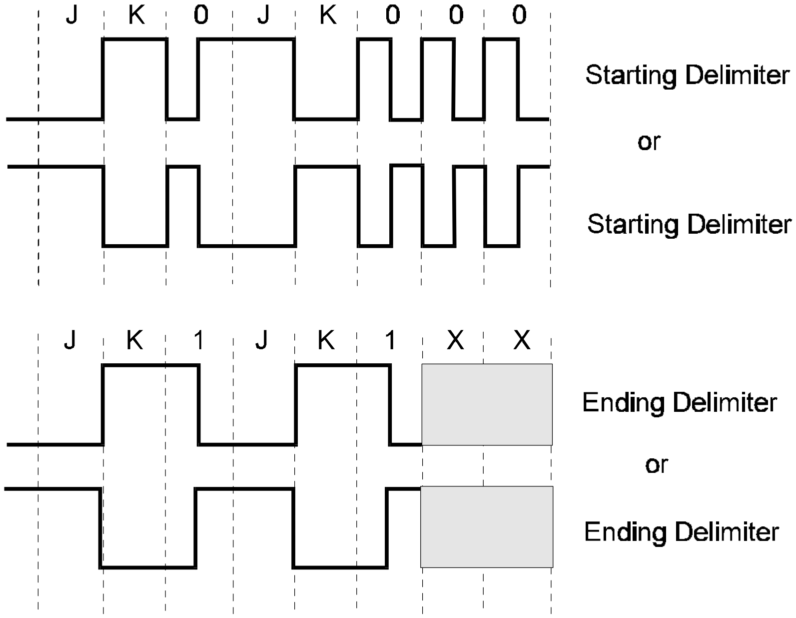

Figure 28—Starting and ending delimiter patterns

Starting delimiters (SD) and ending delimiters (ED) shall be detected even if they do not occur on a symbol boundary. If an ending delimiter is detected that does not occur on the symbol boundary established by the previous SD, then an ED_alignment_error is indicated to the MAC (the MAC uses this signal to verify that the ending delimiter falls exactly on an octet boundary). Once a delimiter has been detected, it shall establish the symbol and octet boundary for the decoding of the received signal elements until another delimiter is detected.

5.7 Signal acquisition/loss and timing synchronization

5.7.1 Signal acquisition/loss

When the signal at the receiver input meets the requirements for amplitude, frequency, jitter, etc., as specified in clause 7, the receiver's clock recovery circuit shall synchronize with the received signal such that no signal elements are added or dropped from the repeated data. The station indicates signal acquisition to the MAC interface as follows:

a) The station shall, upon receipt of a signal that is within specification, acquire phase synchronization within 1.5 ms and indicate Signal_acquired to the MAC.

b) The PHY should indicate Signal_loss to the MAC interface when the receiver is not able to provide accurate data due to lack of input signal quality including loss of timing synchronization.

5.7.2 Frequency error

A difference in frequency between the received data and the crystal clock may cause an overflow or underflow of the elastic buffer. This overflow or underflow is optionally indicated as Frequency_error to the MAC as follows: 1) When operating in the Crystal_transmit mode as described in 5.8.3 or 2) when not in the Crystal_transmit mode, by monitoring the rate at which the elastic buffer (which is not in the data path) underflows or overflows. If this or another method is used to detect frequency error, Frequency_error shall not be indicated if

a) The input jitter is below the value B specified in 5.8.3, and
b) The frequency error over 1 ms is less than 1.76% at 16 Mbit/s or 0.7% at 4 Mbit/s.

Frequency_error, if implemented, shall be indicated if the received bit rate, averaged over 32 ms, deviates more than ±10% from the data signaling rate specified in 5.2.

5.8 Latency

Latency is the time, expressed in number of symbols, it takes for a signal element to pass through a ring component. This clause specifies two latency buffers, a fixed latency buffer and an elastic buffer, that are inserted in the ring signal path as requested by the MAC protocol. To ensure stations do not introduce excessive latency, it is recommended that the average PHY latency should not exceed 6.25 ms to prevent expiration of TRR (3.4.2.13). It is further recommended that this time should not exceed 2.5 ms for performance/access reasons.

5.8.1 Ring latency

Ring latency, defined as the time it takes for a signal element to travel around the entire ring, is equal to the cumulative latency of the ring plus the latency of the active monitor. Cumulative latency is the time it takes for a signal element to travel from the active monitor’s transmitter output to its receiver input. The latency of the active monitor consists of two buffers, elastic buffer and fixed latency buffer, in addition to the normal station latency. The active monitor station uses the elastic buffer to compensate for variations in cumulative latency and uses the fixed latency buffer to provide the required latency to ensure ring operation even when only one station is on the ring.

When the ring is in a normal operating state, the MAC protocol ensures that only one station, the active monitor, will have Crystal_transmit asserted, which inserts the fixed latency and elastic latency buffers in the ring path. The active monitor transmits data to the ring timed from its local clock, which provides the master timing for the ring. Although the mean data signaling rate around the ring is established by the active monitor, segments of the ring can, for short periods of time, operate at a slightly higher or lower signaling rate than the signaling rate of the active monitor. The cumulative effect of these signaling rate variations may cause variations in the cumulative latency.

The elastic buffer in the active monitor compensates for the variation in the cumulative latency as follows. When the frequency of the received signal at the active monitor is slightly faster than the crystal clock, the elastic buffer will expand, as required, to maintain a constant ring latency. Conversely, if the frequency of the received signal is slightly slower, the elastic buffer will contract to maintain a constant ring latency. Constant ring latency is a requirement to avoid adding or dropping signal elements from the data stream. The elastic buffer compensates for dynamic variations in latency due to jitter. The elastic buffer does not compensate for some changes such as step changes in latency caused by stations becoming or ceasing to be stacking stations.

The fixed latency buffer is provided by the active monitor to provide the latency required for token circulation as described in 5.8.2.

5.8.2 Fixed latency buffer

In order for the token to continuously circulate around the ring, the ring must have a minimum ring latency of at least 24 symbols (the number of symbols in a token). To ensure this minimum ring latency, a fixed latency buffer of at least 24 symbols shall be inserted into the data path of the station when MAC asserts Crystal_transmit.

5.8.3 Elastic buffer

An elastic buffer shall be present in the station's data path when the MAC asserts Crystal_transmit. This elastic buffer latency is in addition to the latency provided by the fixed latency buffer.

5.8.3.1 Operation of the elastic buffer

Cumulative latency variations are based on the jitter considerations discussed in annex C. The elastic buffer shall accommodate without error a minimum cumulative latency variation of positive or negative B symbols where

Data rate	=	4	16	Mbit/s
B	=	3	15	symbols

The elastic buffer in the active monitor shall be initialized to the center of the elastic latency range as described in 5.8.3.2.

The active monitor's elastic buffer shall also be reinitialized to the center of the elastic latency range as described in 5.8.3.2 whenever the MAC indicates Token_received (station is repeating a token).

If the elastic buffer exceeds full elasticity (underflow/overflow), it shall be reinitialized as described in 5.8.3.2. In addition, the signal Frequency_error may optionally be indicated as described in 5.7.

5.8.3.2 Centering of the elastic buffer

The length of the elastic buffer is initialized (reinitialized) to at least the minimum allowable elastic latency (B symbols) and no more than the actual total elastic latency minus the minimum allowable elastic latency (i.e., with a maximum elastic buffer size x then after centering, the latency will be a minimum of B and a maximum of $x–B$). After the elastic buffer is centered, the station shall accommodate without error a minimum cumulative latency variation of positive or negative B symbols as specified in 5.8.3.1.

The centering process may cause signal elements to be added or removed from the repeated data stream. In addition, the reinitializing of the elastic buffer may cause an inversion of the repeated data produced by the adding or removing of an odd number of signal elements. The total combined effect of reinitializing the elastic buffer shall not introduce more than a single isolated Non-data_J or Non-data_K symbol into the data stream. This symbol shall not immediately precede a starting delimiter. During re-initialization, no portion of the token shall be altered. The active monitor station may alter signal elements following the token (recommended) or signal elements preceding the token (i.e., the interframe gap).

5.8.3.3 Modification of the interframe gap

Although the centering process may cause signal elements to be added or removed from the interframe gap, the length of the interframe gap has been specified to permit the centering of the active monitor's elastic buffer with an elastic buffer length of 2B symbols as specified in 5.8.3.1. When the interframe gap is equal to or larger than the length specified in 3.2.10, the active monitor station reinitializing its elastic buffer shall only alter those symbols within the interframe gap or the fill following the token. Except for the active monitor, a station shall not cause the interframe gap to be decreased to less than the minimum specified in 3.2.10. When the interframe gap is equal to or larger than the length specified in 3.2.10 and the cumulative latency variation is equal to or less than *B* symbols specified in 5.8.3.1, then the active monitor station shall not reduce the length of the interframe gap to less than *C* symbols as specified:

Data rate	=	4	16	Mbit/s
C	=	5	25	symbols

If the active monitor station alters an interframe gap, it shall only alter the interframe gap preceding a token and may alter all of the signal elements in the interframe gap provided that the altered signal elements result in only Data_zero and Data_one symbols.

A station is not required to receive a frame preceded by an interframe gap with a length less than 8 symbols at 4 Mbits/s or 24 symbols at 16 Mbits/s.

5.9 Ring access control

The Ring access control (RAC) receives two signals from the MAC to control the physical insertion of the station into the ring. As defined in 7.2.1.1, when the MAC indicates Insert_station, the RAC signals “ring insertion” to the concentrator’s TCU. Normally, this action will attach the upstream station’s transmitter output to this station’s receiver input and also attach this station’s transmitter output to the downstream station’s receiver input. Conversely, when the MAC indicates Remove_station, the station’s RAC signals “ring bypass” to the TCU. The TCU is expected to connect the upstream station’s transmitter output to the downstream station’s receiver input. The station’s transmitter output will also be connected to the station’s receiver input utilizing as much of the physical medium as possible from the station to the TCU and still isolate the station from the ring. This “wrapping” of the station’s output to its input allows the station to perform a comprehensive self test verifying as much of the lobe data path and station’s components as possible.

The Wire_fault attribute is provided by the PMC to indicate to the MAC that the station has or has not detected a wire fault condition specified in 7.2.1.2. This signal generally indicates whether the station is or is not properly connected to the TCU. Detection of Wire_fault by the PMC varies with the medium used to attach the station to the ring. The MAC protocol specifies the actions taken (i.e., if Wire_fault is detected, the station eventually removes from the ring as specified in table 7).

6. Token ring management

This clause specifies the primitives necessary to manage operation of a token ring station. This clause also defines the managed objects specific to a token ring station.

6.1 Station management primitives

This clause specifies the primitives necessary to manage the operation of the MAC and PHY as a station and also to support additional management functions. Specification of these primitives do not specify any particular implementation but rather are supplied solely to define the management of the MAC and PHY entities.

The support of the parameters, events, and actions accessible through the primitives is necessary only to support management applications. The only required primitives are the MGT_ACTION.request(acOpen) and the MGT_ACTION.request(acClose). The remaining primitives and their attributes are defined for interoperability with other standards and thus are only required as specified by other applications.

6.1.1 Token ring actions

MGT_ACTION.request { action
parameters }

MGT_ACTION.response { action
status }

The action may be acOpen or acClose, which are defined as follows:

- **acOpen:** Management instructs the MAC to join the station to the ring. This action issues the Connect.MAC event to the MAC protocol. Upon completion (success or failure) of the join ring process, an Open_status is returned to the MGT interface. The parameters passed with this action are the individual MAC address (IndMACAddress) and the ring rate (RingRate). If the MAC address is not specified, then the MAC protocol will use the MAC address assigned by the manufacturer (a universally administered). The MAC address parameter provides the means for management to set a locally administered address. If the specified address (1) is not an individual locally administered address or (2) if it is not specified and no universally administered address is assigned, then the status returned in the MGT_ACTION.response is Invalid_address. If the requested data rate is not supported, then the status returned is Rate_not_supported.

  The MGT_ACTION.response is generated when the station has successfully joined the ring or when the join process has failed. The returned status is defined as follows:

Table 8—MGT_ACTION.response definition

Open status	Table 7 REF	Definition
sNotOperational	n/a	Rejected—device not ready, internal error, failure, etc.
sInvalidAddress	n/a	Rejected—MAC address is invalid (group or not specified)
sRateNotSupported	n/a	Rejected—the requested data rate is not supported
sSuccess	100, 131, 188	Successful
sFailGeneral	n/a	Failure in activation—general
sFailLobeTest	276	Failure in activation—lobe test failure
sFailDataRate	238	Failure in activation—medium rate error
sFailDuplicate	127	Failure in activation—duplicate address
sFailDATTimeout	282	Failure in activation—DAT time-out—TJR expired
sFailNNTimeout	284	Failure in activation—neighbor notification time-out
sFailInitTimeout	286, 317	Failure in activation—initialization time-out
sFailRingBeaconing	067, 068, 069, 070, 071	Failure in activation—beacon received
sFailPurgeTimeout	318, 319, 320, 321	Failure in activation—purge failure—TRP expired
sFailClaimTimeout	265, 269, 270	Failure in activation—claim failure—TCT expired
sFailBeaconTimeout	263	Failure in activation—beacon failure—TBT expired
sFailSignalLoss	209, 210, 211	Failure in activation—signal loss
sFailInternal	224, 225, 226, 228, 229	Failure in activation—internal error
sFailTxFault	255, 256, 258, 259	Failure in activation—station transmit failure
sRemove	156, 157, 158, 160, 161	Failure in activation—remove frame received

— **acClose:** Management instructs the MAC protocol to remove the station from the ring. There are no parameters associated with this action. This action issues the Disconnect.MAC event to the MAC protocol. When the station has completed removing from the ring (i.e., transitions to bypass (JS=BP) as per table 7 REF 325), the MGT_ACTION.response is indicated with the status of Complete.

6.1.2 Token ring events

MGT_EVENT.indication { event,
event_status }

An event report is an information unit that is used to inform MGT of a significant activity in the operation of the resource. A description of the events defined for token ring follows the list of events:

evRingNonOperational
evRingOperational
evRingBeaconing

evStationFailure
evProtocolError

6.1.2.1 The **evRingNonOperational** event is generated when the monitor FSM transitions into recovery (MS=TCT, RCT, TBN, RBN). The Event_status is defined as follows:

Table 9—evRingNonOperational Event_status

Event_status	Table 7 REF	Definition
sRingNonOp	n/a	Ring not operational—cause not specified
sPurgeFailure	322	Claim—purge failure
sRingClaiming	108, 115, 116, 117, 118	Claim—claim received
sClaimNoMonitor	026, 310, 336, 337,	Claim—no active monitor (TJR, TNT, TSM=E)
sClaimSignalLoss	214, 215	Claim—signal loss
sRingBeaconing	075, 078	Beacon received

6.1.2.2 The **evRingOperational** event is generated when the ring transitions into normal repeat (MS=RPT) as per table 7 REF 169, 175, 177. The Event Status is defined as follows:

Table 10—evRingOperational Event_status

Status	Table 7 REF	Definition
sRingNonOp	175, 177, 169	Ring operational—monitor state not specified
sStandbyMonitor	175, 169	This station is a standby monitor (FAM=0)
sActiveMonitor	177	This station is the active monitor (FAM=1)

6.1.2.3 The **evRingBeaconing** event is generated when the monitor FSM transitions into a beacon state (MS=TBN, RBN) as per table 7 REF 076, 078, 073, 075, 273, 271, 272, 266. The Event_status is defined as the SA, beacon type, reported UNA, and reported physical drop number from the received beacon frame if MS=RBN or the transmitted beacon frame if MS=TBN.

6.1.2.4 The **evStationFailure** event is generated when a station fault causes the station to be removed from the ring. The Event_status is defined as follows:

Table 11—evStationFailure Event_status

Status	Table 7 REF	Definition
sRingNonOp	n/a	Station removed—cause not specified
sRemove	161	Remove frame received
sWireFault	345	Wire fault
sFailInternal	227, 257	Internal failure
sFailTest	278	Beacon test failure

6.1.2.5 The **evProtocolError** event is generated when the MAC protocol detects an error. The event status is defined as either sActiveMonitorError or sNotificationIncomplete.

Table 12—evProtocolError Event_status

Status	Table 7 REF	Definition
sActiveMonitorError	061, 104, 171, 322	Active monitor error
sNotificationIncomplete	260	Neighbor notification incomplete

The status for the active monitor error includes the error code (EC) and the status for the notification incomplete includes the address of the last station to participate (LMP).

6.1.3 Attributes

MGT_GET.request { attributeID }

MGT_REPLACE.request { attributeID,
newValue }

MGT_GET.request: The MGT_GET is issued to request information (attributes) about the token ring station. The effect of the request is that the requested attribute is returned. There is only one parameter associated with the Get request and it indicates the attributes group to be supplied in the response.

MGT_REPLACE.request: The MGT_REPLACE is issued to change certain characteristics (attributes) that effect the operation of the token ring station. The parameters indicate the attribute to change and the desired value of the attribute. The effect of the request is that the attribute is modified and a response is generated containing the value of the attribute after modification. Note that some values of attributes listed as GET-REPLACE may not be supported by the station and therefore may not change.

The attributes of a resource indicate its state (present or past) and control its operation (in the future). The Token Ring station attributes are specified in groups. Attributes are classified as follows:

a) *Characteristics.* Operational information that describes some aspect of the resource’s capabilities. In general, characteristics affect the operation of the resource at some future time. Characteristics may be specifically defined to be GET (read-only) or GET-REPLACE (read-write) with respect to remote management access.

b) *Status.* Dynamic information about the resource’s present state. Status attributes are read-only.

c) *Statistics.* Information about the resource’s past behavior. Statistical attributes are typically a form of an event log, such as counters. The only type of statistics defined for this standard are counters that are read-only with no reset control.

6.1.3.1 Characteristics

Characteristics are the reference data of a resource that may be either necessary or useful to operate or manage the resource.

6.1.3.1.1 Station characteristics

These attributes specify the operating characteristics of the station in general. This group represents attributes that are characteristic of network operation and some of these parameters may optionally be set by management while others can not be altered by management. The characteristics of the station are

IndMACAddress—This parameter specifies the address used in the SA field of the MAC frames originated by this station. This attribute may only be set by the Open command and thus cannot be changed while the MAC is active.

UniversallyAdministeredAddress—This parameter specifies the universally administered address (if any) assigned to this station. This value can not be set by management.

FunctionalAddresses—This specifies the set of functional addresses active in the station (i.e., MAC will set the A bits and copy the frame). This parameter provides the means for management to activate functional addresses.

GroupAddresses—This parameter specifies a list of group addresses recognized by the station (i.e., MAC will set the A bits and copy the frame). This parameter provides the means for management to establish group addresses.

MicrocodeLevel—This parameter specifies the value of the X'23' Ring Station Version Number subvector (3.3.4) used in the Report Station State MAC frame. This value can not be set by management.

ProductInstanceId—This parameter specifies the value of the X'22' Product Instance ID subvector (3.3.4) used in the Report Station Attachment and Report New Active Monitor MAC frames. This value can not be set by management.

MaxFrameSize—The value of this parameter specifies the size, in octets, of the largest frame that may be transmitted and received by the station.

6.1.3.1.2 Protocol configuration options

These attributes specify the operating characteristics of the MAC protocol (refer to clause 3.5). This group represents attributes that may optionally be set by management.

Burst4Option—This parameter specifies the value of the flag FBEO indicating whether the station compensates for burst4 errors when stripping frames. The value is either Compensates or No_compensation.

BeaconHandling—This parameter specifies the value of the flag FBHO indicating whether the station waits until JS=JC before fully participating ring recovery. The value is either Before_initialization or After_initialization

ClaimContender—This parameter specifies the value of the flag FCCO indicating whether the station always actively participates in the claiming process. The value is either Contender or No_contender.

EarlyTokenRelease—This parameter specifies the value of the flag FETO indicating if early token release is selected. The value is either Enabled or Not_enabled.

RingRate—This parameter specifies the value of the station’s operating speed use to set the flag FMRO. The value is either 4_Mbit/s or 16_Mbit/s. This attribute may only be set by the Open command and thus cannot be changed while the MAC is active.

MultipleFrameTransmission—This parameter specifies the value of the flag FMFTO indicating whether the station honors the reservation field when transmitting multiple frames. The value is either Relinquish or Continue.

RejectRemove—This parameter specifies the value of the FRRO indicating whether the station will honor the Remove MAC frame. The value is either Removes or Rejects.

TokenErrorDetection—This parameter specifies the value of the FTEO indicating whether the station counts errors in Tokens. The value is either Counted or Not_Counted

TokenHandling—This parameter specifies the value of the FTHO indicating how the station detects lost tokens. The value is either Active_timer or Periodic_inspection.

ErrorCountingMethod—This parameter specifies the value of the FECO indicating how the station manages the error report timer. The value is either Triggered or Free_running.

6.1.3.2 Status

A status attribute is one that indicates something about the current state of the resource. A status attribute is distinguished from a characteristic in that it is modified internally by the resource rather than by an external management entity. Status attributes are read-only.

6.1.3.2.1 Station operating status

NAUN—This parameter specifies the MAC address and physical drop number of this station's nearest upstream neighbor. The parameter values are derived from the AMP or SMP frame (table 7 REF 064, 199, 200).

RingNumber—This parameter specifies the value set by the X'03' Local Ring Number subvector (3.3.4) from the Change Parameters or Initialize Station MAC frame.

AccessPriority—This parameter specifies the value set by the X'07' Authorize Access Priority subvector (3.3.4) from the Change Parameters MAC frame.

ErrorReportTimer—This parameter specifies the value of the timer TER as set by the X'05' Error Timer Value subvector (3.3.4) from the Change Parameters or the Initialize Station MAC frame.

AuthorizeFunctionClasses—This parameter specifies the value set by the X'06' Authorize Function Classes subvector (3.3.4) of the Change Parameters MAC frame.

PhysicalDropNumber—This parameter specifies the value set by the X'04' Assign Physical Drop Number subvector (3.3.4) of the Change Parameters or the Initialize Station MAC frame.

RingStationStatus—This parameter specifies the value of the X'29' Ring Station Status subvector (3.3.4) used in the Report Station State MAC frame.

6.1.3.2.2 MAC Protocol Status

JoinState—This parameter specifies the present state of the Join FSM. The value will be one of the following: NotSpecified, Bypass, LobeTest, DetectMonitorPresent, AwaitNewMonitor, DuplicateAddressTest, NeighborNotification, RequestInitialization, JoinComplete, or BypassWait.

MonitorState—This parameter specifies the present state of the Monitor FSM. The value will be one of the following: NotSpecified, StandbyMonitor (MS=RPT & FAM=0), ActiveMonitor (MS=RPT & FAM=1), RingPurge (MS=TRP), ClaimTransmit (MS=TCT), ClaimRepeat (MS=RCT), BeaconTransmit (MS=TBN), BeaconRepeat (MS=RBN), or BeaconTest (MS=BNT).

BeaconInfo—These parameters specify the MAC address of the beaconing station, the beacon type, the MAC address of the beaconing station's upstream neighbor, and the physical drop number of the beaconing station. These parameters are the value of the SA, the X'01' Beacon Type subvector, the X'02' UNA subvector, and the X'0B' Physical Drop Number subvector in the last Beacon MAC frame transmitted or received.

TxBeaconType—This parameter specifies the value of the X'01' Beacon Type subvector in the last Beacon MAC frame transmitted.

6.1.3.3 Statistics

Statistics are attributes that contain a record of events over some period of time. The statistics defined for this standard are counters with no reset control. Access to the counterValue parameter is read-only. Refer to 3.6 for definition of each of the error counters.

6.1.3.3.1 Isolating error counters group

The isolating error counters group lists the counters and their definitions. These errors are those that can be isolated to a particular fault domain (this station, its upstream neighbor, and the wire between them). The following indicates the parameters reported and the FSM transition that causes the counter to increment:

LineErrorCounter—This counter is incremented each time CLE is incremented (table 7 REF 204, 205, 308, 309).

BurstErrorCounter—This counter is incremented each time CBE is incremented (table 7 REF 001, 002).

AcErrorCounter—This counter is incremented each time CACE is incremented (table 7 REF 194, 195, 196, 197).

AbortErrorCounter—This counter is incremented each time CABE is incremented (table 7 REF 252, 253, 302, 303).

InternalErrorCounter—This counter is incremented each time CIE is incremented (table 7 REF 222, 223).

6.1.3.3.2 Non-isolating error counters group

The non-isolating error counters group lists the counters and their definitions. These errors are those that cannot be isolated to a particular fault domain. The following indicates the parameters reported and the FSM transition that causes the counter to increment:

LostFrameErrorCounter—This counter is incremented each time CLFE is incremented (table 7 REF 031, 032).

ReceiveCongestionErrorCounter—This counter is incremented each time CRCE is incremented (table 7 REF 152, 153).

FrameCopiedErrorCounter—This counter is incremented each time CFCE is incremented (table 7 REF 041, 042).

FrequencyErrorCounter—This counter is incremented each time CFE is incremented (table 7 REF 147, 148)

TokenErrorCounter—This counter is incremented each time CTE is incremented (table 7 REF 052, 053, 291, 292, 339, 340, 342, 343).

6.2 Token ring station managed objects

This clause contains the managed objects for the token ring station. These managed objects are specified in conformance with ISO 10742 : 1994. The service and protocol elements that are involved in the management of token ring stations are defined in ISO/IEC 15802-2 : 1995.

A managed object is defined as the smallest accessible entity that contains managed attributes, actions. and event notifications that are independent and unrelated to duplicate copies of the same attributes, actions, and event notifications that exist in another managed entity.

The MAC sublayer is defined by the TokenRingLayer2MAC managed object.

6.2.1 TokenRingLayer2MAC MANAGED OBJECT CLASS

```
oTokenRingLayer2MAC MANAGED OBJECT CLASS
    DERIVED FROM “10165-2 : 1992”:top;mACDLE;
    CHARACTERIZED BY
        pTokenRingMAC;
    CONDITIONAL PACKAGES
        pTrStationCharacteristics  PRESENT IF supported,
        pTokenRingMACOptions  PRESENT IF supported,
        pTrCharacteristics  PRESENT IF supported,
        pTokenRingStatus  PRESENT IF supported,
        pTokenRingStatistics  PRESENT IF supported;
```

```
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) managedobjectclass(3) tokenringlayer2mac(0)};
```

6.2.2 TokenRingLayer2MAC NAME BINDING

```
nTokenRingLayer2MAC NAME BINDING
    SUBORDINATE OBJECT CLASS TokenRingLayer2MAC;
    NAMED BY
    SUPERIOR OBJECT CLASS macDLE;
    WITH ATTRIBUTE ID UniversallyAdministratedAddress
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) namebinding(6)
        tokenringlayer(0)};
```

6.2.3 TokenRingLayer2MAC PACKAGES

6.2.3.1 TokenRingMAC PACKAGE

```
pTokenRingMAC  PACKAGE
    BEHAVIOR  bTokenRingLayer2MAC  BEHAVIOR
    DEFINED AS !This package identifies the various MAC addresses related to the station and provides
        the general management operation.!;
    ATTRIBUTES
                         aIndMACAddress  GET,                 -- 0
                       aUniversalAddress  GET,                -- 1
                    aFunctionalAddresses  GET-REPLACE,        -- 6
                         aGroupAddresses  GET-REPLACE,        -- 7
                                   aNAUN  GET;                -- 22
    ATTRIBUTE GROUPS
                   "gTrStationAddresses"                      -- 0
    ACTIONS
                                 acOpen,                      -- 0
                                acClose;                      -- 1
    NOTIFICATIONS
                       evStationFailure;                      -- 3
    REGISTERED AS {iso(1) memberbody(2) us(840) 802dot5(10033) package(4)
        tokenringmac(0)};
```

6.2.3.2 TrStationCharacteristics PACKAGE

```
pTrStationCharacteristics  PACKAGE
    BEHAVIOR   bTokenRingLayer2MAC   BEHAVIOR
    DEFINED AS !This package identifies the characteristics of the ring station.!;
    ATTRIBUTES
                       aMicrocodeLevel  GET,                  -- 2
                    aProductInstanceId  GET,                  -- 3
                         aMaxFrameSize  GET;                  -- 4
    ATTRIBUTE GROUPS
           "gTrStationCharacteristics"                        -- 1
    REGISTERED AS {iso(1) memberbody(2) us(840) 802dot5(10033) package(4)
        trstationcharacteristics(1)};
```

6.2.3.3 TokenRingMACOptions PACKAGE

```
pTokenRingMACOptions  PACKAGE
    BEHAVIOR   bTokenRingLayer2MAC   BEHAVIOR
    DEFINED AS !This package identifies the protocol configuration of the ring station.!;
    ATTRIBUTES
                                    aBurst4Option   GET-REPLACE,          -- 8
                                  aBeaconHandling   GET-REPLACE,          -- 9
                                 aClaimContender   GET-REPLACE,          -- 10
                               aEarlyTokenRelease   GET-REPLACE,          -- 11
                                        aRingRate   GET,                  -- 5
                       aMultipleFrameTransmission   GET-REPLACE,          -- 12
                                    aRejectRemove   GET-REPLACE,          -- 13
                              aTokenErrorDetection   GET-REPLACE,          -- 14
                                   aTokenHandling   GET-REPLACE,          -- 15
                             aErrorCountingMethod   GET-REPLACE;          -- 16
    ATTRIBUTE GROUPS
                              "gTrProtocolOptions"                         -- 2

    REGISTERED AS {iso(1) memberbody(2) us(840) 802dot5(10033) package(4)
        tokenringmacoptions(2)};
```

6.2.3.4 TrCharacteristics PACKAGE

```
pTrCharacteristics  PACKAGE
    BEHAVIOR   bTokenRingLayer2MAC   BEHAVIOR
    DEFINED AS !This package identifies the ring characteristics for the station.!;
    ATTRIBUTES
                                      aRingNumber   GET{-REPLACE},        -- 17
                                 aErrorReportTimer   GET{-REPLACE},        -- 18
                           aAuthorizeFunctionClass   GET{-REPLACE},        -- 19
                                   aAccessPriority   GET{-REPLACE},        -- 20
                               aPhysicalDropNumber   GET{-REPLACE},        -- 21
                                aRingStationStatus   GET;                  -- 23
    ATTRIBUTE GROUPS
                              "gTrCharacteristics"                         -- 3
    REGISTERED AS {iso(1) memberbody(2) us(840) 802dot5(10033) package(4)
        trcharacteristics(3)};
```

6.2.3.5 TokenRingStatus PACKAGE

```
pTokenRingStatus  PACKAGE
    BEHAVIOR   bTokenRingLayer2MAC   BEHAVIOR
    DEFINED AS !This package identifies the protocol state of the ring station.!;
    ATTRIBUTES
                                       aJoinState   GET,                  -- 24
                                    aMonitorState   GET,                  -- 25
                                     aBeaconInfo,   GET,                  -- 26
                                   aTxBeaconType,   GET;                  -- 27
```

```
    ATTRIBUTE GROUPS
                                  gTrMACStatus                        -- 4

    NOTIFICATIONS
                           evRingNonOperational,                      -- 0
                              evRingOperational,                      -- 1
                               evRingBeaconing,                       -- 2
                                evProtocolError;                      -- 4
    REGISTERED AS {iso(1) memberbody(2) us(840) 802dot5(10033) package(4)
        tokenringstatus(4)};
```

6.2.3.6 TokenRingStatistics PACKAGE

```
pTokenRingStatistics  PACKAGE
    BEHAVIOR   bTokenRingLayer2MAC   BEHAVIOR
    DEFINED AS !This package provides the error statistics of the ring station.!;
    ATTRIBUTES
                            aLineErrorCounter   GET,                  -- 28
                           aBurstErrorCounter   GET,                  -- 29
                              aAcErrorCounter   GET,                  -- 30
                       aAbortTransmitCounter   GET,                  -- 31
                        aInternalErrorCounter   GET,                  -- 32
                       aLostFrameErrorCounter   GET,                  -- 33
                  aRcvCongestionErrorCounter   GET,                  -- 34
                    aFrameCopiedErrorCounter   GET,                  -- 35
                      aFrequencydErrorCounter   GET,                  -- 36
                           aTokenErrorCounter   GET;                  -- 37
    ATTRIBUTE GROUPS
                              gTrIsolatingErrors                      -- 5
                           gTrNonIsolatingErrors                      -- 6
    REGISTERED AS {iso(1) memberbody(2) us(840) 802dot5(10033) package(4)
        tokenringstatistics(5)};
```

6.2.4 TokenRingLayer2MAC BEHAVIOR

```
bTokenRingLayer2MAC  BEHAVIOR
    DEFINED AS !This object class defines the MAC sublayer for token ring.  The operation of the
        Token Ring MAC is defined in the ISO/IEC 8802-5 : 1995 Standard (802dot5).!;
```

6.2.5 Attributes

6.2.5.1 Individual MAC Address ATTRIBUTE

```
aIndMACAddress ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.MACAddress;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bIndMACAddress  BEHAVIOR
    DEFINED AS !Specifies the MAC address of the token ring station.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        indmacaddress (0)};
```

6.2.5.2 Universal Address ATTRIBUTE

```
aUniversalAddress ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.MACAddress;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bUniversalAddress BEHAVIOR
    DEFINED AS !Specifies the universal address assigned to the device.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        universaladdress(1)};
```

6.2.5.3 Microcode Level ATTRIBUTE

```
aMicrocodeLevel  ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.MicrocodeLevel;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bMicrocodeLevel  BEHAVIOR
    DEFINED AS !Specifies the microcode level of this device.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        microcodelevel  (2)};
```

6.2.5.4 Product ID ATTRIBUTE

```
aProductInstanceId  ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.ProductInstanceId;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bProductInstanceId  BEHAVIOR
    DEFINED AS !Identifies Product type of the Token Ring device.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        productinstanceid(3)};
```

6.2.5.5 Maximum Frame Size ATTRIBUTE

```
aMaxFrameSize  ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.MaxFrameSize;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bMaxFrameSize BEHAVIOR
    DEFINED AS !Specifies the maximum frame size supported by the device.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        maxframesize(4)};
```

6.2.5.6 Ring Media Rate ATTRIBUTE

```
aRingRate  ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.RingRate;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bRingRate  BEHAVIOR
    DEFINED AS !Specifies the current data transmission rate of this device.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7) ringrate(5)};
```

6.2.5.7 Functional Addresses ATTRIBUTE

```
aFunctionalAddresses ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.FunctionalAddresses;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bFunctionalAddresses BEHAVIOR
    DEFINED AS !Specifies the active functional addresses of the token ring station.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        functionaladdresses(6)};
```

6.2.5.8 Group Addresses ATTRIBUTE

```
aGroupAddresses ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.GroupAddresses;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bGroupAddresses BEHAVIOR
    DEFINED AS !Specifies the group addresses recognized by the token ring station.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        groupaddresses(7)};
```

6.2.5.9 Burst4 Option ATTRIBUTE

```
aBurst4Option ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.Burst4Option;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bBurst4Option BEHAVIOR
    DEFINED AS !Specifies whether the station compensates for burst error when stripping.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        burst4option(8)};
```

6.2.5.10 Beacon handling ATTRIBUTE

```
aBeaconHandling ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.BeaconHandling;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bBeaconhandling BEHAVIOR
    DEFINED AS !Indicates when the station will fully participate in beaconing.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        beaconhandling(9)};
```

6.2.5.11 Claim contender ATTRIBUTE

```
aClaimContender ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.ClaimContender;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bClaimContender BEHAVIOR
    DEFINED AS !Indicates whether the station will claim when it is not the station that detected the
        protocol error.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        claimcontender(10)};
```

6.2.5.12 Early token release ATTRIBUTE

```
aEarlyTokenRelease ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.EarlyTokenRelease;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bEarlyTokenRelease BEHAVIOR
    DEFINED AS !Specifies if the station releases the token before waiting for the reception of the
        transmitted frame (early token release option).!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        earlytokenrelease(11)};
```

6.2.5.13 Multiple frame transmission fairness ATTRIBUTE

```
aMultipleFrameTransmission ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.MultipleFrameTransmission;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bMultipleFrameTransmission BEHAVIOR
    DEFINED AS !Indicates whether the station stops transmitting when a higher priority reservation is
        received.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot.5(10033) attribute(7)
        multiframetxopt(12)};
```

6.2.5.14 Reject remove frame ATTRIBUTE

```
aRejectRemove ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.RejectRemove;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bRejectRemove BEHAVIOR
    DEFINED AS !Indicates whether the station will reject the Remove Station MAC frame.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        rejectremove(13)};
```

6.2.5.15 Token error detection ATTRIBUTE

```
aTokenErrorDetection ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.TokenErrorDetection;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bTokenErrorDetection BEHAVIOR
    DEFINED AS !Indicates whether the station counts errors in tokens.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        tokenerrordetect(14)};
```

6.2.5.16 Token handling option ATTRIBUTE

```
aTokenHandling ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.TokenHandling;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bTokenHandling BEHAVIOR
    DEFINED AS !Indicates how the station detects lost tokens.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        tokenhandling(15)};
```

6.2.5.17 Error counting method ATTRIBUTE

```
aErrorCountingMethod ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.ErrorCountingMethod;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bErrorCountingMethod BEHAVIOR
    DEFINED AS !Indicates how the station manages the error report timer.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        errorcounting(16)};
```

6.2.5.18 Ring Number ATTRIBUTE

```
aRingNumber ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.RingNumber;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bRingNumber BEHAVIOR
    DEFINED AS !Indicates the ring number assigned to the device.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        ringnumber(17)};
```

6.2.5.19 Error Report Timer ATTRIBUTE

```
aErrorReportTimer ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.ErrorReportTimer;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bErrorReportTimer BEHAVIOR
    DEFINED AS !Specifies the time-out value for the station’s error report timer.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        errorreporttimer(18)};
```

6.2.5.20 Authorize FunctionClasses ATTRIBUTE

```
aAuthorizeFunctionClasses ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.AuthorizeFunctionClasses;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bAuthorizeFunctionClasses BEHAVIOR
    DEFINED AS !Specifies the function classes allowed for this station.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        authorizefunctionclasses(19)};
```

6.2.5.21 Access Priority ATTRIBUTE

```
aAccessPriority ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.AccessPriority;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bAccessPriority BEHAVIOR
    DEFINED AS !Specifies the maximum priority of a token that may be used for LLC traffic.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        accesspriority(20)};
```

6.2.5.22 Physical Drop Number ATTRIBUTE

```
aPhysicalDropNumber ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.PhysicalDropNumber;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bPhysicalDropNumber BEHAVIOR
    DEFINED AS !Specifies the station’s physical location.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        physicaldropumber(21)};
```

6.2.5.23 Nearest Active Upstream Neighbor ATTRIBUTE

```
aNAUN ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.NAUN;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bNAUN BEHAVIOR
    DEFINED AS !Specifies the individual MAC address of this station’s upstream neighbor.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7) naun(22)};
```

6.2.5.24 Ring Station Status ATTRIBUTE

```
aRingStationStatus ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.RingStationStatus;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bRingStationStatus BEHAVIOR
    DEFINED AS !Indicates the current state of the station’s MAC protocol.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        ringstationstatus(23)};
```

6.2.5.25 Join State ATTRIBUTE

```
aJoinState ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.JoinState;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bJoinState BEHAVIOR
    DEFINED AS !Indicates the current state of the MAC protocol’s Join process.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7) joinstate(24)};
```

6.2.5.26 Monitor State ATTRIBUTE

```
aMonitorState ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.MonitorState;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bMonitorState BEHAVIOR
    DEFINED AS !Indicates the current state of the MAC protocol’s Monitor process.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7) monitorstate(25)};
```

6.2.5.27 Beacon Information ATTRIBUTE

```
aBeaconInfo ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.BeaconInfo;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bBeaconInfo BEHAVIOR
    DEFINED AS !Indicates the fault domain.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7) beaconinfo(26)};
```

6.2.5.28 TxBeacon Type ATTRIBUTE

```
aBeaconType ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.BeaconType;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bTxBeaconType BEHAVIOR
    DEFINED AS !Indicates the reason for beaconing.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7) txbeacontype(27)};
```

6.2.5.29 Line Error Counter ATTRIBUTE

```
aLineErrorCounter ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.EventCounter;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bLineErrorCounter BEHAVIOR
    DEFINED AS !The number of line errors (CLE). This counter is incremented when a ring station
        detects an error within frame or token that has not previously been detected.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        lineerrorcounter(28)};
```

6.2.5.30 Burst Error Counter ATTRIBUTE

```
aBurstErrorCounter ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.EventCounter;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bBurstErrorCounter BEHAVIOR
    DEFINED AS !The number of burst errors (CBE). This counter is incremented when a ring
        station detects the absence of transitions for 5 signal element times.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        bursterrorcounter(29)};
```

6.2.5.31 AC Error Counter ATTRIBUTE

```
aAcErrorCounter ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.EventCounter;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bAcErrorCounter BEHAVIOR
    DEFINED AS !The number of AC errors (CACE). This counter is incremented as an indication that
        this station or its upstream neighbor is not properly setting or recognizing the A&C bits.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        acerrorcounter(30)};
```

6.2.5.32 Abort Transmit Counter ATTRIBUTE

```
aAbortTransmitCounter ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.EventCounter;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bAbortTransmitCounter BEHAVIOR
    DEFINED AS !The number of Abort Transmit errors (CABE). This counter is incremented when this
        station transmits an abort delimiter.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        aborttransmitcounter(31)};
```

6.2.5.33 Internal Error Counter ATTRIBUTE

```
aInternalErrorCounter ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.EventCounter;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bInternalErrorCounter BEHAVIOR
    DEFINED AS !The number of Internal errors (CIE). This counter is incremented when this
        station recognizes a recoverable internal error.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        internalerrorcounter(32)};
```

6.2.5.34 Lost Frame Error Counter ATTRIBUTE

```
aLostFrameErrorCounter ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.EventCounter;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bLostFrameErrorCounter BEHAVIOR
    DEFINED AS !This counter contains the number of Lost Frame errors (CLFE). This counter is
        incremented when this station is transmitting and its return-to-repeat (TRR) timer expires.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        lostframeerrorcounter(33)};
```

6.2.5.35 Receive Congestion Error Counter ATTRIBUTE

```
aRcvCongestionErrorCounter ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.EventCounter;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bRcvCongestionErrorCounter BEHAVIOR
    DEFINED AS !This counter contains the number of Receiver Congestion errors (CRCE). This counter
        is incremented when this station recognizes that the frame being received is addressed to this
        station but has no buffer space available to copy the frame.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        rcvcongestionerrorcounter(34)};
```

6.2.5.36 Frame Copy Error Counter ATTRIBUTE

```
aFrameCopiedErrorCounter ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.EventCounter;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bFrameCopiedErrorCounter BEHAVIOR
    DEFINED AS !This counter contains the number of Frame Copy errors (CFCE). This counter is
        incremented when this ring station receives a MAC frame addressed to this station’s individual
        address, but the A & C bits indicate another station also recognized the address as its own,
        indicating a possible duplicate address. LLC frames or frames with routing information may or
        may not be counted.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        framecopiederrorcounter(35)};
```

6.2.5.37 Frequency Error Counter ATTRIBUTE

```
aFrequencydErrorCounter ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.EventCounter;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bFrequencydErrorCounter BEHAVIOR
    DEFINED AS !This counter contains the number of Frequency errors (CFE). This counter is
        incremented when this station detects the received signal stream’s data signaling rate is out of
        tolerance.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        frequencyderrorcounter(36)};
```

6.2.5.38 Token Error Counter ATTRIBUTE

```
aTokenErrorCounter ATTRIBUTE
    WITH ATTRIBUTE SYNTAX
        MACDefinitions.EventCounter;
    MATCHES FOR EQUALITY;
    BEHAVIOR
        bTokenErrorCounter BEHAVIOR
    DEFINED AS !This counter contains the number of Token errors (CTE). This counter is incremented
        when an active monitor recognizes an error condition resulting in the need to issue a new token.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attribute(7)
        tokenerrorcounter(37)};
```

6.2.6 Attribute groups

6.2.6.1 Station Addresses ATTRIBUTE GROUP

```
gTrStationAddresses ATTRIBUTE GROUP
    GROUP ELEMENTS
        IndMACAddress,
        UniversalAddress,
        FunctionalAddress,
        GroupAddress,
        NAUN;
    DESCRIPTION    !Specifies the addresses of the token ring station!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attributeGroup(Group(8))
        stationcharacteristics(0)};
```

6.2.6.2 Station Characteristics ATTRIBUTE GROUP

```
gTrStationCharacteristics ATTRIBUTE GROUP
    GROUP ELEMENTS
        MicrocodeLevel,
        ProductInstanceId,
        MaxFrameSize;
    DESCRIPTION    !Specifies the characteristics of the token ring station!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attributeGroup(8)
        stationcharacteristics(1)};
```

6.2.6.3 MAC Protocol Options ATTRIBUTE GROUP

```
gTrProtocolOptions  ATTRIBUTE GROUP
    GROUP ELEMENTS
        Burst4Option,
        BeaconHandling,
        ClaimContender,
        EarlyTokenRelease,
        RingRate,
        MultipleFrameTransmission,
        RejectRemove,
        TokenErrorDetection,
        TokenHandling,
        ErrorCountingMethod;
    DESCRIPTION   ! Specifies the token ring MAC protocol operating characteristics.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attributeGroup(8)
        protocoloptions(2)};
```

6.2.6.4 Station’s Ring Characteristics ATTRIBUTE GROUP

```
gTrCharacteristics  ATTRIBUTE GROUP
    GROUP ELEMENTS
        RingNumber,
        ErrorReportTimer,
        AuthorizeFunctionClass,
        AccessPriority,
        PhysicalDropNumber,
        RingStationStatus;
    DESCRIPTION    !Specifies the current status of the token ring station.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attributeGroup(8)
        stationstatus(3)};
```

6.2.6.5 MAC Status ATTRIBUTE GROUP

```
gTrMACStatus  ATTRIBUTE GROUP
    GROUP ELEMENTS
        JoinState,
        MonitorState,
        BeaconInfo,
        TxBeaconType;
    DESCRIPTION    !Specifies the current status of the token ring MAC protocol.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attributeGroup(8)
        macstatus(4)};
```

6.2.6.6 Isolating Error Counters ATTRIBUTE GROUP

```
gTrIsolatingErrors  ATTRIBUTE GROUP
    GROUP ELEMENTS
        LineErrorCounter,
        BurstErrorCounter,
        ACErrorCounter,
        AbortErrorCounter,
        InternalErrorCounter;
    DESCRIPTION   !Specifies the current values of the isolating error counters.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attributeGroup(8)
        isolatingerrors(5)};
```

6.2.6.7 Non Isolating Error Counters ATTRIBUTE GROUP

```
gTrNonIsolatingErrors  ATTRIBUTE GROUP
    GROUP ELEMENTS
        LostFrameErrorCounter,
        RcvCongestionErrorCounter,
        FrameCopiedErrorCounter,
        FrequencyErrorCounter,
        TokenErrorCounter;
    DESCRIPTION  !Specifies the current values of the non isolating error counters.!;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) attributeGroup(8)
        nonisolatingerrors(6)};
```

6.2.7 Actions

6.2.7.1 Open ACTION

```
acOpen  ACTION
    BEHAVIOR
        bOpen  BEHAVIOR
    DEFINED AS !This action requests that the station starts the join ring process.!;
    MODE   CONFIRMED;
    WITH INFORMATION SYNTAX MACDefinitions.OpenData :
    WITH REPLY SYNTAX MACDefinitions.OpenRspData;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) action(9) open(0)};
```

6.2.7.2 Close ACTION

```
acClose  ACTION
    BEHAVIOR
        bClose  BEHAVIOR
    DEFINED AS !This action requests that the station removes itself from the ring.!;
    MODE   CONFIRMED;
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) action(9) close(1)};
```

6.2.8 Notifications

6.2.8.1 Ring Non Operational NOTIFICATION

```
evRingNonOperational NOTIFICATION
    BEHAVIOR
        bRingNonOperational BEHAVIOR
    DEFINED AS ! The ring is in a non-operational state (purging, claiming, or beaconing).!;
    WITH INFORMATION SYNTAX   MACDefinitions.RingNonOpStatus
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) notification(10)
        ringnonop(0)};
```

6.2.8.2 Ring Operational NOTIFICATION

```
evRingOperational NOTIFICATION
    BEHAVIOR
        bRingOperational BEHAVIOR
    DEFINED AS ! The ring has returned to a normal operating status (normal repeat).!;
    WITH INFORMATION SYNTAX   MACDefinitions.RingOpStatus
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) notification(10)
        ringop(1)};
```

6.2.8.3 Ring Beaconing NOTIFICATION

```
evRingBeaconing NOTIFICATION
    BEHAVIOR
        bRingBeaconing BEHAVIOR
    DEFINED AS ! Ring has failed and a station is transmitting beacon frames.!;
    WITH INFORMATION SYNTAX    MACDefinitions.BeaconInfo
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) notification(10)
        ringbeaconing(2)};
```

6.2.8.4 Station Failure NOTIFICATION

```
evStationFailure NOTIFICATION
    BEHAVIOR
        bStationFailure BEHAVIOR
    DEFINED AS ! Station has failed and is no longer connected to the ring.!;
    WITH INFORMATION SYNTAX    MACDefinitions.StationFailure
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) notification(10)
        stationfailure(3)};
```

6.2.8.5 Protocol Error NOTIFICATION

```
evProtocolError NOTIFICATION
    BEHAVIOR
        bProtocolError BEHAVIOR
    DEFINED AS ! Protocol error has been detected.!;
    WITH INFORMATION SYNTAX    MACDefinitions.ProtocolError
    REGISTERED AS {iso(1) member-body(2) us(840) 802dot5(10033) notification(10)
        protocolerror(4)};
```

6.2.9 MAC Definitions

```
MACDefinitions
DEFINITIONS::= BEGIN

IMPORTS
    MACAddress
    FROM IEEE802CommonDefinitions{iso(1) memberbody(2) us(840) ieee802-1F(10011)
    asn1Module(2) commondefinitions(0) version(0)}
-- Imported implicitly from ISO/IEC 8824 : 1990

-- Definitions in Alphabetical Order
AccessPriority::= INTEGER  -- Range 0 to 7

AuthorizedFunctionClasses::= OCTETSTRING(SIZE(2))  -- (see 3.3.4 X'06' subvector)

BeaconHandling::= CHOICE {
                beforeInitialization  [0] IMPLICIT NULL,     -- FBHO=0
                 afterInitialization  [1] IMPLICIT NULL  }   -- FBHO=1

BeaconInfo::= SEQUENCE {
                    beaconSourceAddr  [0] IMPLICIT BeaconSA,
                          beaconType  [1] IMPLICIT BeaconType,
                           beaconUNA  [2] IMPLICIT BeaconRUA,
                          beaconDrop  [3] IMPLICIT BeaconPDN }

BeaconPDN::= OCTET STRING (SIZE 4) -- PhysicalDropNumber reported in beacon frame
                                   -- (part of BeaconInfo

BeaconRUA::= MACAddress  -- UNA subvector field reported in beacon frame
                                   -- (part of BeaconInfo)

BeaconSA::= MACAddress  -- SA field of beacon frame
                                   -- (part of BeaconInfo)

BeaconType::= CHOICE {  --  Beacon type subvector reported in beacon frame
                         typeUnknown  [0] IMPLICIT NULL,     -- unknown
                               type1  [1] IMPLICIT NULL,     -- reconfiguration
                               type2  [2] IMPLICIT NULL,     -- signal loss
                               type3  [3] IMPLICIT NULL,     -- symbol streaming
                               type4  [4] IMPLICIT NULL  }   -- claim streaming

Burst4Option ::= CHOICE {
                      noCompensation  [0] IMPLICIT NULL,     -- FBEO=1
                         compensates  [1] IMPLICIT NULL }    -- FBEO=0

ClaimContender::= CHOICE {
                         noContender  [0] IMPLICIT NULL,     -- FCCO=0
                           contender  [1] IMPLICIT NULL  }   -- FCCO=1

EarlyTokenRelease::= CHOICE {
                          notEnabled  [0] IMPLICIT NULL,     -- FETO=0
                             enabled  [1] IMPLICIT NULL }    -- FETO=1
```

```
ErrorCountingMethod::= CHOICE {
                       triggered  [0] IMPLICIT NULL,     -- FECO=0
                     freeRunning  [1] IMPLICIT NULL }    -- FECO=1

ErrorReportTimer::= INTEGER  -- Number in 0.01 s increments, range 0 to 655.35 s

EventCounter::= INTEGER  -- 32-bit counter

FunctionalAddresses::=  OCTETSTRING(SIZE(4)) -- (see 3.2.4.1)

GroupAddresses::= LIST OF { MACAddress}

JoinState::= CHOICE {
                    notSpecified  [0] IMPLICIT NULL,     -- Not specified or unknown
                          bypass  [1] IMPLICIT NULL,     -- JS=BP
                        lobeTest  [2] IMPLICIT NULL,     -- JS=LT
               detectMonPresent  [3] IMPLICIT NULL,     -- JS=DMP
                awaitNewMonitor  [4] IMPLICIT NULL,     -- JS=ANM
              duplicateAddrTest  [5] IMPLICIT NULL,     -- JS=DAT
           neighborNotification  [6] IMPLICIT NULL,     -- JS=NN
          requestinitialization  [7]  IMPLICIT NULL,    -- JS=RQI
                   joinComplete  [8] IMPLICIT NULL,     -- JS=JC
                     bypassWait  [9] IMPLICIT NULL }    -- JS=BPW

LastMonitorAddress::= MACAddress   -- value of X'0A' subvector - SA of last AMP/SMP
                                 -- (used in ProtocolErrorStatus part of PrototcolError)

LocalAddress::= MACAddress  -- Local assigned address, value of all zeros indicates not specified
                            -- (part of openData)

MaxFrameSize::= INTEGER -- Range 133 to 65 535 octets
                                   -- The maximum allowable values for MaxFrameSize is
                                   -- dependent on ring speed as follows:
                                       --    1) 4550 at 4 Mbit/s
                                       --    2) 18 200 at 16 Mbit/s

MicrocodeLevel::=OCTET STRING (SIZE 1..32)

MonitorErrorCode::= CHOICE  { -- value of X'30' Error Code subvector
                                 -- (used in ProtocolErrorStatus part of PrototcolError)
                  claimReceived  [0] IMPLICIT NULL,     -- Claim Frame received
               duplicateMonitor  [2] IMPLICIT NULL,     -- Another Active Monitor detected
               duplicateAddress  [3] IMPLICIT NULL }    -- Another station with same address

MonitorState::= CHOICE {
                   notSpecified  [0] IMPLICIT NULL,     -- Not specified or unknown
                      ringPurge  [1] IMPLICIT NULL,     -- MS=TRP
                  claimTransmit  [2] IMPLICIT NULL,     -- MS=TCT
                 beaconTransmit  [3] IMPLICIT NULL,     -- MS=TBN
                     beaconTest  [4] IMPLICIT NULL,     -- MS=BNT
                   beaconRepeat  [5] IMPLICIT NULL,     -- MS=RBN
                    claimRepeat  [6] IMPLICIT NULL,     -- MS=RCT
                  activeMonitor  [7] IMPLICIT NULL,     -- MS=RPT & FAM=1
                 standbyMonitor  [8] IMPLICIT NULL }    -- MS=RPT & FAM=0
```

```
MultipleFrameTransmission::= CHOICE {
                         relinquish  [0] IMPLICIT NULL,      -- FMFTO=0
                           continue  [1] IMPLICIT NULL  }    -- FMFTO=1

NAUN::= SEQUENCE {
                                sua  [0] IMPLICIT SUA
                      uPhysicalDrop  [1] IMPLICIT PhysicalDropNumber  } -- from AMP/SMP frame

OpenData::= SEQUENCE {  -- (used in Open Action)
                       localaddress  [0] IMPLICIT LocalAddress
                           ringrate  [1] IMPLICIT RingRate  }

OpenRspData::=  SEQUENCE {  -- (used in Open Action)
                          productId  [0] IMPLICIT ProductInstanceId
                           micLevel  [1] IMPLICIT MicrocodeLevel
                         openStatus  [2] IMPLICIT OpenStatus
                           ringRate  [3] IMPLICIT RingRate
                                una  [4] IMPLICIT NAUN    }

OpenStatus::= CHOICE {  -- (used in OpenRspData)
                 sSuccess  [0]  IMPLICIT NULL,  --Successful
          sNotOperational  [1]  IMPLICIT NULL,  --Internal error - HW or SW failure, etc.
          sInvalidAddress  [2]  IMPLICIT NULL,  --MAC address is invalid (group or not specified)
        sRateNotSupported  [3]  IMPLICIT NULL,  --The requested data rate is not supported
             sFailGeneral  [4]  IMPLICIT NULL,  --Failure in activation - General
            sFailLobeTest  [5]  IMPLICIT NULL,  --Failure in activation - Lobe Test Failure
            sFailDataRate  [6]  IMPLICIT NULL,  --Failure in activation - Medium Rate Error
           sFailDuplicate  [7]  IMPLICIT NULL,  --Failure in activation - Duplicate Address
          sFailDATTimeout  [8]  IMPLICIT NULL,  --Failure in activation - DAT Timeout - TJR expired
           sFailNNTimeout  [9]  IMPLICIT NULL,  --Failure in activation - Neighbor Notifi. timeout
         sFailInitTimeout [10]  IMPLICIT NULL,  --Failure in activation - Initialization timeout
       sFailRingBeaconing [11]  IMPLICIT NULL,  --Failure in activation - Beacon received
        sFailPurgeTimeout [12]  IMPLICIT NULL,  --Failure in activation - Purge failure - TRP expired
        sFailClaimTimeout [13]  IMPLICIT NULL,  --Failure in activation - Claim failure - TCT expired
       sFailBeaconTimeout [14]  IMPLICIT NULL,  --Failure in activation - Beacon failure - TBT expired
          sFailSignalLoss [15]  IMPLICIT NULL,  --Failure in activation - Signal Loss
            sFailInternal [16]  IMPLICIT NULL,  --Failure in activation - Internal Error
             sFailTxFault [17]  IMPLICIT NULL,  --Failure in activation - Station transmit failure
                  sRemove [18]  IMPLICIT NULL }  --Failure in activation - Remove Frame received

PhysicalDropNumber::=  OCTET STRING (SIZE(4))
ProductInstanceId::= OCTET STRING (SIZE(1..31))

ProtocolError::= SEQUENCE  {  -- (used in ProtocolError Event)
                    protocolErrType  [0] IMPLICIT ProtocolErrType,
                  protocolErrStatus  [1] IMPLICIT ProtocolErrStatus  }

ProtocolErrType::= CHOICE   {  -- (part of ProtocolError used in ProtocolError Event)
                sActiveMonitorError  [0] IMPLICIT NULL,      -- Active Monitor Error
           sNotificationIncomplete  [1] IMPLICIT NULL}      -- Neighbor Notification Incomplete

ProtocolErrStatus:: SEQUENCE  {  -- (part of ProtocolError used in ProtocolError Event)
                          errorCode  [0] IMPLICIT MonitorErrorCode,
                 lastMonitorAddress  [0] IMPLICIT LastMonitorAddress }
```

```
RejectRemove::= CHOICE {
                          removes  [0] IMPLICIT NULL,     -- FRRO=0
                          rejects  [1] IMPLICIT NULL  }   -- FRRO=1

RingNonOpStatus::= CHOICE  {  -- (used in RingNonOperational Event)
               sRingNonOp [0]  IMPLICIT NULL, -- Claim - Purge Failure
            sPurgeFailure [1]  IMPLICIT NULL, -- Claim - Purge Failure
           sRingClaiming  [2]  IMPLICIT NULL, -- Claim - Claim Received
        sClaimNoMonitor   [3]  IMPLICIT NULL, -- Claim - No Active Monitor (TJR, TNT, TSM=E)
        sClaimSignalLoss  [4]  IMPLICIT NULL, -- Claim - Signal Loss
          sRingBeaconing  [5]  IMPLICIT NULL } -- Beacon received

RingNumber::= OCTET STRING (SIZE(2))

RingOpStatus::= CHOICE   {  -- (used in RingOperational Event)
                  sRingOp [0]  IMPLICIT NULL, -- Ring operational - Monitor status not specified
         sStandbyMonitor  [1]  IMPLICIT NULL,  --This station is a standby monitor (FAM=0)
          sActiveMonitor  [2]  IMPLICIT NULL} --This station is the Active Monitor (FAM=1)

RingRate::= CHOICE {
                        rate4Mbps  [0] IMPLICIT NULL,     -- FMRO=0
                       rate16Mbps  [1] IMPLICIT NULL  }   -- FMRO=1

RingStationStatus::= OCTET STRING (SIZE(6))  --   (See 3.3.4 X'29' subvector)

StationFailure::= CHOICE   {  -- (used in StationFailure Event)
             sStationDown [0]  IMPLICIT NULL,           -- Station Failure - cause not specified
                  sRemove [0]  IMPLICIT NULL,           --Remove Frame received
               sWireFault [1]  IMPLICIT NULL,           --Wire Fault
             sFailInternal [2] IMPLICIT NULL,           --Internal Failure
                 sFailTest [3] IMPLICIT NULL}           --Beacon Test Failure

SUA::= MACAddress -- Address of upstream neighbor

TokenErrorDetection::= CHOICE {
                       notCounted  [0] IMPLICIT NULL,     -- FTEO=1
                          counted  [1] IMPLICIT NULL  }   -- FTEO=0

TokenHandling::= CHOICE {
                      activeTimer  [0] IMPLICIT NULL,     -- FTHO=0
               periodicInspection  [1] IMPLICIT NULL  }   -- FTHO=1
```

7. Station attachment specifications

This clause defines the physical medium components (PMC) of the station attachment PHY layer as well as restrictions on the jitter accumulation in a ring caused by any PHY. This includes specifications for jitter accumulation, a transmitter, a receiver, a transmission channel, and medium interface connectors (MICs). The objective of these specifications is to provide a transmission system which

a) Ensures compatibility of independently developed physical and electrical interfaces.

b) Provides a communication channel that consists of a transmitter, receiver, and channel with an equivalent bit error rate (BER) of less than or equal to 10^{-9} and a ring with a ring equivalent BER of 10^{-8}.[8]

The analyses used in determining these specifications are based on linear transfer functions for the timing recovery circuits.

Subclause 7.1, which applies to all PHYs, provides the media-independent specifications that limit the jitter buildup in rings to a level that all PHYs can tolerate, regardless of media type. Clause 7.2 provides the media-dependent specifications of the PMC for transmission over STP and UTP media.[9]

7.1 Media-independent PHY specifications

This subclause describes the media independent PHY[10] requirements. These specifications are designed to limit jitter buildup in rings to levels that all PHYs, regardless of media type, can tolerate.

7.1.1 Accumulated jitter

Accumulated jitter is the jitter that builds up as it passes through each ring PHY subsequent to the active monitor. Since each ring PHY receiver is unable to fully compensate for the signal distortion introduced by the sending transmitter, interconnecting channel and its own receiver, jitter is introduced at each PHY. Since each PHY subsequent to the active monitor transmits with the clock derived from the jittered signal at the receiver, a filtered form of the jitter at each PHY is transferred to subsequent PHYs, causing the jitter to accumulate as the signal travels around the ring. The accumulated jitter is measured as the time displacement of the transmitted waveform driving a conformant channel with respect to the active monitor station crystal clock as any valid data pattern passes through the ring PHYs. The accumulated jitter has a total magnitude (AJ) and a magnitude of the rate of phase change called Accumulated Phase Slope (APS). The accumulated jitter is typically dominated by the jitter that is correlated to data pattern "correlated jitter") but also has components from sources not related to the data pattern ("uncorrelated jitter"). Limits on AJ and APS, specified in this clause, control the combination of the correlated and uncorrelated jitter.

The PHYs must track the accumulated jitter[11] and the elastic buffer (5.8.3) in the Active Monitor Station must be able to compensate for this jitter. To ensure this, the total magnitude of the Accumulated Jitter

---

[8] Equivalent BER is used because the effect of both single and multiple bit errors within a frame is to cause a frame error and a retransmission. Equivalent BER can be determined by assuming that all frame errors are caused by a single bit error and computing accordingly. For example, the worst-case allowable station packet error rate for a sequence of packets, each 1250 octets long (10 kbit long), is 10^{-5}.

[9] Alternate media are the subject of future study.

[10] Since the specifications in 7.1 are a function not only of PMC parameters but also of PSC parameters such as station delay, the term PHY is used in 7.1 instead of PMC.

[11] *Accumulated jitter* is equivalent to the *cumulative latency variations* of 5.8.3.

(AJ) is limited to an amount smaller than the elasticity[12] of the elastic buffer in the active monitor station. AJ is the peak-to-peak magnitude of the jitter after the last PHY in the ring.

Accumulated phase slope is defined as the phase slope (rate of phase change) that is accumulated by the PHYs in relation to the active monitor's imbedded crystal clock. The magnitude of the filtered accumulated phase slope (FAPS)[13] must be kept below a maximum value in order for adequate remaining eye-opening margin to be available to tolerate jitter from reflections, NEXT, and other sources under worst-case data patterns.

To aid interoperability, the phase slope changes due to jitter not present in the test channel for the FAPS specification described below but present in actual channels (such as the jitter caused by crosstalk and reflections) must be limited. This is accomplished by limiting the allowable increase in phase slope in a ring per unit of increased jitter at each PHY by specifying a parameter termed DFAPS (Delta FAPS), which is the increase in phase slope caused by a 1 ns increase in jitter at each PHY.

The AJ, FAPS, and DFAPS values are defined after the minimum number of PHYs detailed in table 7-4 (for STP and UTP cases). Where PHYs with simple well-behaved clock recovery systems are added to the ring, measurements on a lower number of PHYs may allow extrapolation to the values expected after the full number of PHYs. For these measurements, each PHY has the same Test_Channel. Each measurement must be performed on a single sequence of adjacent zero-crossings of the data over any valid data pattern. The zero-crossing timing data is filtered by a low-pass filter with a single pole at F6 for FAPS. Between each measurement the symbol polarity shall be randomized at each station. These measurements shall be repeated at least 15 times and the average of the maximum magnitudes (filtered in the case of FAPS) used to determine AJ and FAPS.

The PHY shall meet the following AJ and FAPS specifications for any test channel conformant with Test_Chan_A (specified in 7.2.4.4 for STP and UTP cases) during the transmission of any valid data pattern. The phase slope generated by the chain of PHYs shall not increase by more than DFAPS for each 1 ns increase in jitter at each PHY over the range of jitter possible at the PHY for any channel conformant with the channel specifications in 7.2.4. (The range of jitter and the phase slope response to it is implementation-dependent and may require a verification method specific to the implementation. A method suitable for a number of common implementations is outlined in C.2.5 for informational purposes.)

Table 13—Accumulated jitter specifications

	4 Mbit/s	**16 Mbit/s**
|AJ|	< 600 ns	< 750 ns
|FAPS|	< 0.0035 ns/ns	< 0.0082 ns/ns
DFAPS	0.0012 ns/ns	0.0048 ns/ns
F6	90 kHz	360 kHz

---

[12] The elasticity is the quantity *B* in 5.8.3.1.

[13] *Filtered* means that the phase-slope waveform is filtered by a single-pole low-pass filter to remove high-frequency jitter, which is typically not accumulated, from the measurement. If the phase slope is an analog waveform, it should be filtered by a standard single-pole filter. Care must be taken to ensure that the band-limiting of the demodulation process is significantly above the pole of the low-pass filter. If the phase slope is a digitally sampled waveform, such as that derived from a time interval analyzer, the waveform is usually interpolated to a waveform containing equally spaced sample points and then filtered by the digital equivalent of a single-pole analog filter. Note that the phase slope waveform is filtered to remove high-frequency jitter components, which would give misleadingly high values of the phase slope when taking the derivative.

7.1.2 Accumulated uncorrelated jitter

Accumulated uncorrelated jitter is the component of the jitter that is not correlated to the data pattern. Typical sources of uncorrelated jitter include time-dependent transmit duty cycle distortion, high-frequency cyclic jitter, and phase noise in the local oscillator's output. Accumulated uncorrelated jitter is specified[14] in terms of the accumulated uncorrelated jitter alignment (AUJA) error it would produce in an ideal first-order clock-recovery circuit with a 3 dB bandwidth of F6 receiving the uncorrelated jitter built up from a chain of similar retiming elements where the alignment error is defined as the difference between the input and output jitter waveforms of the clock-recovery circuit at any particular point in time. This measurement effectively high-pass filters the accumulated jitter, resulting in an alignment error with a zero mean. The minimum length of this chain is detailed in table 13 (for STP and UTP cases). Where PHYs with simple well-behaved clock-recovery systems are added to the ring, measurements on a lower number of PHYs may allow extrapolation to the values expected after the full number of PHYs. For informational purpose, C.2.7 provides an ideal first-order clock-recovery model that may be used to determine AUJA error.

The PHY shall meet the following AUJA specifications for any test channel conformant with Test_Chan_A (STP and UTP cases) for data streams composed of all ones and data streams composed of all zeros. The probability of the zero-to-peak alignment error at each symbol element interval exceeding AUJA shall be less than 0.3%.

Table 14—Accumulated uncorrelated jitter specifications

	4 Mbit/s	**16 Mbit/s**
AUJA	10 ns	1.5 ns
F6	90 kHz	360 kHz

7.1.3 PHY net delay

Interoperability between PHYs with different implementations on the same ring requires that the derivative of the magnitude of the FAPS with respect to the net delay time (NDT; as defined in C.2.2) of each node at the operating point shall be greater than or equal to zero (i.e., increasing the NDT does not decrease the magnitude of the FAPS).[15]

7.2 Media-dependent PMC specifications

This subclause describes the media-dependent specifications for STP and UTP station PMCs. All resistive loads discussed in this clause have a ±1% tolerance unless otherwise noted. The station port signal names used in this specification map to the MIC_S and MIC_U contacts as shown in table 15.

---

[14] Note that accumulated uncorrelated jitter is also indirectly limited by the accumulated jitter specifications.

[15] C.2.6 contains a clarification of the PHY net delay specification.

Table 15—Station port signal to MIC_S and MIC_U contact mapping

MIC_U contact	MIC_S contact	Station port signal
1	—	Not used by this standard
2	—	Not used by this standard
3	B	TX-A
4	R	RX-A
5	G	RX-B
6	O	TX-B
7	—	Not used by this standard
8	—	Not used by this standard

Table16 details the minimum number of PMCs that must be supported on STP and UTP at 4 and 16 Mbit/s.

Table 16—Ring PMC counts

Media	Data rate	Min. # of supported PMCs
STP	4 Mbit/s	250
STP	16 Mbit/s	250
UTP	4 Mbit/s	144
UTP	16 Mbit/s	250

Note that each retiming port in the ring counts as a single PMC, so that, for example, a 4 Mbit/s UTP ring using all ARC ports supports a minimum of 72 stations connected to 72 ARC ports for a total of 144 PMCs.

7.2.1 Phantom signaling and wire fault detection

Stations that use STP and UTP transmission media and are controlled by a ring access control (RAC) function (see 5.9) shall use the signaling method described below to effect access to the ring through the concentrator. They use the wire fault detection method described below to detect certain fault conditions in the transmission channel wiring.

As described in 5.9, when the RAC receives Insert_station and Remove_station requests from the MAC, it shall signal *ring insertion* and *ring bypass* requests, respectively, to the TCU in the concentrator. Phantom signaling, which is transparent to the Manchester data signals on the channel, shall be used to send the ring insertion and ring bypass requests to the TCU. The TCU effects ring insertion and ring bypass as described in 8.3.1. An example of a phantom signaling circuit is shown in figure 29. It consists of two circuits whose function, described in 7.2.1.1 and 7.2.1.2, is designed to ensure that the station can detect a single open wire and certain short-circuit conditions in the channel between the station and the TCU.

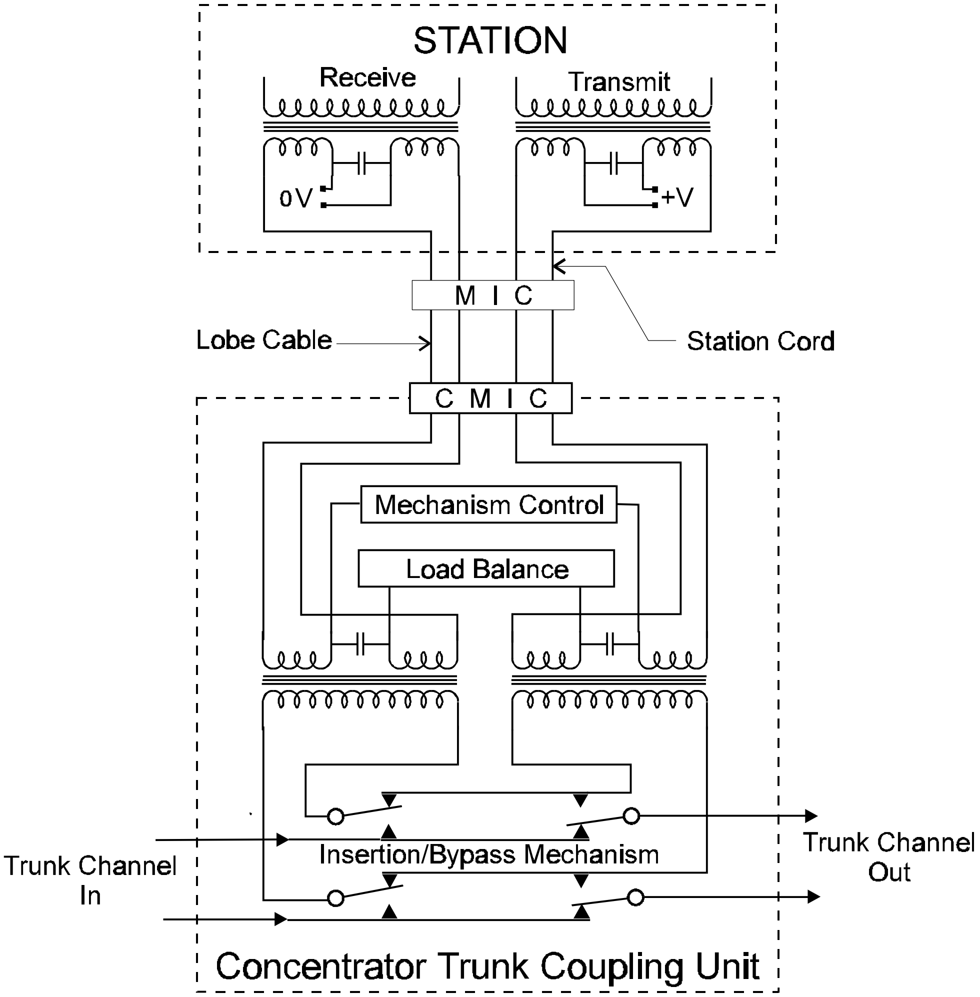

Figure 29—Example of a station-to-TCU connection

7.2.1.1 Ring insertion and ring bypass

When the RAC receives an Insert_station request from the MAC, the RAC shall signal *ring insertion* to the TCU by applying *power* to both phantom circuits. Unless directed by the MAC, the RAC should not signal ring insertion. Under any circumstances, the RAC shall not signal *ring insertion* for more than 25 ms over any 5 s period unless the station is opening on the ring.[16] In response to Remove_station, the RAC shall signal *ring bypass* to the TCU by applying *no power* to both phantom circuits. The time from the receipt of Remove_station to the indication of *no power* shall not exceed a time that is the station's Remove Wait Time (see 3.4.2.11, TRW timer) minus 185 ms when the station is connected to a resistive load of 5.4 kΩ on each circuit.

The phantom circuits shall have the following sources and returns:

a) Source TX-A
b) Source TX-B

---

[16] The intent of this statement is to avoid glitches in the phantom signal that could cause the concentrator to insert when the station is not ready to repeat data and thus break the ring until the concentrator removes. An allowance of 25 ms is made for glitches that might occur during station power-up; however, it is desirable that the station not glitch the phantom signal at all. See 8.3.2 for recommendations on concentrator glitch tolerance.

c) Return RX-A
d) Return RX-B

The differential dc voltage between any source and any return shall not exceed 7.0 V. In this clause, *voltage* is defined as the difference in potential between a circuit's source and return MIC terminals. Furthermore, *current* is defined as the electrical current flowing out of the source and back into the return.

Power is defined as follows: Voltage is always impressed on both circuits at the same time where *the same time* means within 0.1 s of each other. This voltage shall not exceed 7 V under any compliant load conditions. The current shall not exceed 20 mA under any compliant load condition. When the current is less than 2 mA, the voltage shall be greater than 3.5 V under compliant load conditions. When the current is less than 1 mA, the voltage shall be greater than 4.1 V under compliant load conditions. Power shall be applied to both load circuits. Under identical compliant load conditions, the lower of the two circuit voltages shall be within 5% of the higher circuit voltage. The station PMC shall not be required to supply more than 2 mA to achieve insertion. Each phantom drive source shall maintain the currents and voltages specified above to either phantom return. Whenever the load impedance presented to the station at its MIC terminals does not allow the voltage to reach a level of 3.5 V, the station shall supply at least 2 mA phantom current to its MIC terminals while power is applied.

No-power is defined as follows: The voltage shall not exceed 1 V under any compliant load condition. The current shall not exceed 340 µA under any compliant load condition.

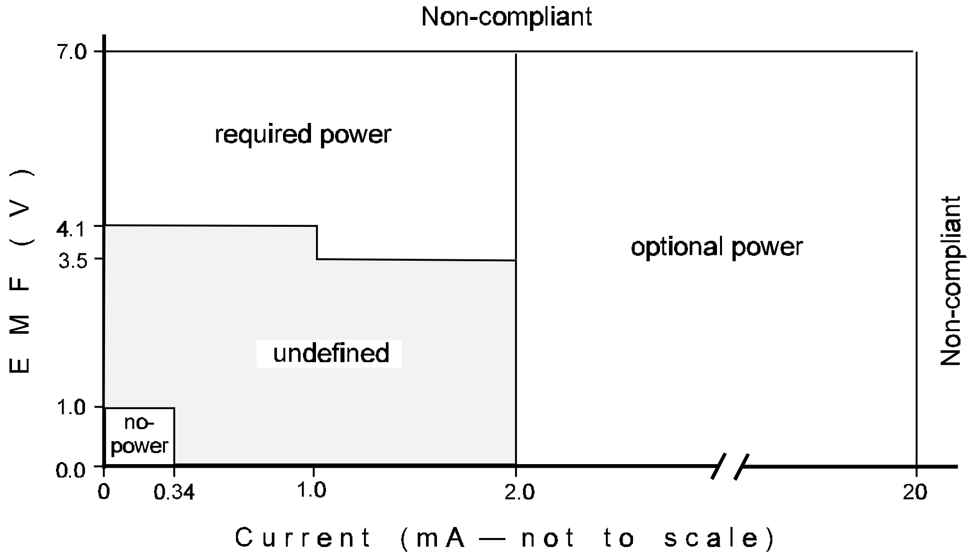

Figure 30—Phantom power/no-power regions

Compliant load conditions are defined as follows. Conditions are illustrated in figure 31.

a) *MIC loopback* is defined as an independent passive load to each circuit under *power* and *no-power* conditions. The resistance of each of these passive loads shall be less than 50 Ω.

b) *TCU presence* is defined as an independent passive load to each circuit under power conditions. The resistance of each of these passive loads shall be greater than 2.9 kΩ and less than 5.4 kΩ (100 Ω for cabling plus 5.3 kΩ for TCU). This load may be highly capacitive and nonlinear but should not be inductive.

c) *TCU loopback* is defined as an independent passive load to each circuit under *no-power*. The resistance of each of these passive loads shall be greater than 2.9 kΩ and less than 5.4 kΩ (100 Ω for cabling plus 5.3 kΩ for TCU). This load may be highly capacitive and nonlinear but should not be inductive.

d) *Inserted* is defined as an independent passive load to each circuit under *power* conditions. The resistance of each of these passive loads shall be greater than 2.9 kΩ and less than 5.4 kΩ. Furthermore, the resistance of these two loads shall be matched within ± 15%.

e) *No load* is defined as an independent passive load to each circuit under *power* and *no-power* conditions. The resistance of each of these passive loads shall be greater than 50 kΩ.

The definition of resistance is *V*/*I* at the measurement point. The leakage resistance between the two loads shall be greater than 100 kΩ.

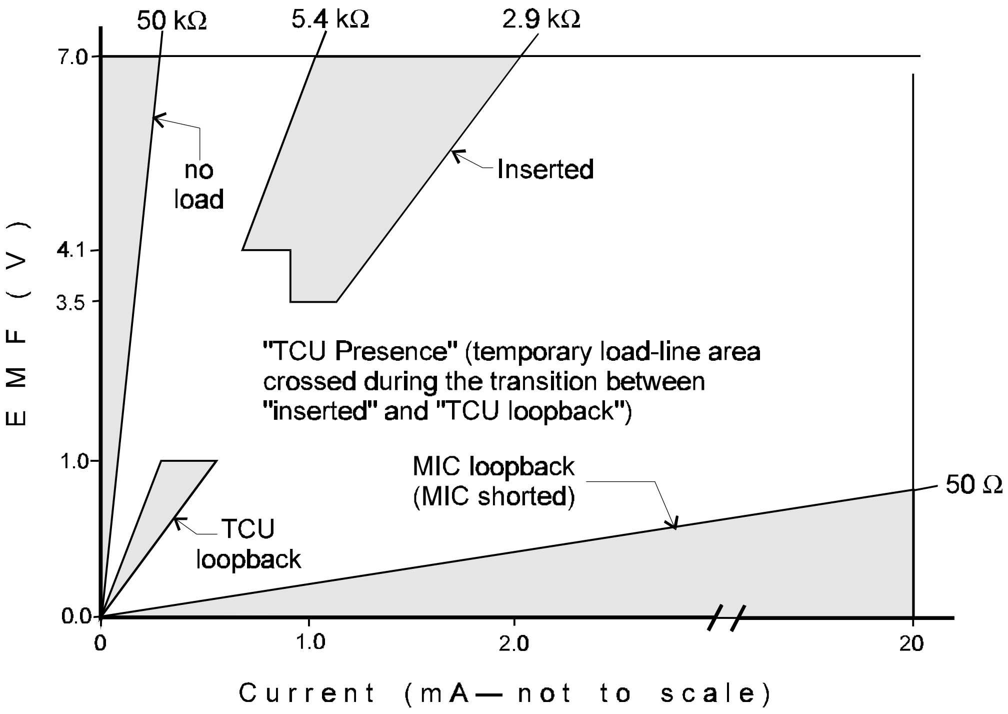

Figure 31—Operating regions for compliant load conditions

7.2.1.2 Wire_fault detection

To detect some opens and shorts in the channel between the station and the TCU, as well as some opens and shorts in the TCU, the phantom signaling circuit shall consist of two independent circuits as shown in figure 29. The pairing of sources to returns in the figure is only an example and either permutation of pairing may be used.

While the PMC is requesting “ring insertion” as described previously, the PMC shall indicate “Not detected” for Wire_fault as described in 5.1.4.1 when the resistances of both phantom circuits are within the range of 2.9–5.4 kΩ.

While the PMC is requesting “ring insertion,” the PMC shall indicate “detected” for Wire_fault as described in 5.1.4.1 for the following conditions.

a) The resistance of either phantom circuit is less than 50 Ω.
b) The resistance of either phantom circuit is greater than 50 kΩ.

The indications for conditions outside the above resistance ranges are unspecified.

7.2.2 Transmitter specification

The PMC shall provide a transmitter meeting the specifications of this clause. Phantom signaling for ring access control is described separately in 7.2.1. The phantom signals (if provided) shall be in a powered and loaded state for all measurements of the transmitter specified in this clause. The dc phantom load used to test the transmitter specifications shall be equivalent to a pair of resistors with a 15% mismatch. The smaller resistor shall be 2.9 kΩ. The ac transmitter load between TX-A and TX-B used to test the transmitter specifications shall be 100 Ω for UTP and 150 Ω for STP. These specifications are not intended to cover electromagnetic compatibility concerns (see E.3.1). All measurements of transmitter properties are made at the appropriate reference MIC.

7.2.2.1 Differential output signal duty cycle distortion

Transmitter duty cycle distortion (TDCD) is defined as one half the difference between the arithmetic averaged positive and negative pulse widths of the data signal. The pulse widths are defined from the ac-coupled zero-crossings of the differential waveform measured at the appropriate MIC with a resistive load. The averaged widths are determined by averaging at least 128 measurements of each of the positive and negative widths.

Transmitter duty cycle distortion shall be measured at the transmitter output of a PMC with the PMC retiming data from its timing recovery circuit and transmitting a constant stream of all Data_zero (see 5.3) or all Data_one symbols (to avoid intersymbol interference disturbance with the measurement). The transferred jitter from the transmitter transmitting to the PMC under test should be minimized. The absolute value of the duty cycle distortion for either data stream shall be within the following specifications:

Table 17—Transmitter duty cycle distortion (TDCD) specifications

	4 Mbit/s	**16 Mbit/s**
TDCD (STP 150 Ω load)	< 6 ns	< 1.5 ns
TDCD (UTP 100 Ω load)	< 6 ns	< 1.5 ns

7.2.2.2 Differential output waveform

The differential waveform is measured at the appropriate MIC with a resistive load. The transmitter waveform shall meet the criteria in the following two subclauses:

7.2.2.2.1 Waveform edge jitter

For a data stream alternating between a stream of Data_zero symbols and a stream of Data_one symbols, Tdiff01 is the difference in transmitter data edge placement (relative to an internal transmit clock or conceptual equivalent) between a zeros data pattern and a ones data pattern. Referring to figure 32, define Tdiff01 as

Tdiff01 = (Tone – Toclk) – (Tzero – Tzclk)

where Tone is the ac-coupled zero-crossing point of the Data_one's edge being measured, Toclk is the ac-coupled zero-crossing point of the transmit clock's edge that latches out the Data_one's edge being measured, Tzero is the ac-coupled zero-crossing point of the Data_zero's edge being measured, and Tzclk is the ac-coupled zero-crossing point of the transmit clock's edge that latches out the Data_zero's edge being measured. Positive Tdiff is defined as (Tone–Toclk) being greater than (Tzero–Tzclk); that is, the Data_zeros pattern edges occur closer to their corresponding ideal clock edges than the Data_ones pattern edges.

To avoid introducing the effects of transmitter duty cycle distortion into this measurement, Tdiff01 is calculated only with pairs of Tone and Tzero measurements made either on the rising edges of data only or on the falling edges of data only. Tdiff01 shall be calculated for both rising and falling edges. Tdiff01 may show variation across repeated measurements due to phase noise. This variation may be removed by using the mean value of repeated measurements on rising edges or on falling edges. The mean of the rising edge measurements and the mean of the falling edge measurements shall both independently meet the Tdiff01 specification.

Note that in the absence of the relevant transmit clock, the output of a sufficiently narrow band PLL or other timing recovery circuit receiving the transmitted data may be used in place of the transmit clock. Sufficiently narrow band means low enough to filter out the pattern-dependent jitter in the particular data patterns used for these tests. The station must be transmitting using its local crystal oscillator in this case.

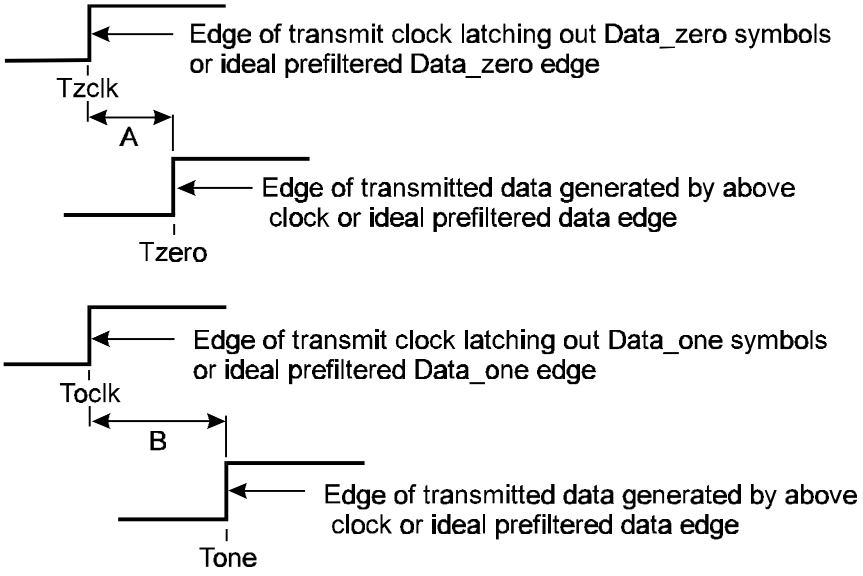

Tdiff01 = (Tone - Toclk) - (Tzero - Tzclk) = B - A

NOTE—Tone and Tzero must either (1) both be measured on rising edges of data or (2) both be measured on falling edges of data to ensure that duty cycle distortion does not affect the measurement.

Figure 32—Tdiff measurement illustration

Tdiff01 shall meet the following specification:

a) at 4 Mbit/s –0.5 ns < Tdiff01 < 0.5 ns
b) 16 Mbit/s 0.3 ns < Tdiff01 < 1.0 ns

The Tdiff01 is allowed to vary over a range only to allow for manufacturing tolerances. To aid interoperability it is important that manufacturers center their transmitter designs in the center of the given range.

Similarly, for any valid data stream consisting of random data, start delimiters (SDs), and end delimiters (EDs):

Tdiffmax = (TR1 – Tclk1) – (TR2 – Tclk2)

where TR1 is the ac-coupled zero-crossing point of the first data edge being measured, Tclk1 is the ac-coupled zero-crossing point of the transmit clock's edge that latches out the first data edge being measured, TR2 is the ac-coupled zero-crossing point of the second data edge being measured, and Tclk2 is the ac-coupled zero-crossing point of the transmit clock's edge that latches out the second data edge being measured. To avoid introducing the effects of transmitter duty cycle distortion into this measurement, Tdiffmax is calculated only with pairs of TR1 and TR2 measurements made either on the rising edges of data only or on the falling edges of data only. Tdiffmax shall be calculated for both rising and falling edges.

Tdiffmax shall meet the following specification:

a) at 4 Mbit/s -2.0 ns < Tdiffmax < 2.0 ns
b) at 16 Mbit/s -2.0 ns < Tdiffmax < 2.0 ns

7.2.2.2.2 Waveform pulse shape

For a data stream consisting of Data_zero symbols, the voltage waveform for positive and negative pulses shall meet the template described in figure 33 and table 18 (16 Mbit/s) or figure 36 and table 21 (4 Mbit/s). For comparison to the template, time 0% is the time of the ac-coupled zero-crossing at the start of the waveform; time 100% is the time of the ac-coupled zero-crossing at the end of the waveform. Both positive and negative pulses are normalized by *Vpp*/2. The amplitude *Vpp* is measured as the peak-to-peak amplitude of the fundamental sine wave component of the Data_zero symbols waveform.

For a data stream consisting of Data_one symbols, the voltage waveform for positive and negative pulses shall meet the template described in figure 34 and table 19 (16 Mbit/s) or figure 36 and table 22 (4 Mbit/s) when normalized by *Vpp*/2 measured on the Data_zero symbols waveform.

For a data stream consisting of a starting delimiter (SD) preceded by a minimum of 20 Data_zero symbols, the voltage waveform of the first pulse three data signal elements wide shall meet the template described in figure 35 and table 20 (16 Mbit/s) or figure 36 and table 23 (4 Mbit/s) when normalized by *Vpp*/2 measured on the Data_zero symbols waveform.

In the following waveform pulse shape templates, (figures 33 through 36 and tables 18 through 23), the times are relative to the zero-crossing point at the start of the waveform and are scaled relative to the time between the zero-crossing point at the start of the waveform and the zero-crossing point at the end of the waveform. Measured waveforms should be normalized by *Vpp*/2 for comparison to the template. The template consists of the space between the two piecewise linear curves connecting the above sets of data points.

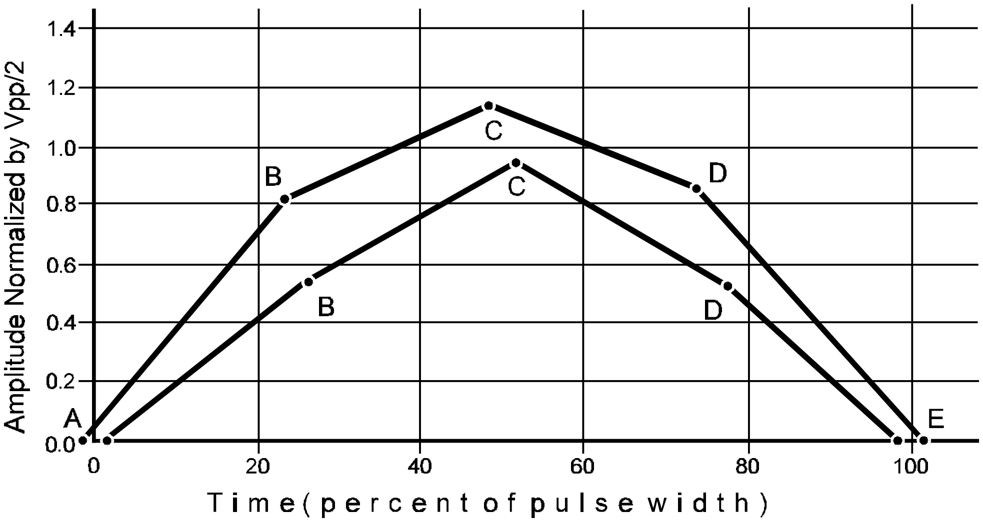

NOTE—See table 18 for point coordinates.

Figure 33—16 Mbit/s template for Data_zero symbols waveform

Table 18—16 Mbit/s Data_zero symbols template

Point	Upper time (%)	Upper amplitude	Lower time (%)	Lower amplitude
A	–1.5	0	1.5	0
B	23.5	0.83	26	0.55
C	48.5	1.15	51.5	0.95
D	73.5	0.862	77.5	0.52
E	101.5	0	98.5	0

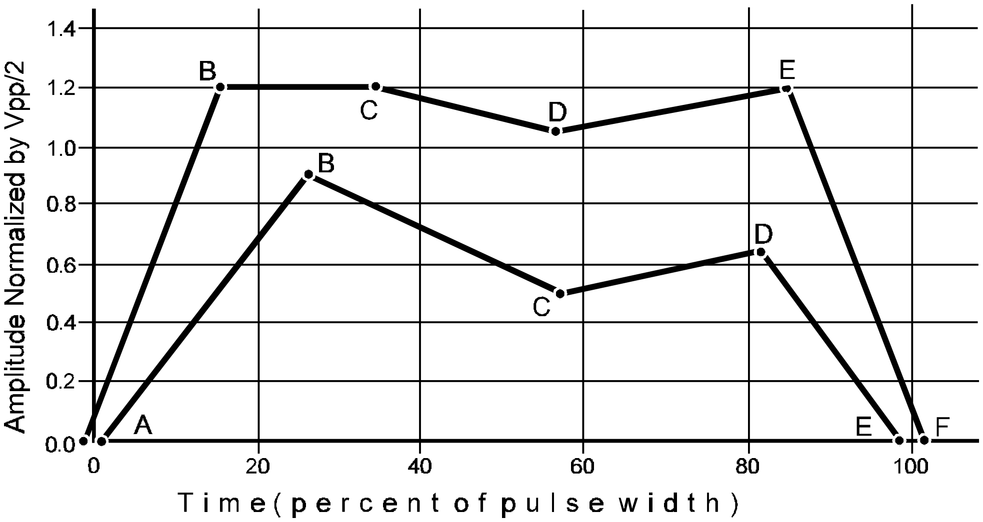

NOTE—See table 19 for point coordinates.

Figure 34—16 Mbit/s template for Data_one symbols waveform

Table 19—16 Mbit/s Data_one symbols template

Point	Upper time (%)	Upper amplitude	Lower time (%)	Lower amplitude
A	–1	0	1	0
B	15.5	1.2	26	0.9
C	34.5	1.2	57	0.5
D	56.5	1.05	81.5	0.65
E	85	1.2	99	0
F	101	0	—	—

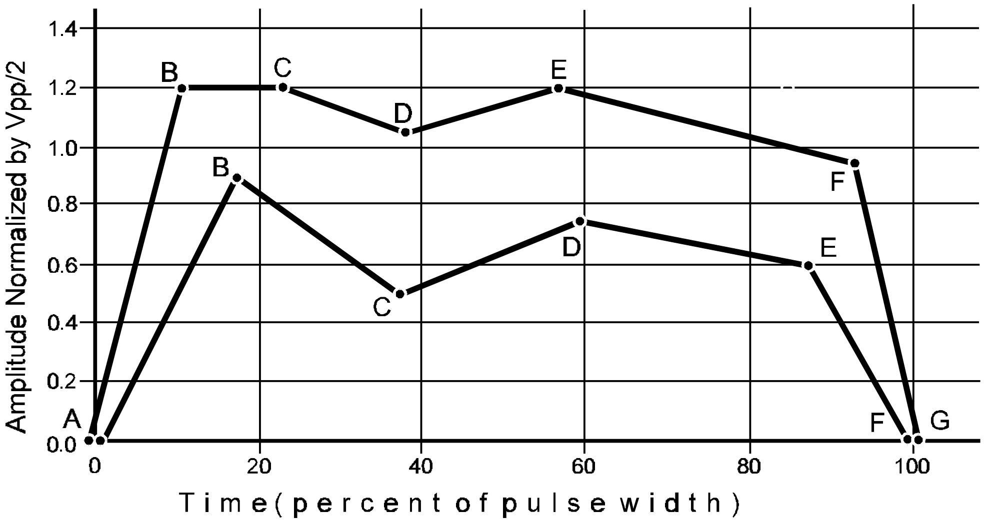

NOTE—See table 20 for point coordinates.

Figure 35—16 Mbit/s template for SD waveform

Table 20—16 Mbit/s SD template

Point	Upper time (%)	Upper amplitude	Lower time (%)	Lower amplitude
A	–0.5	0	0.5	0
B	10.5	1.2	17.5	0.9
C	23	1.2	37.5	0.5
D	38	1.05	59.5	0.75
E	57	1.2	87.5	0.6
F	93	0.95	99.5	0
G	100.5	0	—	—

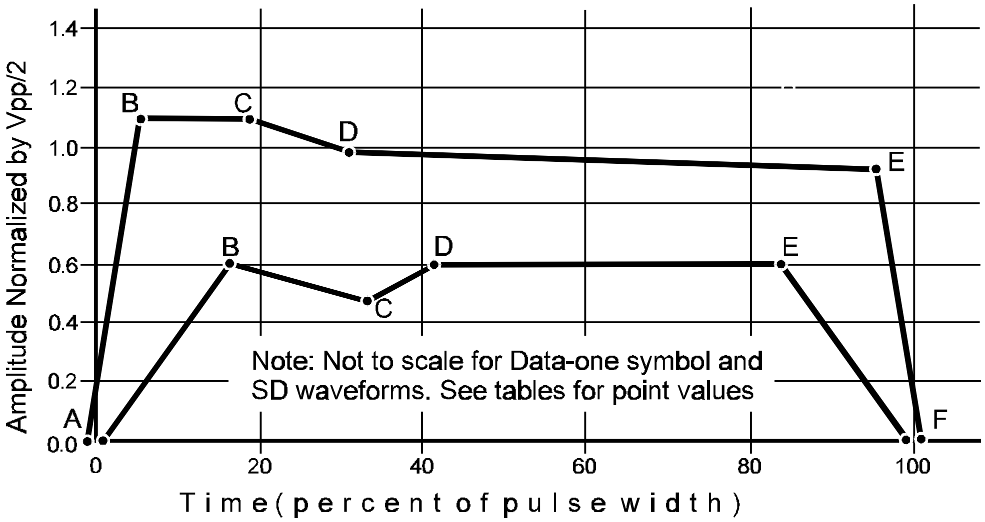

NOTE—See tables 21 through 23 for respective point coordinates.

Figure 36—4 Mbit/s template for Data_zero symbol, Data_one symbol, and SD waveforms

Table 21—4 Mbit/s Data_zero symbols template

Point	Upper time (%)	Upper amplitude	Lower time (%)	Lower amplitude
A	–1	0	1	0
B	5.5	1.10	16.5	0.6
C	19	1.10	33.5	0.48
D	31	0.98	41.5	0.6
E	93.5	0.92	84	0.6
F	101	0	99	0

Table 22—4 Mbit/s Data_one symbols template

Point	Upper time (%)	Upper amplitude	Lower time (%)	Lower amplitude
A	–0.5	0	0.5	0
B	2.5	1.10	8	0.6
C	10	1.10	16.5	0.52
D	15.5	0.98	22.5	0.64
E	100	0.88	90	0.68
F	100.5	0	99.5	0

Table 23—4 Mbit/s SD template

Point	Upper time (%)	Upper amplitude	Lower time (%)	Lower amplitude
A	–0.5	0	0.5	0
B	2	1.10	5.5	0.6
C	6.5	1.10	11	0.52
D	10.5	0.98	15	0.64
E	98.5	0.84	93.5	0.68
F	100.5	0	99.5	0

7.2.2.3 Differential output voltage

The measured value of the differential output voltage *Vpp* at TX-A and TX-B, as defined in 7.2.2.2, shall meet the specification for transmitter differential output voltage (TDOV) shown in table 24.

Table 24—Transmitter differential output voltage specifications

Cable type	4 Mbit/s	16 Mbit/s
TDCD (STP 150 Ω load)	3.0 – 4.5 V	3.0 – 4.5 V
TDCD (UTP 100 Ω load)	2.7 – 3.6 V	2.7 – 3.4 V

7.2.2.4 Transmitter return loss

The transmitter return loss (TXRL), which limits the differential reflected signal due to a differential signal incident upon the transmitter, is specified at the transmitter output TX-A and TX-B. The transmitter, which shall be powered up and transmitting any valid data pattern, shall provide a minimum TXRL, relative to a reference impedance of 150 Ω (STP) or 100 Ω (UTP), as shown in table 25.

Table 25—Transmitter return loss specifications

Ring speed	Frequency range	Minimum TXRL
4 Mbit/s	1 MHz to 6 MHz	14 dB
4 Mbit/s	6 MHz to 12 MHz	12 dB
16 Mbit/s	1 MHz to 6 MHz	14 dB
16 Mbit/s	6 MHz to 17 MHz	12 dB
16 Mbit/s	17 MHz to 25 MHz	8 dB

The differential signal incident upon the transmitter shall be a minimum of 0.5 V peak-to-peak (V_i) when measuring the return loss.

The TXRL can be defined in terms of a differential reflected voltage or transmitter impedance as:

$$\mathrm{TXRL} = 20\log\left|\frac{Z_t + Z_{\mathrm{ref}}}{Z_t - Z_{\mathrm{ref}}}\right| = 20\log\left|\frac{V_i}{V_r}\right|$$

where:

Z_t is the impedance of the transmitter
Z_{ref} is the reference impedance (150 Ω STP, 100 Ω UTP)
V_i is the differential voltage incident upon the transmitter
V_r is the differential voltage reflected from the transmitter

7.2.3 Receiver specifications

The PMC shall provide a receiver in accordance with the electrical specifications of this clause. The receiver includes all circuitry from the receive ports RX-A and RX-B to the transfer of data and timing information to the physical signaling symbol decoder and repeat path defined in figure 26. The receiver specifications are defined by the receiver jitter tolerance (7.2.3.1) and the receiver return loss (7.2.3.2).

7.2.3.1 Receiver jitter tolerance

Jitter tolerance (JTOL) is the ability of a receiver to receive frame data at a data rate that has a static frequency error caused by clocking source tolerances and a time-varying frequency error caused by accumulated jitter. Jitter tolerance is specified here for two types of test channels[17]:

a) A test channel (Test_Chan_B) containing attenuation proportional to the square root of the frequency and flat attenuation with no added noise for the JTOL measurement.

b) A test channel (Test_Chan_C) containing attenuation proportional to the square root of the frequency and flat attenuation with added high-frequency noise designed to simulate the eye closure caused by high frequency disturbances such as that caused by crosstalk for the JTOLX measurement.

These test channels are defined in 7.2.4.4. Interoperability requires that the jitter tolerance exceed the specifications in table 26. JTOL is measured for a PMC receiving frames 1250 octets long containing worst-case data patterns from an active monitor station with a conformant transmitter whose clock source has both a static and a modulated frequency error.

JTOL and JTOLX are measured by using a signal from a transmitter clocked by a controlled clock source with a static frequency error of up to ± 0.01% and a time-varying frequency error of variable amplitude. The time-varying frequency error is zero-mean square wave modulated at a modulation frequency of FM. The amplitude (peak) of the square wave modulation of the clock source (in ns/ns) is increased until the frame error rate exceeds 10^{-5}. The amplitude (peak) at which this occurs is the value of JTOL or JTOLX. JTOL and JTOLX shall meet the specifications shown in table 26 under the specified transmitter differential output voltages under any channel conformant with the test channel specifications referenced above and with worst-case static frequency error.

Table 26—Receiver jitter tolerance specifications

Parameter	4 Mbit/s	16 Mbit/s
Vpp (STP)	3.6–3.8 V	3.6–3.8 V
Vpp (UTP)	2.9–3.1 V	2.9–3.1 V
FM	4.67 kHz	9.17 kHz
|JTOL|	> 0.007 ns/ns	> 0.0172 ns/ns
|JTOLX|	> 0.0035 ns/ns	> 0.0082 ns/ns

---

[17] These test channels are defined in 7.2.4.4, table 32.

7.2.3.2 Receiver return loss

The receiver return loss (RRL), which limits the differential reflected signal due to a differential signal incident upon the receiver, is specified at the receiver input ports RX-A and RX-B of the MIC. The receiver shall provide a minimum RRL, relative to a reference impedance of 150 Ω (STP) or 100 Ω (UTP), as shown in table 27.

Table 27—Receiver return loss specifications

Ring speed	Frequency range	Minimum RRL
4 Mbit/s	1 MHz to 12 MHz	15 dB
16 Mbit/s	1 MHz to 17 MHz	15 dB
16 Mbit/s	17 MHz to 25 MHz	8 dB

The RRL can be defined in terms of a differential reflected voltage or receiver impedance as

$$\mathrm{RRL} = 20\log\left|\frac{Z_r + Z_{\mathrm{ref}}}{Z_r - Z_{\mathrm{ref}}}\right| = 20\log\left|\frac{V_i}{V_r}\right|$$

where

Z_r is the impedance of the receiver
Z_{ref} is the reference impedance (150 Ω STP, 100 Ω UTP)
V_i is the differential voltage incident upon the receiver
V_r is the differential voltage reflected from the receiver

7.2.4 Channel specifications for copper cabling

This clause provides the specifications for two types of channels for connecting conforming devices (stations, concentrators, repeaters, etc.) that will allow the devices to perform as specified in 7.1 and 7.2.3. The two channel types are a passive channel for use with passive concentrators (8.4) and an active channel for use with active concentrators (8.5).

Channel attenuation (7.2.4.2.1), signal-to-crosstalk noise (7.2.4.3.1), and channel impedances (7.2.4.2.2) are specified for both the passive and active channels. These specifications support devices operating exclusively at 4 Mbits/s data rate (only 4 Mbits/s), and devices that can operate at either a 4 or 16 Mbits/s or only 16 Mbits/s data rate (16/4 Mbits/s).

The passive and active channel specifications provide the basis for the test channels specified in 7.2.4.4 and the channel design guidelines in annex B.

7.2.4.1 Channel types

In order to provide channel specifications that permit stations to operate with both passive (8.4) and active (8.5) concentrators, both a passive and active channel are defined.

7.2.4.1.1 Passive channel

The passive channel is the transmission path between the MIC of a transmitting repeater or station and the MIC of the next receiving repeater or station, where the concentrator ports to which the stations are attached are passive. The passive channel, illustrated in figure 37, consists of all the components in the transmission path, which includes the lobe cables and cords, passive concentrator(s), and, where appropriate, the inter-concentrator trunk cable and connectors.

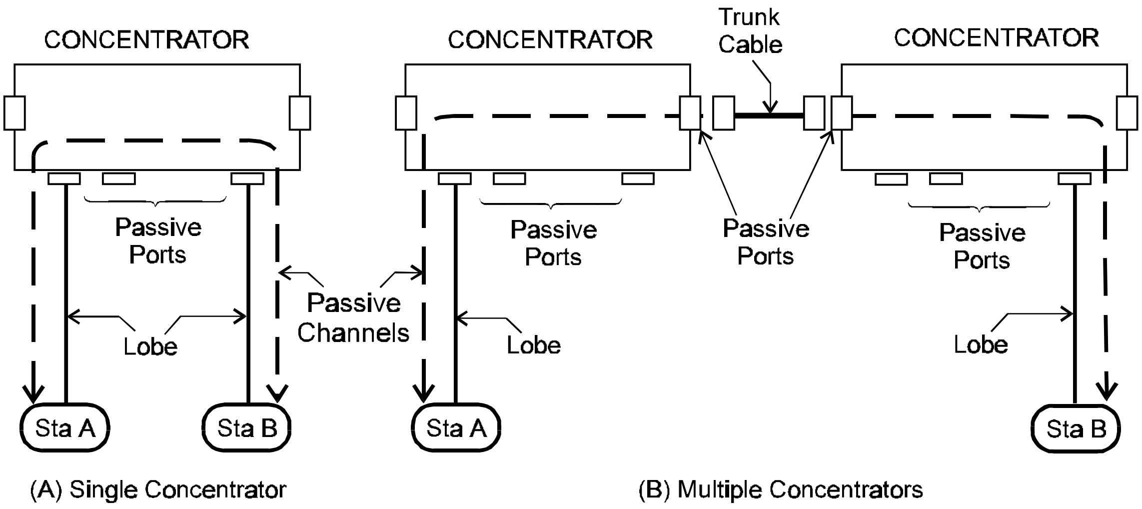

Figure 37—Passive channel[18]

7.2.4.1.2 Active channel

The active channel consists of all the components in the transmission path between the MIC of a station and the MIC of an active concentrator port or all of the components in the transmission path between the MIC of a concentrator ring out (RO) port that is active and the MIC of a concentrator ring in (RI) port that is active. The active channel is illustrated in figure 38.

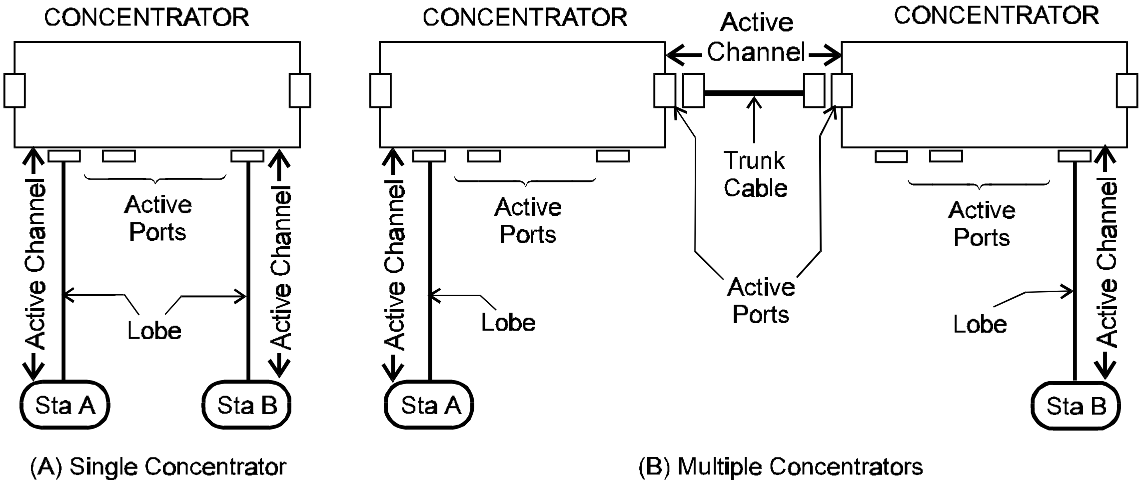

Figure 38—Active channel

---

[18] See figures B-2 and B-3 for channel design using backup path.

7.2.4.1.3 Channel implementation

The passive and active channels shall be implemented using copper shielded twisted pair (STP) cabling or copper unshielded twisted pair (UTP) cabling. STP and UTP shall meet the relevant clauses of ISO/IEC 11801 : 1995. For new installations utilizing STP, Type 1A cable is recommended. For new installations utilizing UTP, Category 5 is recommended. Each of the two channels includes all of the components in the transmission path between the MIC of the transmitting device and the MIC of the receiving device. These components include station cords, station wall outlet connectors, horizontal distribution cable, cross connect fields and associated jumper cords, concentrator patch cords, passive concentrators and inter-concentrator trunk cabling and connectors. This standard addresses both trunk and patch cables as specified in ISO/IEC 11801 : 1995. It is recommended that all UTP patch cords meet Category 5 specifications.

7.2.4.2 Transmission parameters

Two transmission parameters, attenuation and differential impedance, are specified for both the passive and active channels. Crosstalk is also an important characteristic for the channel. The allowed crosstalk for the channel as a whole is not specified directly, but is determined by the requirements for SCNR and NIR (see 7.2.4.3.1 and 7.2.4.3.2). Crosstalk for cables and components is specified in ISO/IEC 11801 : 1995; however, it is specified for individual components in the channel rather than for the channel as a whole.

7.2.4.2.1 Channel attenuation

Channel attenuation consists of the attenuation of the channel cable pairs, connector losses, reflection losses due to impedance mismatches between the various components of the channel, and the passive concentrator loss as measured by insertion loss.

The channel attenuation of a passive channel consists of the loss of all the channel components between a station MIC and the MIC of the next downstream station. As shown in figure 37(A) for a single concentrator, the attenuation of the passive channel consists of the loss between the Station A MIC and the concentrator port MIC, plus the loss between the two concentrator ports (concentrator loss), plus the loss between the concentrator port MIC and the Station B MIC. For the multiple concentrator example of figure 37(B), the attenuation consists of the loss between the Station A MIC and the concentrator port MIC, plus the concentrator loss for the first concentrator, plus the loss of the trunk cable between the two concentrators, plus the loss of the second concentrator, plus the loss between the second concentrator port MIC and the Station B MIC (refer to annex B for further explanation of how to determine the maximum length channel loss).

The channel attenuation of an active channel consists of the loss between a station MIC and its respective active concentrator port MIC, or the loss between an active RO port MIC of one concentrator and the active RI port MIC of a second concentrator. As shown in figure 38(A) for a single concentrator, the channel attenuation of each of the two active channels consists of the loss between the station MIC and its respective active concentrator port MIC. Figure 38(B) shows multiple concentrators with three active channels. The channel attenuation of the two station-to-concentrator active channels consists of the loss between the station MIC and its respective concentrator port. The channel attenuation of the third active channel, which interconnects the two concentrators, consists of loss between the first concentrator RO MIC and the second concentrator RI MIC.

The maximum attenuation of each channel shall not exceed the losses in table 28. The losses shall be met when the channel type channel is terminated in 100 Ω ± 1% for UTP and 150 Ω ± 1% for STP.

Table 28—Channel attenuation specifications

Channel type	Channel attenuation Only 4 Mbit/s	 16/4 Mbit/s
Passive	< 19 dB @ 4 MHz	< 19 dB @ 16 MHz
Active	< 19 dB @ 4 MHz	< 16 dB @ 16 MHz

The passive and active channel attenuation consists primarily of cable loss whose attenuation in decibels and non-linear portion of the phase vary as the square root of frequency. In addition, the passive channel may contain flat loss, a component of attenuation that does not vary with frequency. The passive channel attenuation shown in table 28 shall consist of no more than 2 dB of flat loss for all data rates. Note that the maximum attenuation shown in table 28 applies for channels that just meet the SCNR requirements of table 30. Channels possessing higher SCNR may be able to support slightly higher attenuation. This is a subject for future study.

7.2.4.2.2 Channel differential impedance

The magnitude of the channel differential input and output impedance, for both the passive and active channels, when using STP and UTP media, shall meet the specifications of table 29 when measured according to ASTM D4566-1994 [B1].

Table 29—Channel differential impedance specifications

Media type	Differential impedance	Frequency range Only 4 Mbit/s	 16/4 Mbit/s
UTP	100 ± 15 Ω	1–12 MHz	1–25 MHz[19]
STP	150 ± 15 Ω	1–12 MHz	1–25 MHz

All individual cables, cords, and components in the channel shall be of the same cable type and nominal impedance and shall meet the specifications of table 29 in order to minimize distortion due to reflections at impedance discontinuities.

The impedance ranges specified in table 29 are limits on the nominal characteristic impedance of the media, and are not indicative of the allowable variation in characteristic impedance within a section or within any channel. In addition, it is recommended that in order to limit the impedance variations within a channel, each permanently installed cabling run between telecommunications closets and between a telecommunications closet and a wall outlet should consist of a single homogeneous cable run without splices or intermediate connectors. It is recommended that the return loss of the cabling meet the return loss specified for the class C link in ISO/IEC 11801 : 1995.

---

[19] Category 3 is only specified to 16 MHz but it is assumed that it can be extrapolated to 25 MHz.

7.2.4.3 Noise environment

The level of noise in the passive and active channels must be limited to ensure the system error rate stated in 7.0 is provided. The noise environment associated with the passive and active channels consists primarily of two contributors: (1) crosstalk noise due to signals within the transmission components, and (2) noise induced into the channels from sources external to the transmission components used to provide the channels. Crosstalk noise will be present in all channels and is minimized by using channel components with high crosstalk loss. The primary crosstalk noise in a channel is the result of the transmitter Manchester signals coupling into the receive channel, which appear at the input of a receiver along with the data signals. The level of the crosstalk noise in a channel depends on the level of the transmitter signal(s) and the crosstalk loss between the cable pair(s) carrying the transmitter signal(s) and the cable pair assigned to the receiver.

Channel crosstalk noise specifications are based on the use of a sheathed cable carrying token ring signals for a single station. The sharing of the sheath with other services is beyond the scope of this part of ISO/IEC 8802. While the use of multi-pair cables to carry multiple services is not addressed here, the following guidance is provided.

The result of the total crosstalk noise from all disturbers in a multi-pair sheathed cable should not exceed the specifications of the SCNR of 7.2.4.3.1. For example, sharing of services with analog telecommunications devices may result in unacceptable levels of crosstalk noise.

For guidance on external noise limits and cabling recommendations, see 7.2.4.3.3 and ISO/IEC 11801 : 1995.

7.2.4.3.1 Signal-to-crosstalk noise ratio (SCNR)

To ensure the receiver of 7.2.3 will operate properly in the presence of crosstalk noise generated by the Manchester transmitter signals of 7.2.2, the passive and active channels shall provide the minimum peak SCNR at the channel output shown in table 30 for sinusoidal signals at the data rate.

Table 30—Channel SCNR specifications

Channel type	SCNR Only 4 Mbit/s	 16/4 Mbit/s
Passive	> 15.5 dB	> 13.5 dB
Active	> 15.5 dB	> 12.0 dB

7.2.4.3.2 NEXT loss to insertion loss ratio (NIR)

To meet the SCNR of 7.2.4.3.1 and take into account the transmitter output variation specified in 7.2.1 (3.5 dB for STP and 2 dB for UTP), the channel shall provide an NIR, meeting the specifications in table 31. NIR is defined as the ratio of the minimum Adjusted NEXT Loss over a frequency range to the insertion loss at the reference frequency F_{ref}. For 16/4 and 16 Mbit/s operation, the frequency range is 1–25 MHz and F_{ref} is 16 MHz. For 4 Mbit/s operation, the frequency range is 1–12 MHz and F_{ref} is 4 MHz. The Adjusted NEXT Loss[20] is defined as the NEXT loss of the channel in dB (at a frequency F) plus $15\log(F/F_{\mathrm{ref}})$. The complex distributed nature of the couplings between pairs in cables typically causes channel crosstalk loss characteristics to vary rapidly with frequency. Hence, it is inadequate to characterize the crosstalk loss at a single frequency; therefore, NEXT loss is measured over a frequency

---

[20] Although this adjustment is based on the characteristics of cables (see ISO/IEC 11801 : 1995), it is intentionally used here to specify a limit on the characteristics of the entire channel.

band. Figures 39, 40, and 41 show the relationship between channel Adjusted NEXT Loss and insertion loss at F_{ref}.

Table 31—Channel NEXT loss to insertion loss ratio (NIR) specifications

Channel	Channel NEXT Loss to NIR			
	Only 4 Mbit/s F_{ref} = 4 MHz		16/4 Mbit/s F_{ref} = 16 MHz	
	STP	UTP	STP	UTP
Passive	> 19.0 dB	> 17.5 dB	> 17.0 dB	> 15.5 dB
Active	> 19.0 dB	> 17.5 dB	> 15.5 dB	> 14.0 dB

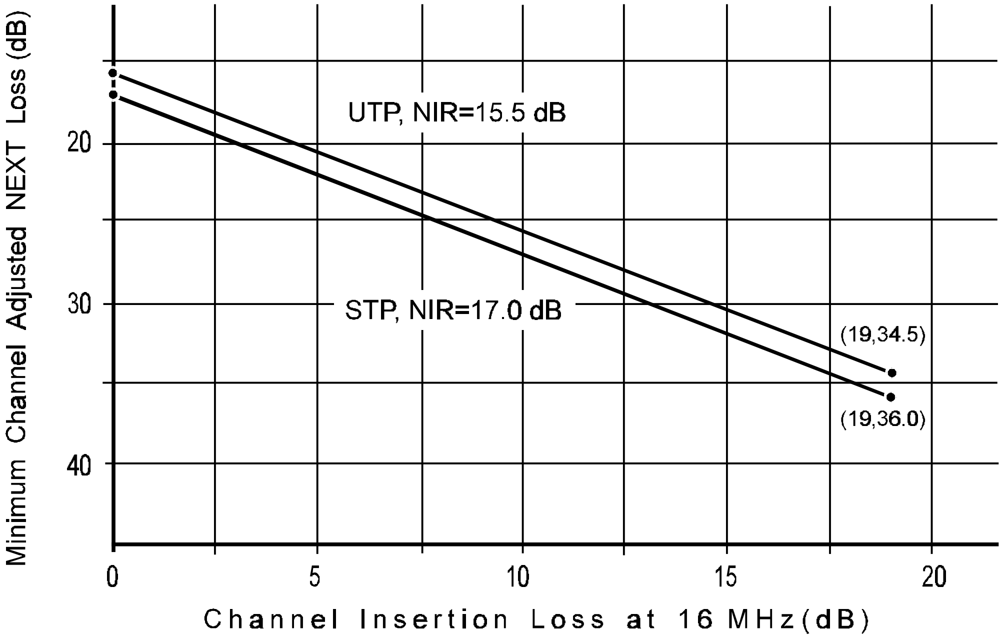

Figure 39—Minimum Adjusted NEXT Loss vs channel insertion loss to meet 16/4 Mbit/s passive channel NIR

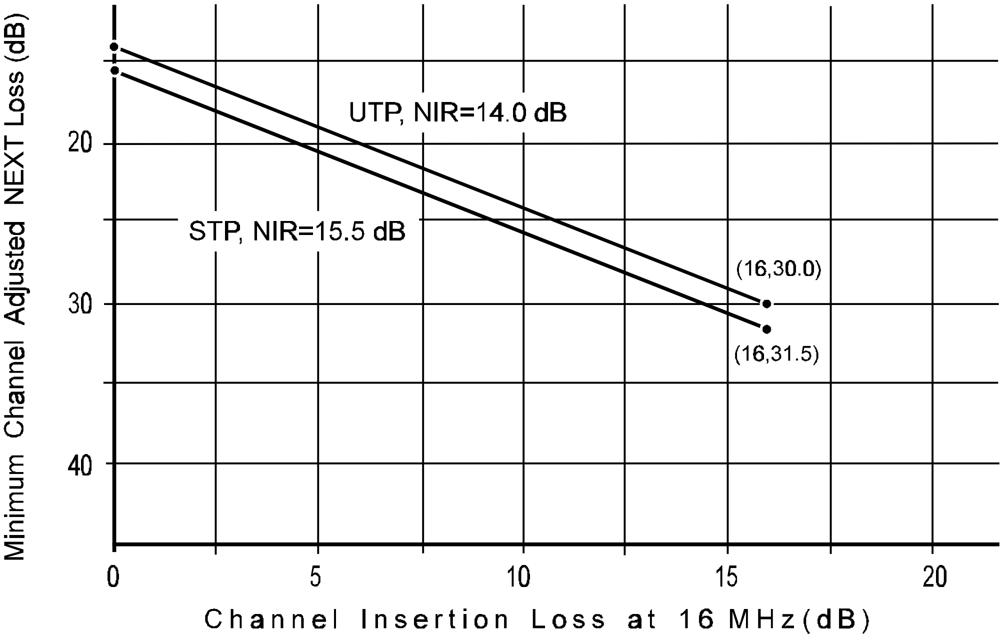

Figure 40—Minimum Adjusted NEXT Loss vs channel insertion loss to meet 16/4 Mbit/s active channel NIR

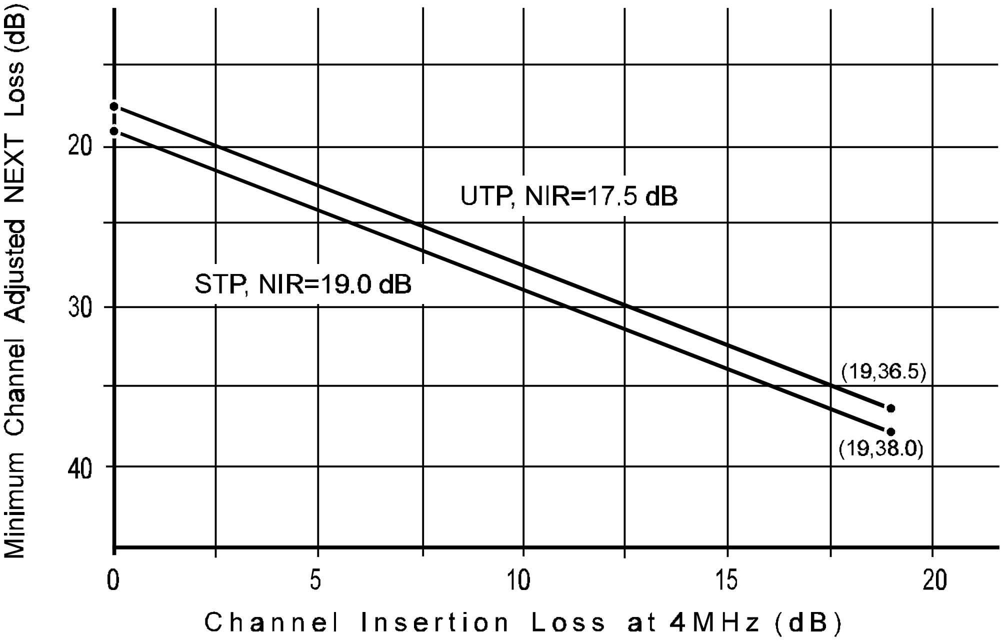

Figure 41—Minimum Adjusted NEXT Loss vs channel insertion loss to meet 4 Mbit/s only passive and active channel NIRs

7.2.4.3.3 External noise recommendations

System characterization of noise from sources external to the channel (i.e., all noise except crosstalk noise) and its impact on the system error rate is a subject for future study. External noise coupling into a channel is dependent on the type of cabling system, cabling environment, and the station and concentrator hardware implementation. The following is a conservative recommendation for externally generated noise for both passive and active channels. The effects of external noise can be neglected when the peak noise at the input to the station receiver is at least 30 dB below the peak level of the data signal. For the minimum transmitter signal level specified in 7.2.2 and the maximum channel attenuation of 7.2.4.2.1, the effects of external noise can be neglected as long as the peak external noise is within the shaded region of figure 42 for STP and UTP. If external noise is outside the shaded region of figure 42, it should not be ignored. In such environments special precautions, such as increasing the minimum acceptable NIR to budget for the effects of external noise, may be warranted to assure that the system error rate is not compromised. Specific recommendations for high external noise environments are beyond the scope of this standard. External noise can be kept within acceptable limits in all applicable environments by the proper choice and installation of the cabling system. STP and UTP Category 5 cables are extremely effective in limiting system errors due to external noise. UTP Category 4 cable is somewhat less effective, and UTP Category 3 is the least effective UTP cable, or most susceptible to coupling external noise signals into the channel. For very high external noise environments, 150 Ω STP cable may be required. For most installations, external noise presents no problem when any of these cable systems are used.

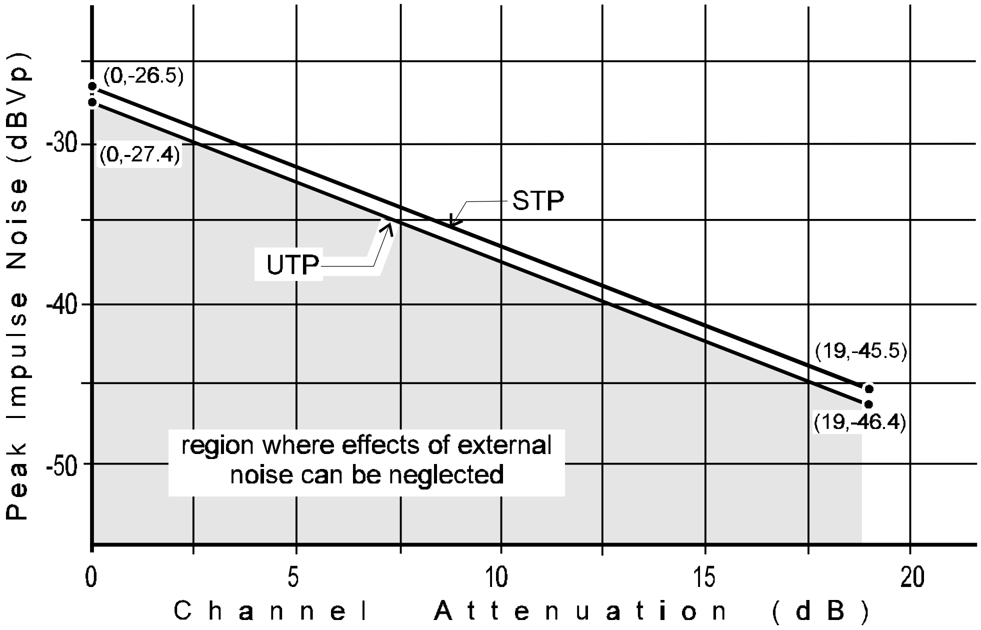

Figure 42—Impulse noise level recommendations for passive and active channels

7.2.4.4 Test channels

Test channels are used to allow the specification of various performance parameters (AJ, FAPS, DFAPS, JTOL, JTOLX) of the PHY under controlled conditions. The generic test channel is a channel containing only attenuation proportional to the square root of frequency and attenuation that is flat across frequency. The total channel attenuation of the test channel may be any value less than TCATT. For active channels this channel attenuation shall be composed entirely of root-frequency attenuation and no flat loss. For passive channels, the total channel attenuation may consist of root-frequency attenuation and up to 2 dB of flat loss. The channel should have minimal crosstalk; however, noise that simulates crosstalk or other noise sources may be added to the channel in a controlled fashion as specified in the next paragraph. The impedance of the test channel shall be within the range from 145–155 Ω for an STP channel and from 95–105 Ω for a UTP channel.

For test channels that include added noise, an external source is used to inject a sinusoidal noise at a frequency F_n into the channel in a manner that causes minimal interference in the characteristics of the channel except for the noise injected. The amplitude of the interference is adjusted to establish a *signal*-to-*interference* ratio of SIR dB at the MIC RX-A and RX-B. The *signal* level is the peak-to-peak amplitude of the fundamental sine wave component of the zeros pattern waveform and the *interference* level is the peak-to-peak amplitude of the sinusoidal noise signal.

Table 32 specifies the values of these parameters for each test channel. These test channels are used in the specifications in 7.1.1 and 7.2.3.1.

Table 32—Test channel specifications

Data rate	Parameter	Test_Cha_A (AJ, FAPS, DFAPS)	Test_Chan_B (JTOL)	Test_Chan_C (JTOLX)
4 Mbit/s	TCATT Fn SIR	13 dB — —	19 dB — —	19 dB 3 MHz 12 dB
16 Mbit/s	TCATT Fn SIR	19(13) dB — —	19(16) dB — —	19(16) dB 12 MHz 12(12) dB

Note that active concentrator PHYs are only tested up to TCATT and SIR conditions shown in parentheses. The channel attenuations for 4 and 16 Mbit/s are referenced to the frequencies 4 and 16 MHz respectively. Test_Chan_A TCATT is lower than that of B and C to allow simpler and less expensive equalizer design at 4 Mbit/s. Since accumulated jitter is the sum of the jitter introduced by all the PHYs in the ring, advantage can be taken of the typical distributions of channel attenuations that result in most of the channels having attenuations well below 13 dB. The presence of a few channels with higher jitter will not substantially affect the accumulated jitter. The same approach cannot be taken with jitter tolerance (test channels B and C) since the presence of even one PHY with poor jitter tolerance can cause errors in the ring.

7.2.5 Medium-dependent interface specifications

This clause defines the station medium-dependent interface MIC_S and MIC_U for attaching the station to STP and UTP media, respectively. The manufacturer is responsible for defining the location of the MIC on the manufacturer’s device and cable. Possibilities include defining the MIC at the device card edge or at the end of a patch cord that attaches to the device. In all cases, the connector shall meet the mechanical specifications given in the following subclauses.

7.2.5.1 MIC_S (STP medium interface connector)

Hermaphroditic connectors meeting the mechanical specifications in the following clauses shall be used as the mechanical interface between the PMC and the STP channel. It is recommended that the connectors also meet the electrical characteristics specified in 7.2.5.1.2.

7.2.5.1.1 MIC_S mechanical specifications

Figure 43 shows an isometric view of the medium interface connector as it would be oriented when it is wall-mounted. It has four signal contacts with a ground contact and is hermaphroditic in design so that two identical units will mate when oriented 180 degrees with respect to each other.

7.2.5.1.2 MIC_S mechanical characteristics

a) Contact force 0.5–1.0 N
b) Minimum insertions without a failure > 500
c) Surface treatment (compatible with the following):
    1) Point-of-pin contact plating with 3 µm of hard gold
    2) Point-of-shield contact plating with 5 µm of tin

The mechanical mateability of the connector is subject to standardization by IEC.

7.2.5.1.3 MIC_S electrical characteristics

It is recommended that the MIC_S meet the specifications in ISO/IEC 11801 : 1995.

7.2.5.1.4 MIC_S medium interface connector–contact detail

Figure 44 shows the details of the signal and ground contacts. When the connector is disconnected, pin R shall be shorted to pin O and pin G shorted to pin B for automatic looping capability. Only those dimensions that are essential to mating are shown.

7.2.5.1.5 MIC_S medium interface connector–locking mechanism detail

Figure 45 shows the locking mechanism of the connector. Only those dimensions that are essential to mating are shown.

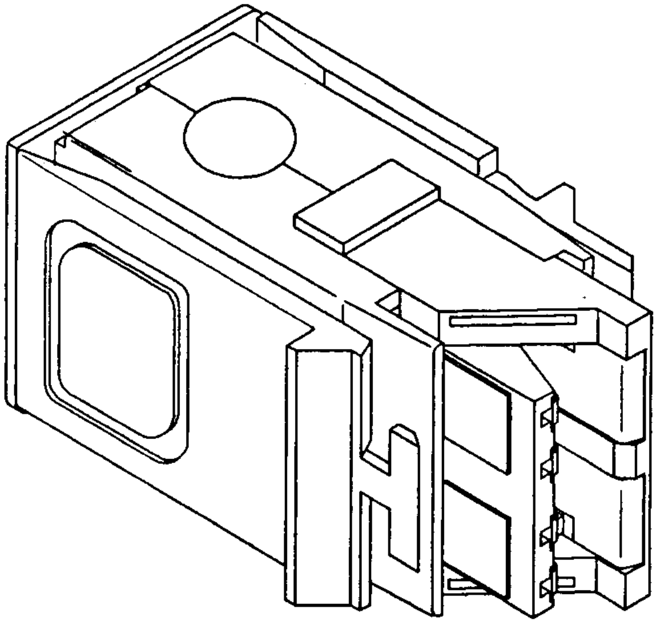

Figure 43—MIC_S—isometric view

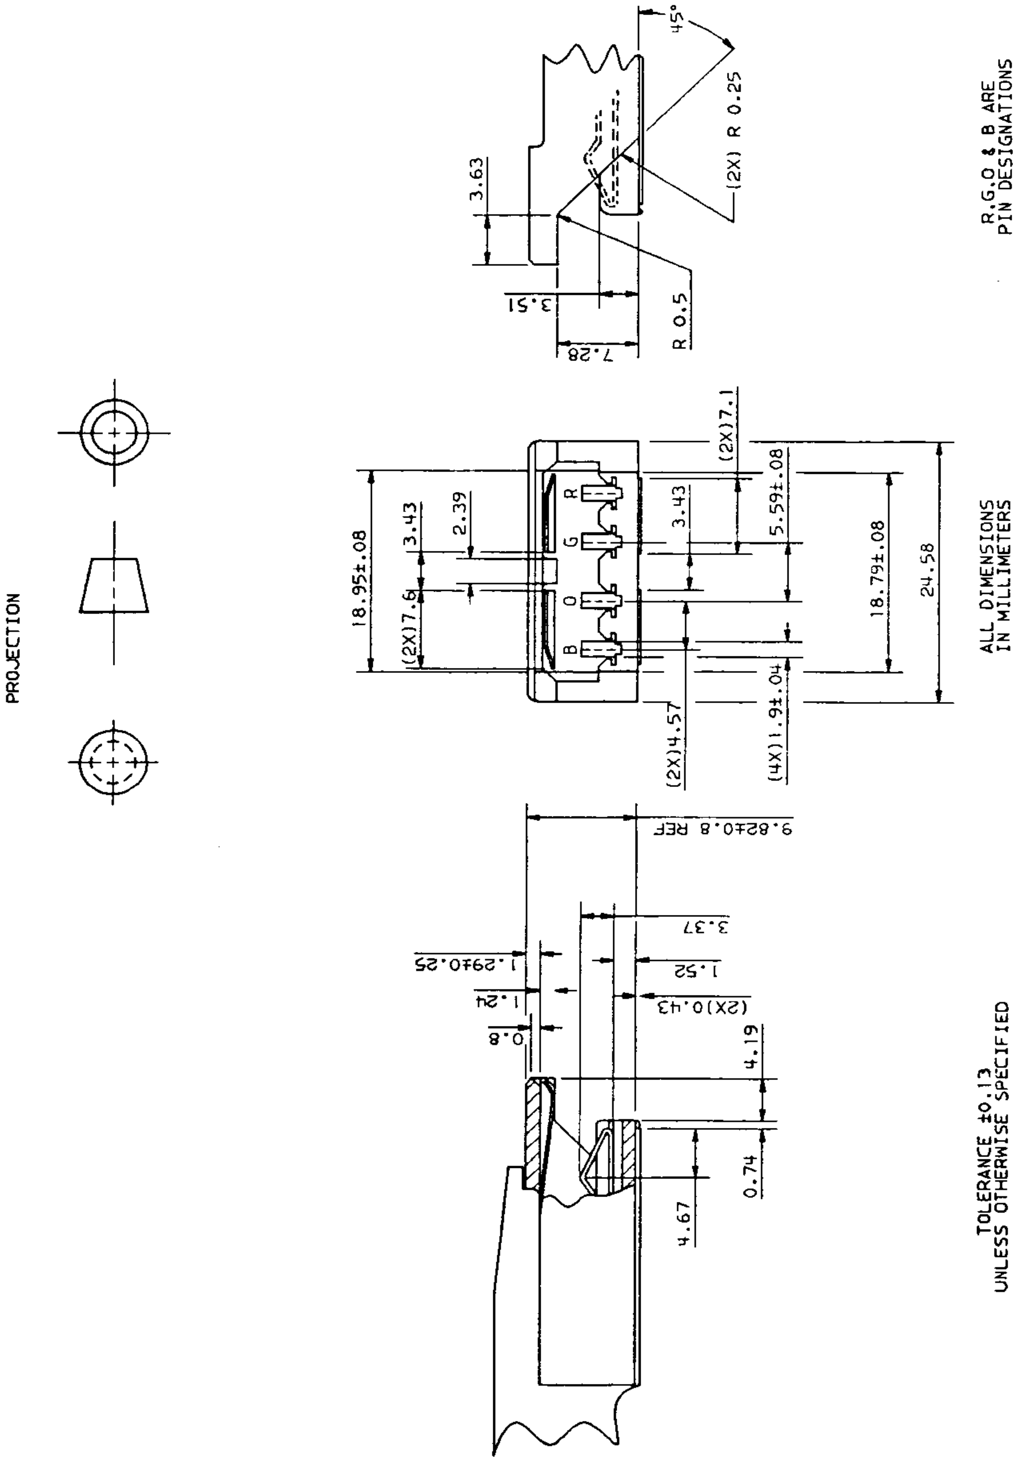

Figure 44—MIC_S—contact detail

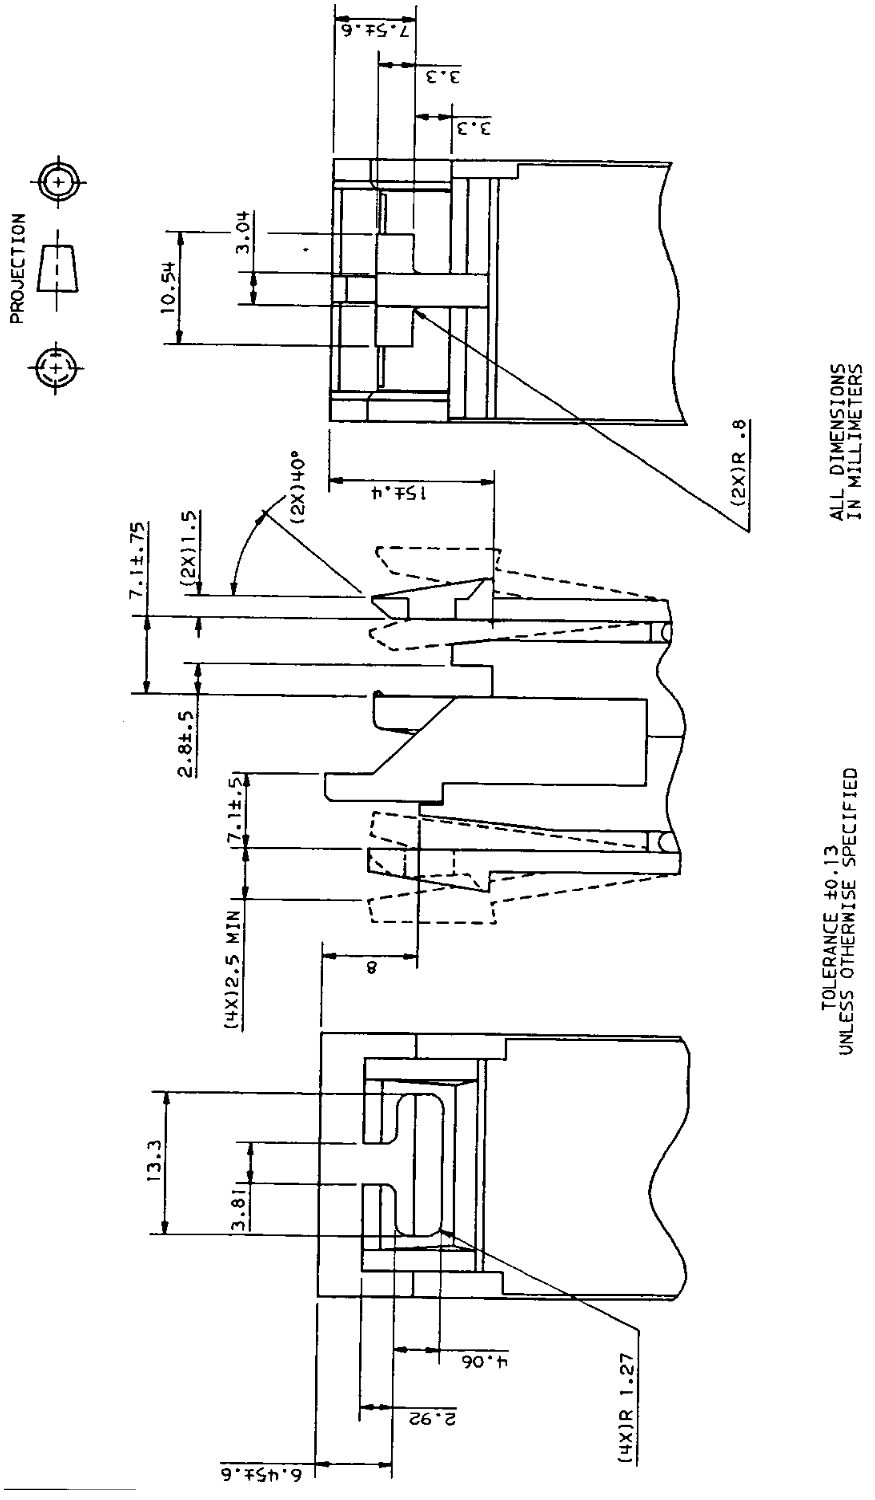

Figure 45—MIC_S—locking mechanism detail

7.2.5.2 MIC_U (UTP medium interface connection)

Either shielded or unshielded eight-pin connectors meeting the mechanical specifications in the following clause shall be used as the mechanical interface between the PMC and the UTP channel.

7.2.5.2.1 MIC_U mechanical specifications

The eight-pin connectors shall meet the mechanical specifications of IEC 603-7 : 1990. Either the plug or jack form of the connector may be used.[21] The plug and jack are illustrated in figure 46.

Unlike the STP connector, the UTP connector is not self-shorting. Other means of ensuring ring integrity in the event of an unplugged cable are necessary.

7.2.5.2.2 MIC_U electrical characteristics

It is recommended that the MIC_U meet the Category 5 electrical specifications defined in ISO/IEC 11801 : 1995.

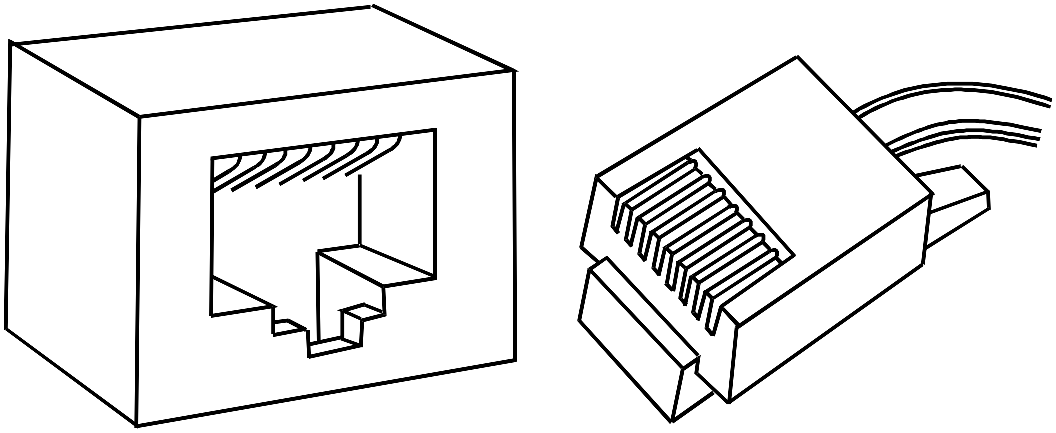

Figure 46—Illustration of MIC_U jack and plug

---

[21] For purposes of conformance testing, the jack form of the connector shall be used.

8. Concentrator specifications

This clause defines the electrical, mechanical, and functional specifications for token ring concentrators. Two types of concentrators are specified, a passive concentrator and an active retiming concentrator. Both concentrator types provide the following functions:

a) Into a ring couple stations attached to lobe cabling to the ring trunk.
b) Interconnect multiple trunk coupling units (TCUs) to form a concentrator to support a star-wired network topology.
c) Support message traffic at data signaling rate of 4 or 16 Mbit/s.
d) Provide for station confirmation of TCU presence and support of station detection of an open wire condition and certain short-circuit conditions in the lobe cabling.
e) May include ring in and ring out trunk attachments for serial connection of multiple concentrators topology over the trunk cable.

The need for specifying two types of concentrators, passive and active retiming, is due to the relationship between the ring speed and the cabling plant. The passive concentrator is a simple device aimed at installations where repeating elements are not required to support channel drive distance. In the case of the passive concentrator, the channel losses include, at a minimum, the two-directional lobe cabling loss, plus the concentrator loss, plus the losses of any additional passive devices in the channel. The active concentrator performs an imbedded repeater function in the lobe port’s data path thereby providing ring segment boundaries at the concentrator lobe port connector (CMIC). Refer to figures 37 and 38 for examples of passive and active channels, respectively. Additional information relative to the relationship that exists between the cabling plant, the ring speed, and the type of concentrator is given in the annex B.

Refer to annex E for recommendations relative to telephony voltages and network safety.

8.1 Concentrator lobe port

The concentrator shall provide a connector (CMIC) to connect a lobe cable from a station to the associated TCU. Two types of connectors are defined; the CMIC_S connector for attaching STP lobe cabling and the CMIC_U connector for attaching UTP lobe cabling. This subclause also defines the transmission path between the concentrator CMIC and the station MIC.

8.1.1 Concentrator lobe connector

The two types of concentrator lobe port connectors are as follows:

a) The CMIC_U connector, used for attaching UTP lobe cabling, shall meet the required specifications for the jack form of the connector in 7.2.5.2.
b) The CMIC_S connector, used for attaching STP lobe cabling, shall meet the required specifications of 7.2.5.1. Table 33 shows the signal and contact assignments for the two connectors.

Table 33—C_MIC contact and signal assignments

CMIC_U contact	CMIC_S contact	Concentrator signal
1	—	Not used by this standard
2	—	Not used by this standard
3	B	LRX-A
4	R	LTX-A
5	G	LTX-B
6	0	LRX-B
7	—	Not used by this standard
8	—	Not used by this standard

The interchange of LTX-A with LTX-B and the interchange of LRX-A with LRX-B will not affect functionality, and is permitted by this standard.

The concentrator receive signal ports (LRX) at the concentrator lobe port connector CMIC_U or CMIC_S are connected via the UTP or STP lobe cabling to the station transmit signal ports (TX) at the station's MIC_U or MIC_S. Similarly, the concentrator transmit signal ports (LTX) are connected to the station receive signal ports (RX). These connections are depicted in figure 47.

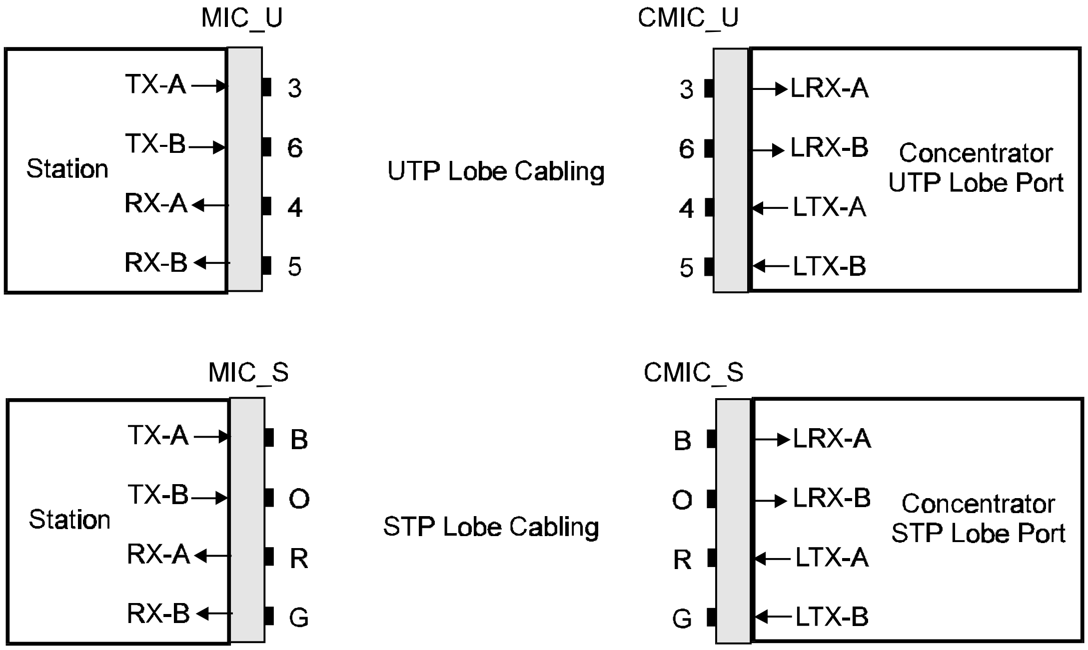

Figure 47—Concentrator UTP and STP lobe connections

8.1.2 Concentrator lobe port indicators

If a visible indicator(s) is provided on a concentrator to indicate the status of a port, it is recommended that the color be green and that the indicator be labeled appropriately. It is further recommended that the indicator be ON when the port has detected phantom current and the attached station is inserted. Guidance for the color and meaning of displays is contained in IEC 73 : 1991.

8.2 Concentrator trunk ports

The concentrator may provide trunk port connections for attaching the concentrator to the trunk cable. The specifications of this clause are based on the assumption that trunk ports exist or that the appropriate internal connections are maintained. At the trunk ports, if alternate media attachment is provided, then these points must be ring segment boundaries. Trunk ports that are not ring segment boundaries will be referred to hereinafter as passive trunk ports. The connectors, referred to as *ring in MIC* (RIMIC) and *ring out MIC* (ROMIC), permit concentrators to be cascaded to expand the size of a ring; RO of one concentrator is connected to RI of the next concentrator. Two connectors are specified for each type: RIMIC_S and ROMIC_S for use with STP media, and RIMIC_U and ROMIC_U for use with Category 5 UTP media. Also, the main ring transmission path and the back-up transmission path through the concentrator are defined.

8.2.1 Concentrator trunk connector

Two concentrator trunk port connectors, if provided, for attaching the concentrator to the trunk cable, shall be as follows:

a) The RIMIC_S and ROMIC_S connectors, used for attaching STP trunk cabling, shall meet the required specifications for the connector in 7.2.5.1.

b) The RIMIC_U and ROMIC_U connectors, used for attaching UTP trunk cabling, shall meet the required specifications for the jack form of the connector in 7.2.5.2. Table 34 shows the signal and contact assignments for the four connectors.

Table 34—U_MIC contact and signal assignments

ROMIC_U/ RIMIC_U contact	ROMIC_S/ RIMIC_S contact	RIMIC signal	ROMIC signal
1	—	Not used by this part of ISO/IEC 8802	
2	—	Not used by this part of ISO/IEC 8802	
3	B	TXBACKUP-A	RXBACKUP-A
4	R	RXMAIN-A	TXMAIN-A
5	G	RXMAIN-B	TXMAIN-B
6	0	TXBACKUP-B	RXBACKUP-B
7	—	Not used by this part of ISO/IEC 8802	
8	—	Not used by this part of ISO/IEC 8802	

The interchange of RIMIC signals RXMAIN-A with RXMAIN-B (TXBACKUP-A with TXBACKUP-B) and the interchange of ROMIC signals TXMAIN-A with TXMAIN-B (RXBACKUP-A with RXBACKUP-B) will not affect functionality, and is permitted by this part of ISO/IEC 8802.

Within the concentrator, the main ring signal transmission path shall be from RIMIC signals RXMAIN-A and RXMAIN-B, through the TCUs, and out ROMIC signals TXMAIN-A and TXMAIN-B. If a station attached to a TCU is inserted, the path within that TCU shall be altered as follows:

a) The incoming signal lines (RXMAIN-A and RXMAIN-B from the RIMIC or LRX-B and LRX-A from an upstream TCU) are connected to the TCU lobe CMIC output contacts LTX-A and LTX-B.
b) The CMIC input contacts LRX-B and LRX-A, receiving signals from the lobe, are connected to either the next downstream TCU or to ROMIC signals TXMAIN-A and TXMAIN-B.

Also provided as part of the trunk interface is a back-up path from ROMIC signals RXBACKUP-B and RXBACKUP-A to RIMIC signals TXBACKUP-B and TXBACKUP-A. All connections described in this subclause are illustrated in figure 48.

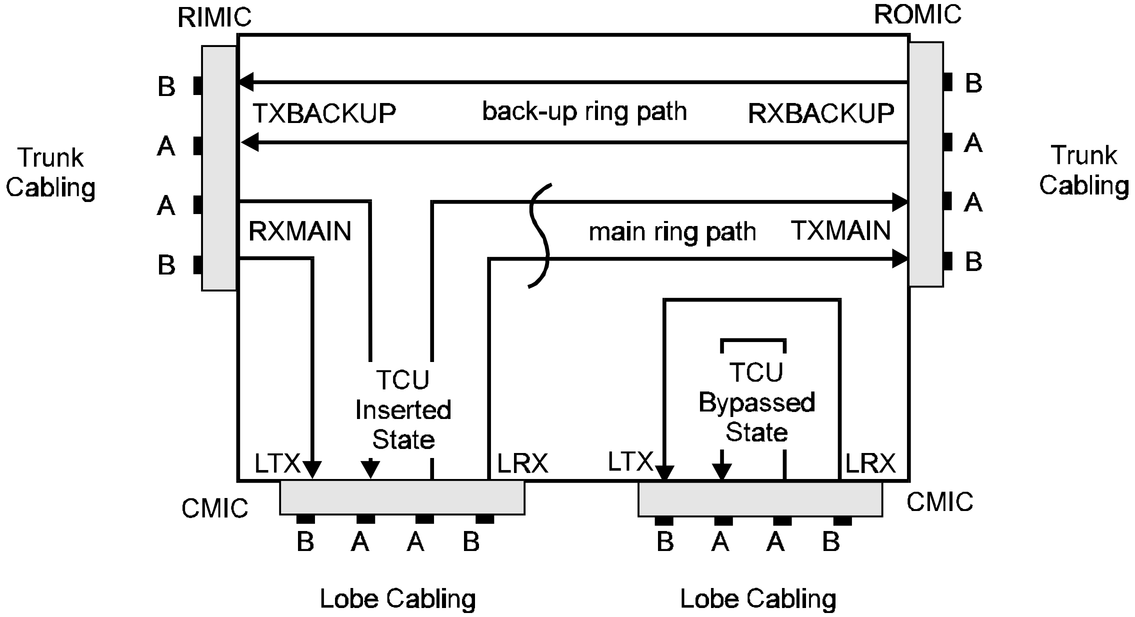

Figure 48—Example of transmission path through a concentrator

The concentrator shall provide a means, at both the RIMIC and ROMIC, by which the main ring path can be looped to the backup path in the absence of an attached trunk cable. The means can be either external or internal to the concentrator and connects RXMAIN-A to TXBACKUP-B (TXMAIN-A to RXBACKUP-B) and RXMAIN-B to TXBACKUP-A (TXMAIN-B to RXBACKUP-A).

8.3 Ring access control function

The concentrator shall provide a ring access control function (RAC) on each lobe port to receive requests from the station for ring insertion and ring bypass, as specified in 7.2.1. The RAC shall receive the station ring insertion and ring bypass requests from the phantom signaling channel as shown in figure 29. Each STP or UTP lobe port shall provide two independent dc isolated paths for the station's phantom signaling. The leakage resistance between the two phantom paths shall be at least 250 kΩ. One path shall be between CMIC signal pin LRX-A and LTX-A. The other path shall be between CMIC signal pin LRX-B and

LTX-B. The mechanism that controls ring insertion and bypass shall be provided on both of the paths or on one of the paths with the other path containing the load balancing circuitry.

8.3.1 Concentrator ring insertion and ring bypass

When a ring insertion request is received by the concentrator from a station, the RAC shall, unless instructed otherwise by higher level management functions,[22] insert the station into the ring. When the concentrator receives the ring bypass request, the RAC shall remove the station from the ring. The station shall be inserted into the ring and bypassed from the ring when the phantom signal from a compliant station, within the required power and no power regions, respectively, as specified in 7.2.1.1, is presented at the CMIC after a maximum lobe cable dc loop resistance of 100 Ω.

In the ring insert state, the RAC shall connect the lobe’s LTX signals to either the RI port receive signals (RXMAIN) or an upstream port’s LRX signals, and the lobe’s LRX signals to either the RO port transmit signals (TXMAIN) or a downstream lobe’s LTX, thereby including the station in the main ring path as shown in figure 48. When the concentrator port is in the ring bypass state, RAC shall connect the CMIC signal pins, LRX-A to LTX-A and LRX-B to LTX-B, to provide a loop back path to the station, which permits the station to perform off-line loop back tests.

8.3.2 Ring insertion/bypass timing

When the concentrator port TCU switches between the ring insertion and ring bypass modes, the port RAC shall ensure that the ring trunk circuit is open no more than 5.0 ms. The time from receipt of request to completion of ring insertion shall be a maximum of 5 s [23] or ring bypass a maximum of 190 ms.[24]

The ring bypass shall meet the above requirements independent of when and how long power was previously applied to the lobe port. To provide protection against phantom glitches causing momentary ring insertions, it is recommended that the TCU require that the ring insertion request be present for at least 50 ms[25] before the TCU switches to the insert state.

8.3.3 Concentrator ring access control loads

The loads presented by each of the two phantom paths shall have and maintain a resistance between 2.9 kΩ and 5.3 kΩ within 5 s [26] after ring insertion until station removal. Given equal and valid phantom voltages, the lower load resistance shall be within 15% of the higher load resistance and the lower load current shall be within 15% of the higher load current.

---

[22] Higher level management function is referring to a concentrator management function that will not be described or specified herein but may be present in a concentrator and is permitted to deny the ring insertion request. Care must be taken by the management function when denying a station’s ring insertion request to ensure the effect on the station’s MAC protocol does not effect the overall ring performance. It is further recommended that if the higher level management function detects a faulty station attachment and denies access it should indicate the possible fault to the user or management function.

[23] This requirement is constrained by the MAC timer insert delay (TID) and the MAC timer join ring (TJR).

[24] This requirement is constrained by the specifications of 7.2.1.1 and the MAC timer remove wait (TRW). This also requires that the concentrator prevent internal energy storage devices (e.g., capacitors) from discharging into the station phantom circuits as the phantom voltage drops in a transition from “power” 0 to “no power”. This is typically accomplished by using diodes to block this discharge.

[25] This recommendation is constrained by 7.2.1.1 and to conditions seen in earlier versions of adapters not compliant with this standard.

[26] This requirement is constrained by the MAC timer wire fault delay (TWFD).

8.4 Passive concentrator

This clause defines the electrical characteristics of the passive concentrator port. The passive concentrator contains no active elements in the lobe path and is included as one of the components in the channel. The passive concentrator may be used where the cabling plant is such that the signal's degradation due to channel impairments are within the drive distance limits supported by a token ring station. Passive concentrators may contain repeaters at the trunk ports to provide ring segment boundaries at the concentrator trunk port connectors. Implementors shall include in the published specification for the passive concentrator whether the RI and RO ports represent ring segment boundaries.

An exemplary implementation of the passive concentrator is shown in figure 49. The TCU is the only functional element in the data path between the LRX connections and the LTX connections in a passive concentrator. In order to ensure interoperability between implementations, the passive concentrator shall meet the specifications in this subclause.

All resistive loads (terminations or sources impedances) discussed in this subclause used to measure concentrator characteristics have a ± 1% tolerance unless otherwise noted.

8.4.1 Passive concentrator return loss

Passive concentrators with or without ring segment boundaries shall meet the return loss requirements specified in 8.4.1.1. Passive concentrators with passive trunks ports shall meet the reflection coefficient requirements specified in 8.4.1.2. Passive concentrators with ring segment boundaries shall meet the return loss requirements specified in 8.4.1.3.

8.4.1.1 Passive concentrator with or without ring segment boundaries

The return loss is measured in the frequency domain at the TX and RX connector pins of each concentrator port excluding paths that include ring segment boundaries. Measurements conducted on the TX connector pins of a port will be terminated at the port's RX connector pins with 100 Ω (UTP) or 150 Ω (STP). Similarly, measurements conducted on a port's RX connector pins will be terminated at the TX connector pins. Two types of measurements shall be conducted 1) with the port in the inserted state with any other port bypassed or inserted and terminated with 100 Ω (UTP) or 150 Ω (STP), and 2) with the port in the bypassed state. The passive concentrator excluding paths that include ring segment boundaries shall provide a minimum return loss, relative to a reference impedance of 150 Ω (STP) or 100 Ω (UTP), as shown in table 35.

Table 35—Minimum return loss for passive concentrators including ring segment boundaries

Frequency range	Minimum return loss
1 MHz to 17 MHz	14 dB
17 MHz to 25 MHz	11 dB

8.4.1.2 Passive concentrators without ring segment boundaries

To ensure interoperability in a network of cascaded passive concentrators that do not contain ring segment boundaries, the reflection coefficient is measured on passive concentrators without ring segment boundaries. The reflection coefficient is measured in the time domain with a 200 kHz trapezoid with a rise and fall times (10% to 90%) of 8 ns issued by a differential generator of impedance 150 Ω (STP) or 100 Ω (UTP). The port at the opposite end from the generator shall be terminated with a load of matched impedance, 150 Ω (STP) or 100 Ω (UTP). Two types of measurements shall be conducted trunk port-to-trunk port and lobe port-to-trunk port. The trunk port-to-trunk port is measured with all lobe ports bypassed and is measured from TXMAIN with RXMAIN terminated and from TXBACKUP with RXBACKUP terminated. The lobe port-to-trunk port is measured at each lobe port with the remaining lobe ports bypassed and is measured from LRX with TXMAIN terminated and LTX with RXMAIN terminated. The passive concentrator without ring segment boundaries shall provide a maximum reflection coefficient shown in table 36.

Table 36—Maximum reflection coefficients for a passive concentrator without ring segment boundaries

Absolute value maximum reflection coefficient (rho)	Description
0.126	Trunk port-to-trunk port
0.178	Lobe port-to-trunk port

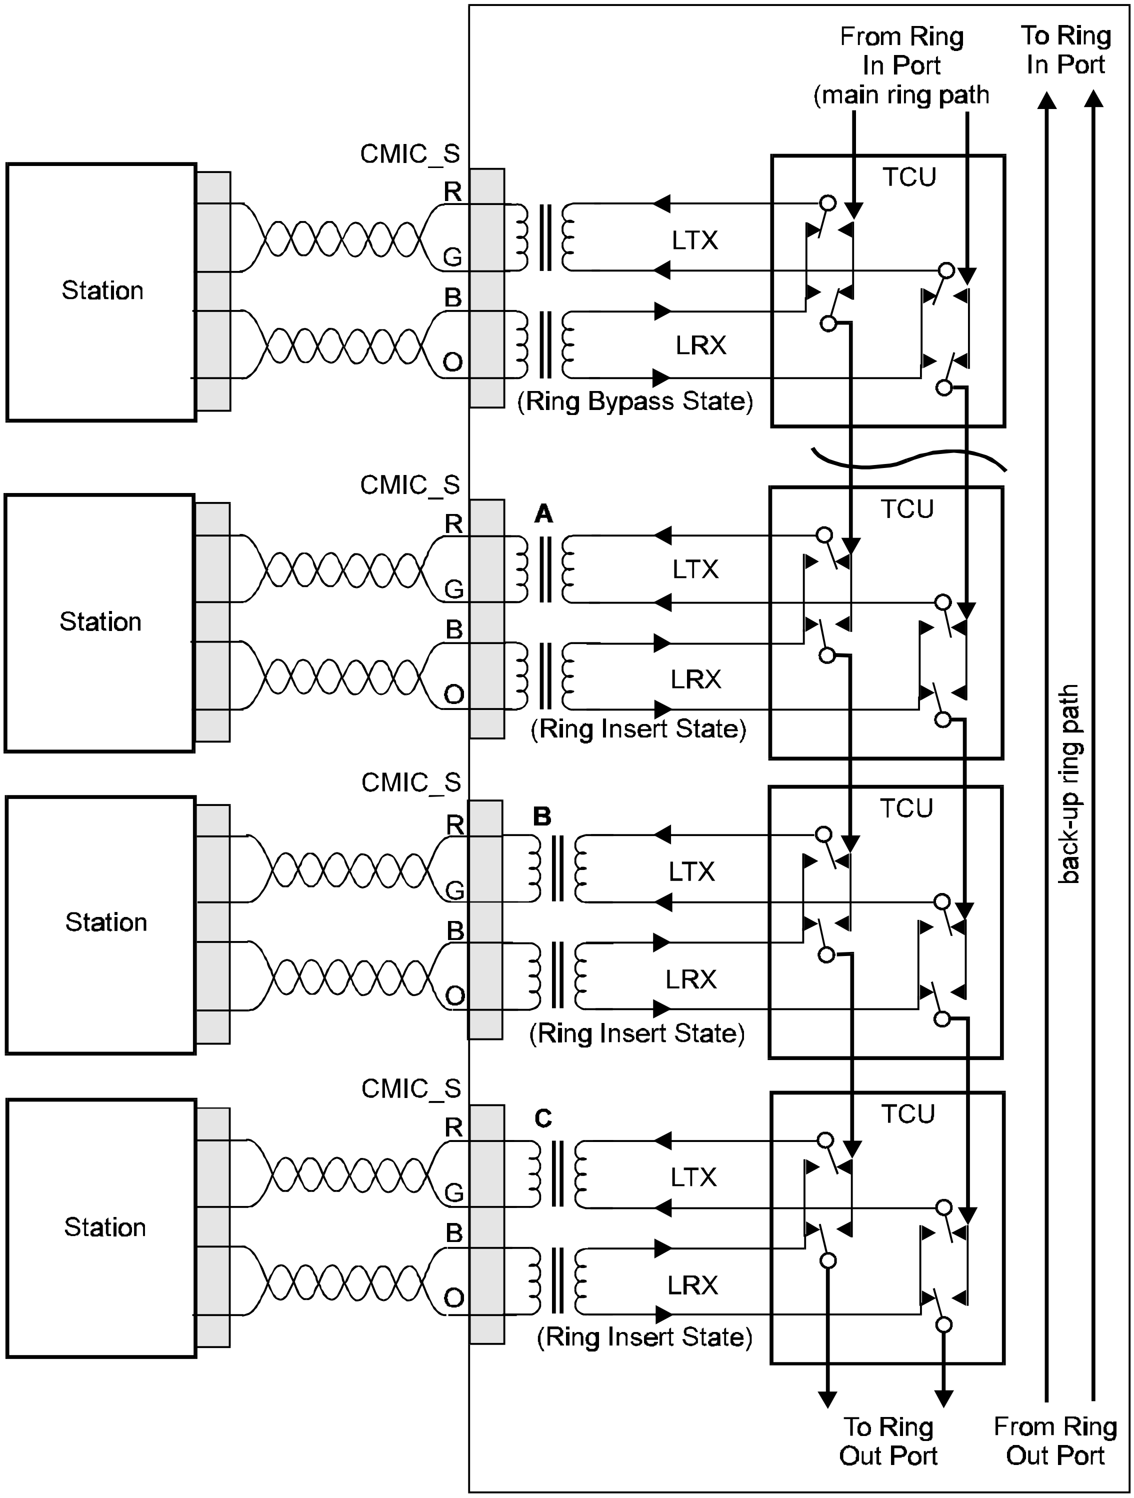

Figure 49—Example of an STP passive concentrator’s signal path

8.4.1.3 Passive concentrator with ring segment boundaries

For paths that include ring segment boundaries, at least one port must be measured wherein the return loss of the ring segment boundaries transmitter shall meet the requirements of 7.2.2.4.[27] Similarly, at least one port must be measured wherein the return loss of the ring segment boundaries receiver shall meet the requirements of 7.2.3.2.[27]

8.4.2 Passive concentrator attenuation

The passive concentrator attenuation consists of four types: lobe square-root-frequency attenuation (CLSQA), lobe flat attenuation (CLATT), trunk square-root-frequency attenuation (CTSQA), and trunk flat attenuation (CTATT). Maximum CLSQA and CLATT must be included in the calculation of the channel losses. CLSQA and CLATT are based on the worst-case specified loss for the transmission path from a CMIC onto the main ring path, to the RO port's self-shorting mechanism, on to the backup path to the RI port's self-shorting mechanism, back on to the main ring path, and back to the same CMIC if there are no ring segment boundaries. CLSQA and CLATT are based on the worst-case measured loss from any TX to any RX connector pins of each concentrator port excluding any retimed path. In addition, if multiple concentrators are cascaded and the trunk ports are passive trunk ports, CTSQA and CTATT must also be included in the channel losses. The implementor shall include in the published specification for the device the maximum CLSQA both at 4 MHz and 16 MHz and CLATT. If the trunk ports are passive trunk ports, the implementor shall also include in the published specification for the device the maximum CTSQA both at 4 MHz and 16 MHz and CTATT. The measurements for CLSQA and CLATT should be taken on the maximum path between the LTX and LRX connector pins on any inserted lobe port. CTSQA and CTATT should be measured between RXMAIN and TXBACKUP with the RO port externally looped (TXMAIN to RXBACKUP). For passive concentrators with small attenuation values it may be necessary to cascade multiple concentrators to obtain reliable results. All measurements shall be based on power loss. In all cases the maximum flat loss shall be

$\text{CLATT} + n{*}\text{CTATT} \leq 2\ \text{dB}$,[28] where n is the maximum number of concentrators in a ring segment.

C*x*ATT and C*x*SQA , where *x* is L or T, shall be determined by the following procedure:

a) The insertion loss in dB, IL_i, is measured at the following frequencies:

   $\text{F}_i = i * 1\text{E}6$ Hz, where i is an integer that ranges from 1 to 30

b) A linear least squares fit of the equation is

$$\text{IL} = \text{CxSQA*SF} + \text{CxATT} \quad \text{where SF} \; = \text{SQRT(F/16E6) for the 16 MHz case}$$
$$= \text{SQRT(F/4E6) for the 4 MHz case}$$

   to the 30 pairs of IL_i and F_i data points is made, yielding

   CxSQA and CxATT.

8.4.3 Passive concentrator low-frequency attenuation response

The passive concentrator low-frequency attenuation cutoff (CxF1), or low-frequency 3 dB point, consists of two specifications: trunk low-frequency cutoff (CTF1) and lobe low-frequency cutoff (CLF1). The

---

[27] It may be necessary to make this measurement on devices modified from standard production hardware because of software, controls, etc. It shall be acceptable to make these tests on such modified devices provided the manufacturer can show that the modifications made do not affect the electrical characteristics of the transmission path being measured.

[28] The channel requirements in 7.2.4 specify a maximum flat loss limit for the channel of 2 dB. The CATT may limit the maximum number of cascaded concentrators within a ring segment, and hence the number of stations that may be connected to that segment. Implementors should ensure that the CATT and CSQA are considered when determining their maximum cascaded configuration in a ring segment. Additional information relative to the role of CATT and CSQA are given in 7.2.4 and in annex B.

measurement shall be taken on the passive trunk and lobe paths as described above in 8.4.2. For passive trunk ports, CTF1 shall meet the following specification:

CTF1 < 1 kHz

For the lobe ports, CLF1 shall meet the following specification:

CLF1 < 25 kHz

8.4.4 Passive concentrator crosstalk

The passive concentrator crosstalk is measured as follows:

a) On each port (individually) from its LTX to its LRX with all ports in the insert state and their LRX and LTX terminated at the port in 100 Ω (UTP) or 150 Ω (STP)
b) From RXMAIN to TXBACKUP with RI and RO inserted and terminated and all lobe ports bypassed except the last lobe port directly upstream of RO is inserted and terminated.
c) The same as b) except from RXBACKUP to TXMAIN with the first lobe port directly downstream of the RI inserted and terminated.

The passive concentrator shall provide a minimum crosstalk loss as shown in table 37.

Table 37—Passive concentrator minimum crosstalk requirements

Frequency range	Minimum crosstalk loss
0.5 MHz to 4 MHz	43 dB
4 MHz to 24 MHz	40 dB

8.5 Active retiming concentrator

This subclause defines the electrical characteristics of the active retiming concentrator (ARC) lobe and trunk ports. The ARC may be used to support lobe cabling lengths beyond the limits provided by a passive concentrator as imposed by signal degradation due to channel impairments. The ARC retimes signals passing through its ports and restores signal amplitude and shape. The data path in the ARC will include three functional blocks: the trunk ports, TCU, and lobe ports. The ports implement a transmitter, including all circuitry in the transmission path between the data retiming (latching) mechanism and the CMIC (RIMIC or ROMIC) transmit connections; and the receiver, including all circuitry between the CMIC (RIMIC or ROMIC) receive connections and the recovered clock used to latch the data. An example of the ARC lobe port transmitter and receiver is shown in figure 50.

8.5.1 ARC port

In order to ensure interoperability between implementations and to provide a physical containment of the channels, each port of the ARC, including the lobe, RI, and RO ports, shall provide a transmitter meeting the specifications of 7.2.2, a receiver meeting the specifications of 7.2.3, and an overall PHY performance meeting the specifications of 7.1. The ARC lobe port transmitter and receiver shall be specified with the worst-case phantom voltage mismatch allowed by 7.2.1.1.

8.5.2 ARC concentrator specific components

The ARC port clock recovery circuit shall be phase synchronized within 1.5 ms after receipt of a valid signal from the upstream station. In the determination of the minimum supported PMC count (see table 32) the ARC port is counted as one PMC. Burst error correction, if implemented, shall only occur on burst6 errors correcting to burst5. The clock recovery circuit may transmit fill, preferably from a crystal clock, while frequency error is occurring. The ARC retiming elements, which are not part of the active monitor, shall not delete bits in the interframe gap if the deletion reduces the size of the interframe gap below the minimum size specified in 3.2.10. To ensure active retimed concentrators do not introduce excessive latency, the average PMC latency should not exceed 100 symbols[29] per port to prevent expiration of timer TRR (3.4.2.13). It is recommended that the port latency be minimized in order to maximize performance/access.

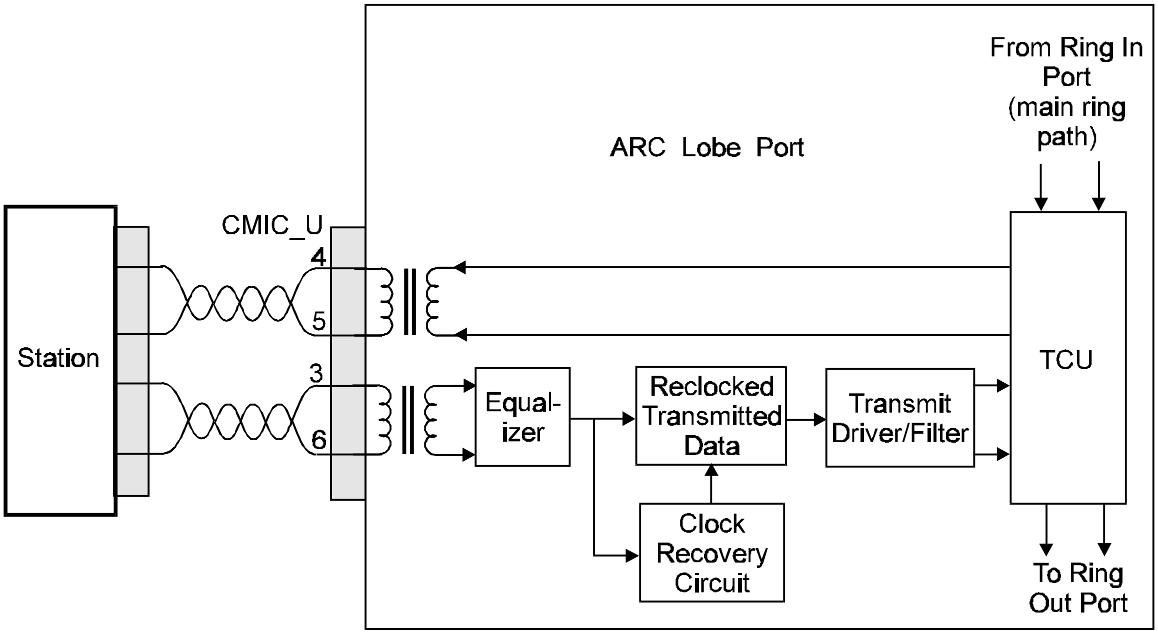

Figure 50—Example of the ARC lobe port transmitter and receiver

---

[29] For 4 Mbit/s operation this number is based upon the assumption that a ring is comprised of a maximum number of ARC ports and stations of 125 each and 100 m lobe length cables (with a propagation speed of 0.7 times the speed of light) and with no trunk cabling.

Annex A

(normative)

Protocol Implementation Conformance Statement (PICS) proforma

A.1 Introduction

The supplier of a protocol implementation that is claimed to conform to this standard shall complete the following Protocol Implementation Conformance Statement (PICS) proforma.

A completed PICS proforma is the PICS for the implementation in question. The PICS is a statement of which capabilities and options of the protocol have been implemented. The PICS can have a number of uses, including use by the following:

a) The protocol implementor, as a check-list to reduce the risk of failure to conform to the standard through oversight;

b) The supplier and acquirer, or potential acquirer, of the implementation, as a detailed indication of the capabilities of the implementation, stated relative to the common basis for understanding provided by the standard PICS proforma;

c) The user, or potential user, of the implementation, as a basis for initially checking the possibility of interworking with another implementation (note that, while interworking can never be guaranteed, failure to interwork can often be predicted from incompatible PICSs);

d) The protocol tester, as the basis for selecting appropriate tests against which to assess the claim for conformance of the implementation.

A.2 Abbreviations and special symbols

A.2.1 Status symbols

The following symbols are used in the PICS proforma:

M	mandatory field/function
O	optional field/function
O.<n>	optional field/function indicating mutually exclusive or selectable options among a set
X	prohibited field/function
<item>:	simple-predicate condition, dependent on the support marked for <item>

A.2.2 Abbreviations

N/A	Not applicable

A.3 Instructions for completing the PICS proforma

A.3.1 General structure for the PICS proforma

The first part of the PICS proforma, implementation identification and protocol summary, is to be completed as indicated with the information necessary to identify fully both the supplier and the implementation.

The main part of the PICS proforma is a fixed-format questionnaire divided into subclauses, each containing a group of items. Answers to the questionnaire items are to be provided in the right-most column, either by simply marking an answer to indicate a restricted choice (usually Yes, No, or Not Applicable), or by entering a value, or a set or a range of values. (Note that there are some items where two or more choices from a set of possible answers can apply; all relevant choices are to be marked.)

Each item is identified by an item reference in the first column; the second column contains the question to be answered; the third column contains the reference or references to the material that specifies the item in the main body of the standard. The remainder of the columns record the status of the item—whether the support is mandatory, optional, or conditional—and provide spaces for the answers; see also A.3.4.

The supplier may also provide, or be required to provide, further information, categorized as either “additional information” or “exception information.” When present, each kind of further information is to be provided in a further subclause of items labeled A<i> or E<i>, respectively, for cross-referencing purposes, where <i> is the unambiguous identification for the item (e.g., simply a numerical); there are no other restrictions on its format or presentation.

A completed PICS proforma, including any additional information or exception information, is the PICS for the implementation in question.

Note that where an implementation is capable of being configured in more than one way, according to the items listed in A.5, a single PICS may be able to describe all such configurations. However, the supplier has the choice of providing more than one PICS, each covering some subset of the implementation’s configuration capabilities, if that would make presentation of information easier and cleaner.

A.3.2 Additional information

Items of additional information allow a supplier to provide further information intended to assist the interpretation of the PICS. It is not intended or expected that a large quantity will be supplied, and the PICS can be considered complete without any such information. Examples might be an outline of the ways in which a (single) implementation can be set up to operate in a variety of environments and configurations; or a brief rationale, based perhaps upon specific application needs, for the exclusion of features which, although optional, are nonetheless commonly present in implementations of the token ring protocol.

References to items of additional information may be entered next to any answer in the questionnaire, and may be included in items of exception information.

A.3.3 Exception information

It may occasionally happen that a supplier will wish to answer an item with mandatory status or prohibited status (after any conditions have been applied) in a way that conflicts with the indicated requirement. No preprinted answer will be found in the Support column for this; instead, the supplier is required to write into the Support column an E<i> reference to an item of exception information, and to provide the appropriate rationale in the exception item itself.

An implementation for which an exception item is required in this way does not conform to this standard.

Note that a possible reason for the situation described above is that a defect in the standard has been reported, a correction for which is expected to change the requirement not met by the implementation.

A.3.4 Conditional status

A.3.4.1 Conditional items

The PICS proforma contains a number of conditional items. These are items for which the status—mandatory, optional, or prohibited—that applies is dependent upon whether or not certain other items are supported, or upon the value supported for other items.

In many cases, whether or not the item applies at all is conditional in this way, as well as the status when the item does not apply.

A conditional symbol is of the form "<pred>: <s>" where "<pred>" is a predicate as described in A.3.4.2, and "<s>" is one of the status symbols M, O, O.<n>, or X.

A conditional symbol of the form "<pred> : :" may be indicated above a particular table. That table shall be completed if and only if the condition evaluates to true.

A.3.4.2 Predicates

A predicate is one of the following:

a) An item-reference for an item in the PICS proforma. The value of the predicate is true if the item is marked as supported, and is false otherwise; or
b) A predicate name for a predicate defined elsewhere in the PICS proforma (usually in the major capabilities section or at the end of the section containing the conditional item), see (a), (b), (c); or
c) The logical negation symbol "¬" prefixed to an item-reference or predicate name. The value of the predicate is true if the value of the predicate formed by omitting the "¬" symbol is false, and vice versa.

The definition for a predicate name is one of the following:

a) An item-reference, evaluated as at (1) above; or
b) A relation containing a comparison operator (i.e., = , < , etc.) with at least one of its operands being an item-reference for an item taking numerical values as its answer. The predicate is true if the relation holds when each item-reference is replaced by the value entered in the Support column as answer to the item referred to; or
c) A Boolean expression constructed by combining simple predicates, as at (a) and (b), using the Boolean operators AND, OR, and NOT, and parentheses, in the usual way. The value of such a predicate is true if the Boolean expression evaluates to true when the simple predicates are interpreted as described previously.

Each item-reference that is used in a predicate or predicate definition is indicated by an asterisk in the Item column. If such item reference is not supported (false), then the support of the item itself will be indicated as N/A (not applicable); otherwise, the support of the item will be indicated as YES.

A.4 Identification

A.4.1 Implementation identification

Supplier	
Contact point for queries about the PICS	
Implementation name(s) and version(s)	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; system name(s)	
NOTES 1—Only the first three items are required for all implementations; other information may be completed as appropriate in meeting the requirements for full identification. 2—The terms *name* and *version* should be interpreted appropriately to correspond with a supplier's terminology (e.g., type, series, model).	

A.4.2 Protocol summary

Protocol version	
Amendments implemented	
Corrigenda implemented	
Have any exception items been required? No [ ] Yes [ ] (See A.3.3; the answer "Yes" means that the implementation does not conform to the standard.)	
Date of statement	

A.5 Major capabilities

Item	Feature	Reference	Status	Support
*DS	Data station	2.4	O.1	Yes [] No []
*ACON	Active retiming concentrator	8.5	O.1	Yes [] No []
*PCON	Passive concentrator	8.4	O.1	Yes [] No []
* DR4	4 Mbit/s data rate	5.3	O.2	Yes [] No []
* DR16	16 Mbit/s data rate	5.3	O.2	Yes [] No []
* STP	Shielded twisted pair cable attachment	7.2	O.3	Yes [] No []
* UTP	Unshielded twisted pair cable attachment	7.2	O.3	Yes [] No []
NOTES— O.1: Support for one and only one of the options is required. O.2 and O.3: Support for at least one of the options is required.				

A.6 PICS proforma for the MAC sublayer

A.6.1 Transmission formats—DS::M

Does the data station support the following frame formats?

Item	Feature	Reference	Status	Support
FF1	Token transmit	3.1.1	M	Yes []
FF1A	Token receive	3.1.1	M	Yes []
FF2	MAC frame transmit	3.1.2	M	Yes []
FF2A	MAC frame receive	3.1.2	M	Yes []
FF3	LLC frame transmit	3.1.2	M	Yes []
FF3A	LLC frame receive	3.1.2	M	Yes []
FF4	Abort sequence transmit	3.1.3	M	Yes []
FF4A	Abort sequence receive	3.1.3	M	Yes []
FF5	Fill transmit	3.1.4	M	Yes []
FF5A	Fill receive	3.1.4	M	Yes []

A.6.2 Frame parameters—DS::M

Does the data station support the following frame parameters?

Item	Feature	Reference	Status	Support
FP1	Starting delimiter transmit	3.2.1	M	Yes []
FP1A	Starting delimiter receive	3.2.1	M	Yes []
FP2	Access control transmit	3.2.2	M	Yes []
FP2A	Access control receive	3.2.2	M	Yes []
FP3	Frame control transmit	3.2.3	M	Yes []
FP3A	Frame control receive	3.2.3	M	Yes []
FP4	Destination address transmit	3.2.4.1	M	Yes []
FP4A	Destination address receive	3.2.4.1	M	Yes []
FP5	Source address transmit	3.2.4.2	M	Yes []
FP5A	Source address receive	3.2.4.2	M	Yes []
FP6	Routing information indicator transmit	3.2.4.2	M	Yes []
FP6A	Routing information indicator receive	3.2.4.2	M	Yes []
*FP7	Routing information field transmit	3.2.5	O	Yes [] No []
FP7A	Routing information field receive	3.2.5	M	Yes []
FP8	RI field length bits transmit	3.2.5	FP7: M	N/A [] Yes []
FP8A	RI field length bits receive	3.2.5	M	Yes []
FP9	MAC frame, info field transmit	3.2.6.2	M	Yes []
FP9A	MAC frame, info field receive	3.2.6.2	M	Yes []
FP10	LLC frame, info field transmit	3.2.6.3	M	Yes []
FP10A	LLC frame, info field receive Minimum 133 octets	3.2.6.3	M	Yes [] ______ octets
FP11	Frame check sequence transmit	3.2.7	M	Yes []
FP11A	Frame check sequence receive	3.2.7	M	Yes []
FP12	Ending delimiter transmit	3.2.8	M	Yes []
FP12A	Ending delimiter receive	3.2.8	M	Yes []
FP13	Frame status transmit	3.2.9	M	Yes []
FP13A	Frame status receive	3.2.9	M	Yes []

A.6.3 MAC frames—DS::M

Does the data station implement the following MAC Frame Station Operation Table transitions?

Item	**Feature**	**Reference**	**Status**	**Support**
DSM	Transitions relating to MAC Frames	3.3, 4	M	Yes []

A.6.4 MAC frame subvectors—DS::M

See A.6.3.

A.6.5 Timers—DS::M

Does the data station support the following timers as specified in the referenced subclause of the standard? Note that the time intervals are applicable only when observable from external behavior. No particular realization is implied.

Item	Feature	Reference	Status	Support
TAM	Active monitor	3.4.2.1	M	Yes []
TBR	Beacon repeat	3.4.2.2	M	Yes []
TBT	Beacon transmit	3.4.2.3	M	Yes []
TCT	Claim token	3.4.2.4	M	Yes []
TER	Error report	3.4.2.5	M	Yes []
TID	Insert delay	3.4.2.6	M	Yes []
TJR	Join ring	3.4.2.7	M	Yes []
TNT	No token	3.4.2.8	M	Yes []
TQP	Queue PDU	3.4.2.9	M	Yes []
TRH	Remove hold	3.4.2.10	M	Yes[]
TRI	Request initialization	3.4.2.12	M	Yes []
TRP	Ring purge	3.4.2.14	M	Yes []
TRR	Return to repeat	3.4.2.13	M	Yes []
TRW	Remove wait	3.4.2.11	M	Yes []
TSL	Signal loss	3.4.2.15	M	Yes []
TSM	Standby monitor	3.4.2.16	M	Yes []
TVX	Valid transmission	3.4.2.17	M	Yes []
TWF	Wire fault	3.4.2.19	M	Yes []
TWFD	Wire fault delay	3.4.2.18	M	Yes []

A.6.6 Station policy flags—DS::M

Note that the following requirements are applicable only when observable from external behavior.

Item	Feature	Reference	Status	Support
FBHO_0	Flag, beacon handling option—flag=0	3.5.1	O.5	Yes [] No []
FBHO_1	Flag, beacon handling option—flag=1	3.5.1	O.5	Yes [] No []
FCCO_0	Flag, claim contender option—flag=0	3.5.2	O.6	Yes [] No []
FCCO_1	Flag, claim contender option—flag=1	3.5.2	O.6	Yes [] No []
FETO_0	Flag, ETR option—flag=0	3.5.3	DR4: M DR16:O.7	N/A [] Yes [] N/A [] Yes [] No []
FETO_1	Flag, ETR option—flag=1	3.5.3	DR16:O.7	N/A [] Yes [] No []
FECO_0	Flag, error counting option—flag=0	3.5.4	O.8	Yes [] No []
FECO_1	Flag, error counting option—flag=1	3.5.4	O.8	Yes [] No []
FMRO_0	Flag, media rate—flag=0	3.5.5	DR4:M	N/A [] Yes []
FMRO_1	Flag, media rate—flag=1	3.5.5	DR16:M	N/A [] Yes []
FMFTO_0	Flag, multiple frame transmission—flag=0	3.5.6	O.9	Yes [] No []
FMFTO_1	Flag, multiple frame transmission—flag=1	3.5.6	O.9	Yes [] No []
FRRO_0	Flag, reject remove option—flag=0	3.5.7	O.10	Yes [] No []
FRRO_1	Flag, reject remove option—flag=1	3.5.7	O.10	Yes [] No []
FTEO_0	Flag, token error detect option—flag=0	3.5.8	O.11	Yes [] No []
FTEO_1	Flag, token error detect option—flag=1	3.5.8	O.11	Yes [] No []
FTHO_0	Flag, token handling option—flag=0	3.5.9	O.12	Yes [] No []
FTHO_1	Flag, token handling option—flag=1	3.5.9	O.12	Yes [] No []
FTUBO_0	Flag, transmit underrun behavior option —flag=0	3.5.9a	O.22	Yes [] No []
FTUBO_1	Flag, transmit underrun behavior option —flag=1	3.5.9a	O.22	Yes [] No []
FWFDO_0	Flag, wire fault detection option	3.5.11	O.23	Yes [] No []
FWFDO_1	Flag, wire fault detection option	3.5.11	O.23	Yes [] No []
FGTO_0	Flag, good token option—flag=0	3.5.10	O.13	Yes [] No []
FGTO_1	Flag, good token option—flag=1	3.5.10	O.13	Yes [] No []
NOTE—Support of at least one of the options shown above is required.				

A.6.7 Counters—DS::M

These counters are considered to be visible externally only through the transmission of the Report Error MAC frame (see 3.3.5, table 4).

Item	Feature	Reference	Status	Support
CABE	Abort error	3.6.1	M	Yes []
CACE	AC error	3.6.2	M	Yes []
CBE	Burst error	3.6.3	M	Yes []
CFCE	Frame-copied error	3.6.4	M	Yes []
CFE	Frequency error	3.6.5	M.21	Yes []
CIE	Internal error	3.6.6	M.21	Yes []
CLE	Line error	3.6.7	M	Yes []
CLFE	Lost frame error	3.6.8	M	Yes []
CRCE	Receive congestion error	3.6.9	M	Yes []
CTE	Token error	3.6.10	M	Yes []
NOTE—M.21: These counters are mandatory, but the indications that cause CFE and CIE to be incremented are optional.				

A.7 PICS proforma for the physical layer

A.7.1 Symbol timing—DS::M, ACON::M

Do the data station and the active concentrator support the following symbol timings?

Item	Feature	Reference	Status	Support
ST1	4 Mbit/s data signaling rate	5.2	DR4: M	N/A[] Yes []
ST2	16 Mbit/s data signaling rate	5.2	DR16: M	N/A[] Yes []
ST3	Acquire phase lock within 1.5 ms	5.7.1	M	Yes []
ST4	Frequency error	5.7.2	O	Yes[] No []
ST5	Signal loss indication	5.7.1	O	Yes[] No []

A.7.2 Symbol encoding and decoding—DS::M

Do the data station and the active concentrator support the following symbol encoding and decoding?

Item	Feature	Reference	Status	Support
SY1	Symbol encoding	5.3	M	Yes []
SY2	Symbol decoding	5.6	M	Yes []

SY3	Burst error/idles transmit	5.4.2	M	Yes []

A.7.3 Station latency—DS::M

Does the data station provide for the following latencies?

Item	Feature	Reference	Status	Support
LB1	A fixed latency buffer of 24 symbols	5.8.2	M	Yes []
LB2	4 Mbit/s latency variation	5.8.3	DR4: M	N/A [] Yes []
LB3	16 Mbit/s latency variation	5.8.3	DR16: M	N/A [] Yes []

A.7.4 Access control

DS::M Does the data station support the following access control?

Item	Feature	Reference	Status	Support
RA1	Perform ring access control	5.9	M	Yes []
RA2	Phantom circuit source/return	7.2.1.1	M	Yes []
RA3	Ring insertion current/voltage	7.2.1.1	M	Yes []
RA4	Ring bypass current/voltage	7.2.1.1	M	Yes []
RA5	Lobe fault indication	7.2.1.2	M	Yes []

ACON::M, PCON::M Does the concentrator support the following access control?

Item	Feature	Reference	Status	Support
CRAC1	Ring insertion max time	8.3.2	M	Yes []
CRAC2	Ring bypass max time	8.3.2	M	Yes []
CRAC3	Phantom dc load	8.3.3	M	Yes []
CRAC4	Max ring open time	8.3.2	M	Yes []
CRAC5	Phantom path leakage resistance	8.3	M	Yes []
CRAC6	Ring access control insert & bypass	8.3.1	M	Yes []

A.8 Attachment specifications

A.8.1 Accumulated correlated jitter—DS::M, ACON::M

Do the data station and the active concentrator support the following specifications?

Item	Feature	Reference	Status	Support
AJ1a	Filtered accumulated phase jitter	7.1.1	DR4: M	N/A [] Yes []
AJ1b	Filtered accumulated phase jitter	7.1.1	DR16: M	N/A [] Yes []
AJ2a	Delta phase accumulated phase jitter	7.1.1	DR4: M	N/A [] Yes []
AJ2b	Delta phase accumulated phase jitter	7.1.1	DR16: M	N/A [] Yes []
AJ3a	Accumulated uncorrelated jitter	7.1.2	DR4: M	N/A [] Yes []
AJ3b	Accumulated uncorrelated jitter	7.1.2	DR16: M	N/A [] Yes []
AJ4	PHY net delay	7.1.3	M	Yes []

A.8.2 Transmitter specification—DS::M, ACON::M

Do the data station and the active concentrator support the following specifications?

Item	Feature	Reference	Status	Support
TR1a	Transmit duty cycle distortion	7.2.2.1	DR4: M	N/A [] Yes []
TR1b	Transmit duty cycle distortion	7.2.2.1	DR16: M	N/A [] Yes []
TR2a	Transmit Tdiff01	7.2.2.2.1	DR4: M	N/A [] Yes []
TR2b	Transmit Tdiff01	7.2.2.2.1	DR16: M	N/A [] Yes []
TR3	Transmit Tdiffmax	7.2.2.2.1	M	Yes []
TR4a	Transmit waveform(zero/one/SDEL)	7.2.2.2.2	DR4: M	N/A [] Yes []
TR4b	Transmit waveform(zero/one/SDEL)	7.2.2.2.2	DR16: M	N/A []Yes []
TR5a	Transmit output voltage	7.2.2.3	STP: M	N/A [] Yes []
TR5b	Transmit output voltage	7.2.2.3	UTP: M	N/A [] Yes []
TR6a	Transmit return loss	7.2.2.4	STP: M	N/A [] Yes []
TR6b	Transmit return loss	7.2.2.4	UTP: M	N/A [] Yes []

A.8.3 Receiver specification—DS::M, ACON::M

Do the data station and the active concentrator support the following receiver specifications?

Item	Feature	Reference	Status	Support
RC1a	Rcvr jitter tolerance (no noise)	7.2.3.1	DR4: M	N/A [] Yes []
TR1b	Rcvr jitter tolerance (no noise)	7.2.3.1	DR16: M	N/A [] Yes []
RC2a	Rcvr jitter tolerance (with noise)	7.2.3.1	DR4: M	N/A [] Yes []
TR2b	Rcvr jitter tolerance (with noise)	7.2.3.1	DR16: M	N/A [] Yes []
TR3a	Rcvr return loss	7.2.3.2	DR4: M	N/A [] Yes []
TR3b	Rcvr return loss	7.2.3.2	DR16: M	N/A [] Yes []

A.8.4 Connector specification

DS::M Does the data station support the following connector specification?

Item	Feature	Reference	Status	Support
MI1	STP media interface connector	7.2.5.1	STP: M	N/A [] Yes []
MI2	UTP media interface connector	7.2.5.2	UTP: M	N/A [] Yes []

ACON::M, PCON::M Does the concentrator support the following specifications?

Item	Feature	Reference	Status	Support
CC1a	STP media lobe connector	8.1.1	STP: M	N/A [] Yes []
CC1b	UTP media lobe connector	8.1.1	UTP: M	N/A [] Yes []
CC2a	Trunk connected STP MIC	8.2.1	O.14	Yes [] No []
CC2b	Trunk connected UTP MIC	8.2.1	O.14	Yes [] No []
CC2c	No trunk connection	8.2.1	O.14	Yes [] No []
CC3	Main ring signal path	8.2.1	M	Yes []
CC4a	Trunk connected backup signal path	8.2.1	O.15	Yes [] No []
CC4b	No trunk connection	8.2.1	O.15	Yes [] No []
CC5	Lobe port indicators	8.2.1	O	Yes [] No []
NOTE—Support of at least one of each option shown above is required.				

A.9 Concentrator specific requirements

ACON::M Does the active concentrator support the following specifications?

Item	Feature	Reference	Status	Support
AC1	Burst error correction	8.5.2	O	Yes [] No []
AC2	Deleting interframe bits	8.5.2	O	Yes [] No []
AC3a	Ring segment trunk port	8.2	O.16	Yes [] No []
AC3b	No trunk port	8.2	O.16	Yes [] No []
NOTE—Support of at least one of each option shown above is required.				

PCON::M Does the passive concentrator support the following specifications?

Item	Feature	Reference	Status	Support
CPA1	Lobe return loss	8.4.1.1	M	Yes []
CPA2a	Trunk reflection coefficient	8.4.1.2	O.17	Yes [] No []
CPA2b	Ring segment boundary return loss	8.4.1.3	O.17	Yes [] No []
CPA2c	No passive trunk	8.4.1	O.17	Yes [] No []
CPA3	Maximum flat loss	8.4.2	M	Yes []
CPA4	Published lobe attenuation values	8.4.2	M	Yes []
CPA5	Crosstalk loss	8.4.4	M	Yes []
CPA6	Lobe low-frequency response	8.4.3	M	Yes []
CPA7a	Passive trunk low-frequency response	8.4.3	O.18	Yes [] No []
CPA7b	No passive trunk	8.4.3	O.18	Yes [] No []
CPA8a	Published passive trunk port	8.4	O.19	Yes [] No []
CPA8b	Published ring segment boundary trunk port	8.4	O.19	Yes [] No []
CPA8c	No trunk port	8.4	O.19	Yes [] No []
CPA9a	Published passive trunk attenuation values	8.4.2	O.20	Yes [] No []
CPA9b	No passive trunk	8.4.2	O.20	Yes [] No []
NOTE—Support of at least one option shown above is required.				

Annex B

(informative)

Channel design considerations

This annex provides channel design examples and cabling guidelines to aid users in implementing the passive and active channels specified in 7.2.4. Implementation of the channels specified in 7.2.4 is intended to result in a jitter budget, as specified in annex C, with sufficient margin to ensure a bit error rate (BER) not exceeding 10^{-9} for a station and 10^{-8} for the ring when connected to equipment that conforms to this standard.

B.1 Channel transmission parameters

When designing a passive or active channel, the primary channel transmission parameters of interest are

a) Insertion loss,
b) NEXT Loss-to-Insertion Loss Ratio (NIR), and
c) Total flat loss for passive channels. The channel impedance is satisfied by selecting the appropriate media.

B.1.1 Insertion loss

As specified in 7.2.4.2.1, the insertion loss of the passive and active channels is the loss between the MIC at a transmitter and the MIC of the next downstream receiver. This loss includes all the components in the channel, which are shown in table B.1 for passive and active channels. As shown in figure 37, the passive channel path includes two lobes since both channel terminations are at the station MICs. Hence, there are at least twice the number of components in a passive channel as there are in an active channel. The concentrator losses are not included in the active channel since the active channel is defined between a station MIC and a concentrator MIC and, between the RI and RO MIC ports of concentrators. Flexible cords used in the channel that do not match the attenuation of the building cable as specified in ISO/IEC 11801 : 1995, should be taken into account when calculating the length of channel components.

Table B.1—Channel components

Component	Passive channel	Active channel
Cords and cables	Station cords Building cables Patch cords Equipment cords	Station cord Building cable Patch cord Equipment cord
Connectors	Wall outlets Equipment and cable termination blocks	Wall outlet Equipment and cable termination blocks
Concentrator	Concentrator(s) Inter-concentrator cord(s)	

B.1.2 Channel noise

As discussed in 7.2.4.3, the noise that appears at the output of a channel consists of two types: crosstalk noise due to signals within the lobe cabling from an adjacent channel and noise induced into the channel from sources external to the channel. Noise that originates external to the channel may be due to a multitude of sources: high intensity fields, noisy motors, etc., and is minimized by the careful planning and installation of cables and placement of concentrators. It is strongly recommended that a token ring channel not share a cable sheath with other services that induce noise into the token ring channel, e.g., analog telephones, RS-232 signals.

B.1.3 Channel input/output impedance

Ring jitter is minimized by using cables, cords, and components that have the same nominal characteristic impedance. This minimizes the impedance discontinuities between components, which in turn minimizes the strength of the signals reflected at the impedance discontinuities. Mixing channel components that have different impedances without impedance matching will result in increased jitter that may degrade network performance and is not recommended. Cables/cords of 100 Ω and 150 Ω should not be mixed in a passive channel. Also, to minimize channel crosstalk noise due to connectors, the NEXT loss of channel connectors should complement the cable type. The electrical specifications for UTP and cabling connectors are specified in ISO/IEC 11801 : 1995.

B.2 Passive channel design example

This clause is intended for manufacturers of passive concentrators as well as technical personnel that design token ring networks. Two passive channel design examples are provided: a ring consisting of a single concentrator and a ring that consists of three concentrators. The three concentrator ring example may be extended to a ring of multiple concentrators by modifying the appropriate parameters. On each ring only a single station is inserted, since it is the single station ring architecture that limits the supportable lobe length (loss). These examples are intended to provide guidance in designing passive channels as well as the application and limitation of the specifications of clause 7 and clause 8.

The results of these examples will be the development of a table that will show the maximum lobe length supported versus the number of passive concentrators in the ring. It is recommended that manufacturers of passive concentrators provide such a table in addition to the ring configuration(s) supported by the concentrator for the applicable media.

B.2.1 Single passive concentrator design example

Given the cabling architecture of figure B.1 and the conditions specified in table B.2, the maximum supportable cable length as well as the maximum supportable lobe length are determined. Also, lengths for the station, jumper, and concentrator cords are suggested. Two configurations are examined, one consisting of a single station, which is assumed to be the limiting configuration, and one consisting of three stations. In each example, the crosstalk noise introduced by the concentrator is considered.

B.2.1.1 Design procedure for ring 1—single concentrator—single station

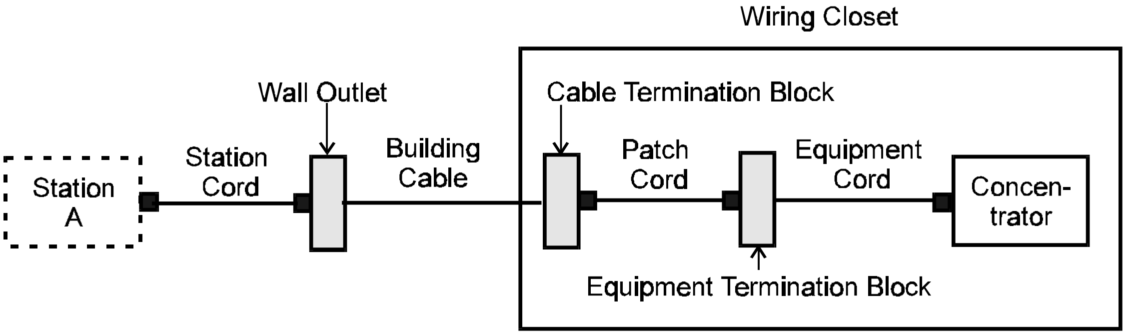

Figure B.1—Passive channel—single concentrator

Table B.2—Single passive concentrator example conditions

Ring speed	Initially 4 Mbit/s, changing to 16 Mbit/s in 6 months
Concentrator type	Passive type: 12 passive UTP ports and passive UTP RO and RI trunk ports Attenuation specifications (16 MHz)[a] CLSQA = 1.0 dB CTSQA = 0.8 dB CLATT = 0.4 dB CTATT = 0.8 dB
Number of users	Ring 1—10 users (single concentrator) Ring 2—30 users (three concentrators)
Media type	Category 5 UTP cabling system Transmission Parameters @ 16 MHz, 20° C Attenuation—8.2 dB/100 m (25 dB/kft) NEXT loss—44 dB (cable only) Characteristic impedance—100 Ω
Media connector type	UTP Category 5: NEXT Loss = 56 dB
Composite NEXT Loss of cable and connectors	CNEXT loss = 40.6 dB[b]

[a] Numbers have been chosen to highlight the impact of CTSQA and CTATT.

[b] For all examples in this subclause, simplifying conservative assumptions are made to compute the crosstalk noise. Specifically, the link NEXT loss, called “Composite NEXT Loss of Cable and Connectors” is assumed to be equal to the voltage sum of the crosstalk from the cable (based on fixed cable crosstalk specification) and the crosstalk from two mated pairs of connectors.

B.2.1.1.1 Design for 4 or 16 Mbit/s?

Although the network will initially be used at 4 Mbit/s, it will be designed for 16 Mbit/s to ensure that when the change to 16 Mbit/s takes place in six months the lobe lengths will support the higher speed.

B.2.1.1.2 Worst-case channel design

The design assumes a single station attached to the ring will experience the maximum channel insertion loss. The station transmit signals must traverse the complete passive channel path and return to its own receiver. The worst-case channel design is for long building cable. Therefore, the effects of crosstalk for channel components in the wiring closet at station A are assumed negligible due to lobe loss.

B.2.1.1.3 Determine flat loss of passive channel

Is the flat loss less than 2 dB? For this example, the only flat loss is CLATT = 0.4 dB; therefore, the flat loss is within the 2 dB specification of 7.2.4.2.1.

B.2.1.1.4 Determine the lobe cable insertion loss assuming the design is insertion loss limited

From 7.2.4.2.1 the maximum insertion loss for a 16 Mbits/s passive channel is 19 dB, and from 7.2.4.3.2 the minimum NIR (NEXT Loss-to-Insertion Ratio) for UTP is 15.5 dB. Figure 39 shows that a passive channel composite NEXT loss of at least 34.5 dB is required for a maximum loss channel of 19 dB. The total concentrator loss is encountered only when a single station is attached to the concentrator as shown in figure B.1.

$$\begin{bmatrix}\text{passive channel}\\ \text{insertion loss}\end{bmatrix}=\begin{bmatrix}\text{loss of cables}\\ \text{and cords}\end{bmatrix}+\begin{bmatrix}\text{loss of}\\ \text{connectors}\end{bmatrix}+\begin{bmatrix}\text{loss of}\\ \text{concentrator}\end{bmatrix}$$

Loss of the connectors can generally be neglected if there are less than 4 in the lobe and the insertion loss of each is less than 0.1 dB. For a single concentrator, since the transmit signals do not pass through the RO and RI ports, the only concentrator losses in the channel are CLSQA and CLATT.

Therefore,

$$\begin{bmatrix}\text{passive channel}\\ \text{insertion loss}\end{bmatrix} = \begin{bmatrix}\text{loss of cables}\\ \text{and cords}\end{bmatrix} + \left(\text{CLSQA}+\text{CLATT}\right)$$

or the maximum allowable loss of the cables/cords is

$$\begin{bmatrix}\text{loss of cables}\\ \text{and cords}\end{bmatrix} = 19 - \left(1.0 + 0.4\right) = 17.6\ \text{dB}$$

The lobe insertion loss is then 17.6/2 = 8.8 dB

B.2.1.1.5 Determine NIR

The requirement is > 15.5 dB to verify design is loss limited.

In order to determine NIR, the sources of crosstalk noise must first be determined, and then the composite NEXT loss of the channel at the station determined. The primary sources of crosstalk noise are the cabling system and the concentrator.

The crosstalk noise generated within the concentrator can be neglected since

a) The station transmitter signal, which is the signal source for the crosstalk noise generated within the concentrator, is 8.8 dB lower when it reaches the concentrator due to the insertion loss of the lobe cable.

b) When the concentrator noise reaches the station receiver, it is again 8.8 dB lower. Since the NEXT loss of the concentrator is 40 dB (8.4.4), the noise is 40 + 17.6 = 57.6 dB lower than the transmitter signal level, which, when compared to the NEXT noise due to the channel (connectorized cords/cables NEXT loss of 40.6 dB), is insignificant.

Therefore, the composite NEXT loss of the channel at the station is 40.6 dB (connectorized cords/cable) and NIR is 40.6 – 19 = 21.6 dB, which exceeds the requirement of 15.5 dB. Therefore, the design is insertion loss limited as assumed in B.2.1.1.4).

B.2.1.1.6 Conclusion

Since the cable and cords are all Category 5, the total cable length is (17.6 dB)/(8.2 dB/100 m) = 214 m or a maximum lobe cable/cord length of 107 m. This length may be divided as follows:

Station cord:	3 m
Patch cord:	3 m
Equipment cord:	3 m
Building cable:	98 m (at 20 °C)

If the attenuation of the cords is 20% higher than the cable attenuation,[30] or 9.8 dB/100 m, the category 5 equivalent electrical attenuation length of each cord is 3.6 m. Therefore, the length of the building cable must be reduced by 1.8–96.2 m. If the environmental temperature is above 20 °C, the cable lengths may need to be derated. (See ISO/IEC 11801 : 1995 for the derating factors for STP cables.)

B.2.1.2 Design procedure for ring 1—single concentrator—multiple stations

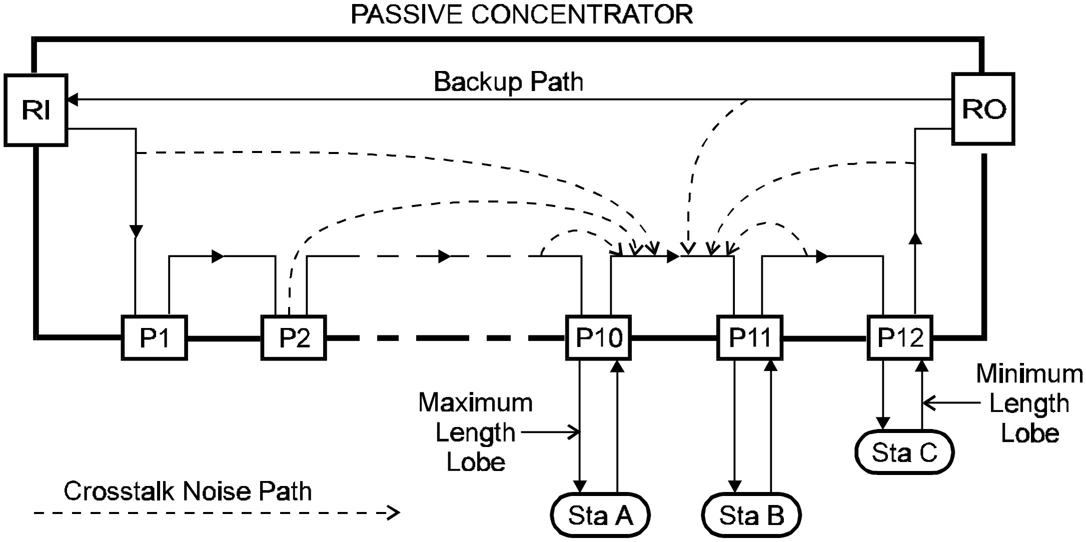

Figure B.2—Passive concentrator with three stations showing examples of crosstalk noise paths

The other single concentrator configuration, consisting of three stations, that may require the lobe cable loss of 8.8 dB for the single station design to be reduced is shown in figure B.2. Stations A and B are at the maximum lobe loss (length) of 8.8 dB, and station C is connected by a short cable to the concentrator. The additional crosstalk noise introduced in the lobe cable connected to port 10, which is due to the disturbing transmitter of station C, will appear at station B.

---

[30] 20% is used in this example. ISO/IEC 11801 : 1995 allows the attenuation of a cord to be 50% higher than that of the cable.

B.2.1.2.1 Determine flat loss of passive channel

Is the flat loss less than 2 dB? For this example, the flat loss of the passive channel between stations A and B is CLATT = 0.4 dB, which is within the 2 dB specification of 7.2.4.2.1.

B.2.1.2.2 Determine the lobe cable insertion loss

From the previous single station example, the loss of station A and station B lobe cables is 8.8 dB. The insertion loss of the passive channel between stations A and B is 2(8.8) + CLATT = 18.0 dB (CLSQA between ports 10 and 11 is assumed to be negligible due to proximity).

B.2.1.2.3 Determine NIR (requirement: > 15.5 dB)

In order to determine NIR, the sources of crosstalk noise must first be determined, and then the composite NEXT loss of the channel at the station determined. The two primary sources of crosstalk noise are as follows:

a) Crosstalk noise within the concentrator—The signal from station C occupies the backup ring path and a major portion of the main ring path. The concentrator NEXT loss is 40 dB (8.4.4). Assuming the loss of the cable that attaches station C to the concentrator port is 1.0 dB, the effective NEXT loss for the concentrator crosstalk noise seen at station B is 40 + 1 + 8.8 = 49.8 dB.

b) Crosstalk noise induced into the receive signal path of the cabling system attached to the next upstream concentrator lobe port (P10)—When port P10 transmits into its lobe cable, a portion of the signal couples (crosstalks) into its own lobe receive channel signal path. The worst-case crosstalk noise is induced when station C is attached directly to the concentrator with a short length lobe cable (1 dB loss). Assume the concentrator loss between port P12 receive and port P10 transmit is 1.4 dB (CLATT + CLSQA). From table B.2 the NEXT loss of the lobe cabling system is 40.6 dB (connectorized cords/cables). Crosstalk noise is coupled into the receive pair of P10, which appears at the transmit of port P11. The loss between these two ports is assumed to be 0 dB. When the noise reaches the receive port of station B, it is reduced by the lobe loss (8.8 dB). Therefore, the worst-case NEXT noise, or equivalent worst-case NEXT loss seen at station B is 40.6 + 1.4 + 0 + 8.8 = 50.8 dB.

The composite worst-case NEXT (CNEXT) loss at station B for the three noise sources is then

$$\begin{aligned}\text{CNEXT loss} &= 20\log\left(10^{(\text{Station\_Cable\_NEXT}/20)} + 10^{(\text{Conc\_NEXT}/20)} + 10^{(\text{P10\_Cable\_NEXT}/20)}\right)\\ &= 20\log\left(10^{(-40.6/20)} + 10^{(-49.8/20)} + 10^{(-50.8/20)}\right)\\ &= 20\log(0.0093 + 0.0032 + 0.0026)\\ &= 36.5\text{ dB}\end{aligned}$$

Therefore, NIR is 36.5 – 18.0 = 18.5 dB, which exceeds the requirement of 15.5 dB by 3.0 dB.

B.2.1.3 Summary of single concentrator design examples

The two single concentrator design examples show that the single station, which must accommodate the full concentrator insertion loss (CLSQA and CLATT) configuration, limits the lobe length. In both examples using category 5 UTP cable there is sufficient NIR margin that the maximum channel insertion loss of 19 dB can be used and the designer may neglect the effect of NEXT noise. Therefore, the maximum channel insertion loss and lobe insertion loss that can be supported for the single concentrator specified in this example are 19 dB and 8.8 dB, respectively.

B.2.2 Multiple passive concentrator design example

This example examines some of the design considerations when a network consists of multiple concentrators. The single concentrator network of figure A.1 is expanded to three concentrators in three different wiring closets. The maximum cable lobe loss and cable length that can be supported will be determined for the conditions specified in table B.2. Also, lengths for the station, patch, and concentrator cords will be suggested.

B.2.2.1 Design procedure for ring 2-multiple concentrators

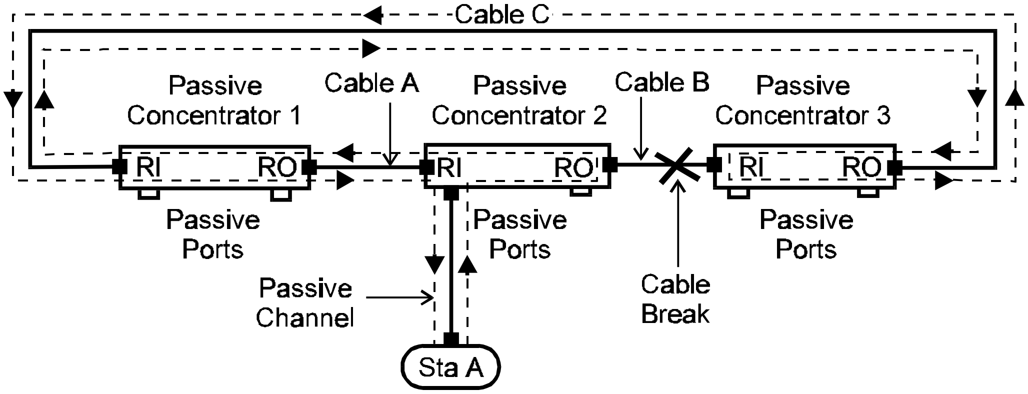

NOTE—To allow the removal and/or replacement of concentrators without disrupting the operation of the remainder of the ring, the maximum length passive channel should include the length of cable B. The example in this figure assumes this tolerance is not required.

Figure B.3—Maximum length passive channel

B.2.2.1.1 Design for 4 or 16 Mbit/s?

Although the network will initially be used at 4 Mbit/s, it will be designed for 16 Mbit/s to ensure that when the change to 16 Mbit/s takes place in six months the lobe lengths will support the higher speed.

B.2.2.1.2 Worst-case channel design

The design assumes a single station (station A) in figure B.3, which is attached to the ring concentrator in wiring closet B, will experience the maximum channel insertion loss when the shortest inter-concentrator cable (cable B) is defective and has been disconnected and bypassed at both ends. In order for a complete ring path to be provided, the RO port of concentrator 2 and the RI port of concentrator 3 must be in the "wrap" states. As shown by the dashed lines in figure B.3, the signals from the station must traverse the complete passive channel path and return to its own receiver.

B.2.2.1.3 Determine flat loss of passive channel

The flat loss of the channel consists of the lobe flat loss (CLATT) plus the trunk RI and RO flat losses (CTATT). As shown in figure B.3, the signal passes through only the trunk paths on two concentrators (Conc1 and Conc3), and through the trunk and lobe paths of one concentrator (Conc2). Since CLATT includes both the lobe and trunk losses of a single concentrator, the total channel flat loss is

$$\text{channel\_flat\_loss} = \text{CLATT} + 2\,(\text{CTATT}) = 0.4 + 2\,(0.8) = 2.0\ \text{dB}$$

This is equal to the 2 dB maximum allowable flat loss specified in 7.2.4.2.1. Since the flat loss equals 2 dB, no additional concentrators can be added to the network.

B.2.2.1.4 Determine the maximum allowable lobe cable insertion loss

From 7.2.4.2.1 the insertion loss for a passive channel operating at 16 Mbit/s is 19 dB and from 7.2.4.3.2, NIR for UTP is 15.5 dB. Figure 39 shows that for NIR of 15.5 dB, a channel composite NEXT loss of at least 34.5 dB is required for a channel insertion loss of 19 dB. The channel insertion loss is

$$\begin{bmatrix}\text{passive channel}\\ \text{insertion loss}\end{bmatrix} = \begin{bmatrix}\text{loss of cables}\\ \text{and cords}\end{bmatrix} + \begin{bmatrix}\text{loss of}\\ \text{connectors}\end{bmatrix} = \ 19 \text{ dB}$$

$$\begin{bmatrix}\text{passive channel}\\ \text{insertion loss}\end{bmatrix} = \begin{bmatrix}\text{2 lobe} + \text{ 2 cable A}\\ + \text{ 2 cable C}\end{bmatrix} + \begin{bmatrix}\text{C2\_SLSQA} + \text{C1\_STSQA}\\ + \text{ C3\_CTSQA} + \text{FLAT}\end{bmatrix} = \ 19 \text{ dB}$$

rewriting

$$\begin{bmatrix}\text{2 lobe} + \text{2 cable A}\\ \text{+2 cable C}\end{bmatrix} = \ 19 \text{ dB} - \begin{bmatrix}\text{C2\_SLSQA} + \text{C1\_STSQA}\\ + \text{ C3\_CTSQA} + \text{FLAT}\end{bmatrix}$$

$$2 \cdot \begin{bmatrix}\text{lobe} + \text{cable A}\\ + \text{ cable C}\end{bmatrix} = \ 19 - (1.0 \ + \ 0.8 \ + \ 0.8 \ + \ 2.0) \ = \ 14.4 \text{ dB}$$

The insertion loss of the lobe and two cables (A and C) is

lobe loss + cable A + cable C = 7.2 dB

and the lobe loss is

lobe loss = 7.2 - (cable A + cable C) = 7.2 – (0.082 dB/m)(8 m + 15 m) = 5.3 dB

B.2.2.1.5 Determine NIR (Requirement: > 15.5 dB)

As discussed in the single concentrator example, in order to determine NIR, the sources of crosstalk noise must first be determined, and then the composite NEXT loss at the station determined. In this example two sources of noise will be considered, the channel NEXT noise and the crosstalk noise within the concentrator. The channel NEXT loss at the station from table B.2 is 40.6 dB (connectorized cable/cord). From 8.4.4, the worst-case NEXT loss for a passive concentrator is 40 dB. From the previous example, the crosstalk noise generated within the concentrator as seen at the station is

conc_NEXT_loss + 2(lobe_loss) = 50.6 dB

Therefore, the channel composite worst-case NEXT at the station is

$$\begin{aligned}\text{CNEXT loss} &= 20\log(\ 10^{\text{(Conn\_Cable\_NEXT/20)}} + 10^{\text{(Conc\_NEXT/20)}}\ )\\ &= 20\log(\ 10^{(-40.6/20)} + 10^{(-50.6/20)}\ )\\ &= 20\log(\ 0.0093 + 0.003)\\ &= -38.2 \text{ dB}\end{aligned}$$

NIR is then 38.2 – 19.0 = 19.2 dB, which exceeds the requirement of 15.5 dB by 3.7 dB.

B.2.2.1.6 Conclusion

The maximum supportable lobe cable length is (5.3 dB)/(8.2 dB/100 m) = 65 m, which may be allocated as

Station cord:	3 m
Patch cord:	3 m
Concentrator cord:	3 m
Building cable:	56 m (at 20 °C)

Therefore, the maximum lobe length that can be supported for a ring consisting of 3 concentrators is 65 m. This distance is considerably less than the single concentrator example since the inter-closet cables (A and C), which total 23 m, are part of the passive channel path. Also, the trunk flat loss CTATT (0.8 dB/concentrator) reduces the lobe length by 5 m/concentrator. If the trunk flat loss (CTATT) is reduced to zero, the maximum lobe length increases from 65–75 m.

B.2.3 Passive concentrator with active trunk ports

The design of a passive channel for use with concentrators with passive ports and active trunk ports is the same as the design for the single passive concentrator example in B.2.1. The two concentrator losses of interest between the two stations are flat loss (CLATT) and frequency-dependent loss (CLSQA).

$$\begin{bmatrix}\text{passive channel}\\ \text{insertion loss}\end{bmatrix} = \begin{bmatrix}\text{loss of cables}\\ \text{and cords}\end{bmatrix} + (\text{SLSQA} + \text{CLATT})$$

The maximum lobe insertion loss that can be supported at 16 Mbits/s is

$$\text{lobe insertion loss} = 1/2\ \{19\ \text{dB} - (\text{CLSQA} + \text{CLATT})\}$$

B.2.4 Passive channel design conclusions

With the concentrator and media specified in table B.2, the maximum lobe length is limited by the single station configuration and the maximum ring size is limited to 3 concentrators by the trunk flat loss (CTATT) flat loss. If the ring is required to support more than 3 concentrators, then one of the following is recommended:

a) Use repeaters in the trunk circuit to add up to three additional concentrators. Two repeaters are required for each additional three concentrators.
b) Use concentrators with minimal flat loss CTATT. This maximizes the number of concentrators that a ring may can support.

The NEXT noise due to the concentrator may be ignored provided the NEXT loss is greater than 40 dB as specified in clause 8.

It is recommended that the concentrator manufacturer provide a table that shows the maximum supported cable length (budget) for the types of media specified in ISO/IEC 11801 : 1995 for passive concentrators. An example of cable budget table is shown in table B.3 for UTP Category 5 and STP Type 1A.

B.2.5 Cable budget table example for passive concentrators

This subclause provides an example of a cable budget table that it is recommended passive concentrator manufacturers provide as part of their documentation. Table B.3 is an example table, which shows the maximum cable budget for category 5 UTP and STP cable systems for the concentrator specified in table

B.2 when CTATT is zero. Figure B.4 shows a four concentrator network used to develop table B.3. Table A.3 includes cables (C1, C2, and C3 of figure B.4) between each concentrator, which have a two-way loss of 0.1 dB. Although a cable is shown between the first and last concentrator (cable C4), which is recommended, the table does not include the cable. Therefore, the table may be used for those configurations that provide cable C4, as well as those that do not provide cable C4. The equation for the table entries is

$$\text{cable\_budget} = [19 - (\text{CLATT}+\text{CLSQA}) - (n-1)(\text{CTATT}+\text{CTSQA}+0.1)]/2 \ \ (CA)$$

where CA is the cable attenuation (dB/m) and n is the number of concentrators.

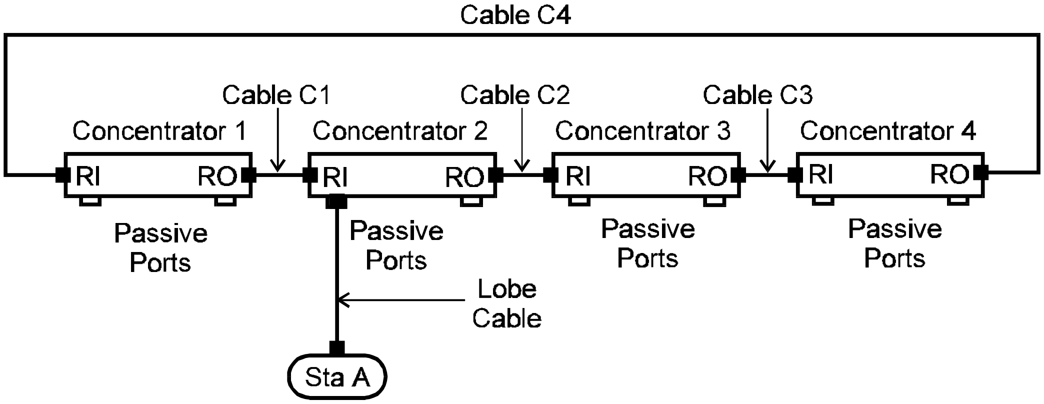

Figure B.4—Four passive concentrator network

Table B.3—Cable budget for passive concentrators @ 16 Mbit/s for UTP category 5 and STP cabling systems

Number of concentrators	1	2	3	4	5	6	7
Number of stations	12	24	36	48	60	72	84
Cable budget (m)—UTP (Category 5)	107	102	96	91	85	80	74
Cable budget (m)—STP[a]	196	185	175	165	155	145	135

[a] Media specifications for STP Type 1A of table B.4 were used to determine the lobe length for STP.

Table B.3 is used as follows:

Assume the network shown in figure B.4 is installed with CONC 1 and CONC 2 in wiring closet 1, CONC 3 in wiring closet 2, and CONC 4 in wiring closet 3. The inter-closet cable lengths are cable C2 = 23 m, cable C3 = 12 m, and cable C4 = 13 m. Assume the shortest cable (C3) is defective. From table B.5, the cable budget for four concentrators, when using UTP Category 5 cable, is 91 m. Therefore, the maximum length lobe cable that can be attached to any of the four concentrators is

$$\text{lobe cable length} = \text{cable\_budget} - (\text{cable\_C2} + \text{cable\_C4}) = 91 - (23 + 13) = 55 \text{ m}$$

Cable C1 is not considered since the loss of C1 is included in the cable budget.

B.3 Active channel design example

This clause provides a design example for an active channel. The design of an active channel is not as complex as the passive channel since the concentrator is not part of the channel. As was recommended for the passive channel, it is recommended that manufacturers of active concentrators provide the supported lobe lengths for the media specified in clause 7.

An active channel is used in two areas of a token ring network: 1) To interconnect a station and an active port of a concentrator and 2) to interconnect active RI and RO concentrator ports.

Figure B.5 shows the station-to-concentrator active channel, which consists of the components shown in table B.1.

B.3.1 Design procedure—active channel

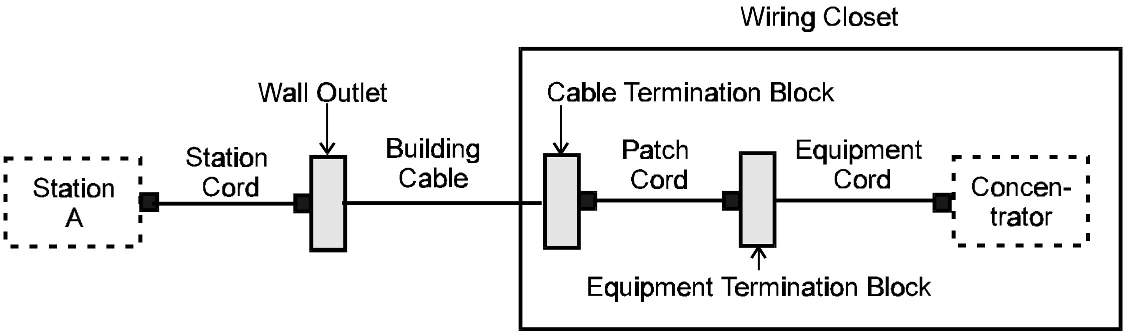

Figure B.5—Station-to-concentrator active channel

B.3.1.1 Design objectives

Design an active channel for 16 Mbit/s and determine the supportable lobe lengths for UTP Categories 3, 4, and 5, and STP cabling systems.

Table B.4—Single active concentrator example conditions

Ring speed	**16 Mbit/s**				
Concentrator Types	1) Active type with 16 UTP active ports 2) Active type with 12 STP active ports				
Media types[a]	STP		UTP		
Parameter	Type 1	Type 1A	Cat 3	Cat 4	Cat 5
Attenuation (dB/100 m)	4.5	4.4	13.1	8.9	8.2
NEXT loss cable (dB)	40.0	50.4	23.2	38.2	44.2
Composite NEXT loss of cable and connectors (dB)	35.7	46.9	22.0[b]	33.0	40.6
Impedance (Ω)	150	150	100	100	100

[a] Transmission parameters @ 16 MHz, 20° C

[b] NEXT loss for connectors is assumed to meet category 4 specifications.

B.3.1.2 Determine supportable lobe length

From 7.2.4.2.1 the insertion loss budget for an active 16 Mbit/s channel is 16 dB, and from 7.2.4.3.2 NIR for STP and UTP is 15.5 dB and 14.0 dB, respectively. To determine the lobe length that can be supported, the maximum supportable insertion loss for each medium is calculated using NIR (insertion_loss = NEXT_loss – NIR). The lesser of 16 dB and the insertion loss using NIR is used to calculate the supportable lobe length shown in table B.5.

Table B.5—Example of supported lobe lengths for active concentrators

Media type	NIR (dB)	NEXT loss (dB)[a]	Insertion loss (NEXT_loss – NIR) (dB)	Supported lobe	
				insertion loss (dB)	length (m)
UTP Cat 3	14	22	22 – 14 = 8	8	61
UTP Cat 4	14	33	33 – 14 = 19	16	180
UTP Cat 5	14	40.6	40.6 – 14 = 26.6	16	195
STP Type1	15.5	35.7	35.7 – 15.5 = 20.2	16	356
STP Type1A	15.5	46.9	46.9 – 15.5 = 31.4	16	364

[a] Composite NEXT loss of connectors and cables/cords.

B.3.1.3 Active channel design conclusions

The lobe lengths that can be supported when using an active concentrator are approximately twice as long as the lobe lengths when using passive concentrators. Four of the five media (UTP Categories 4 and 5, and STP Type 1 and STP Type 1A) are able to support the maximum specified insertion loss of 16 dB. Since active channel insertion loss does not include the concentrator losses, active concentrators can be mixed without penalty. As shown in the passive channel design of B.2, unless the flat loss and frequency-dependent losses of each vendor's concentrators are the same, table B.3 will be different depending on the mix of manufacturers' passive concentrators.

B.4 Use of alternate media—homogeneous 120 Ω cabling

This standard was developed, as specified in clause 7, for copper channel media of 100 and 150 Ω balanced cables. During this process it was demonstrated that homogeneous 120 Ω balanced copper cabling should, with proper restrictions, generally support the operation of ISO/IEC 8802-5 stations and concentrators designed for 100 Ω cabling systems. This clause describes the use of 120 Ω homogeneous cabling system when used with 100 Ω stations and concentrators.

Homogeneous, as used in this clause, implies that all cabling elements have a nominal impedance of 120 ± 15 Ω. The 120 Ω cable parameters are specified by ISO/IEC 11801 : 1995. The 120 Ω cabling system must be terminated by equipment designed for 100 Ω devices. For screened cabling, both the components of the cabling system and the terminating devices must provide screen continuity.

Generally, when using 120 Ω cabling, the reflection induced jitter will be greater than that induced in 100 Ω systems. With active retimed concentrators, this additional jitter is small and can be offset by the greater SCNR resulting from the lower attenuation of 120 Ω cabling for comparable UTP Category 5 channel lengths.

Although 120 Ω token ring operation with passive concentrators may operate satisfactorily, the requirement for robust operation using passive concentrators requires further study.

Annex C

(informative)

Jitter considerations

This annex defines and discusses the components in a token ring that are sources of jitter and is intended to provide information for token ring component designers that is helpful in minimizing ring jitter. These components are shown in figure C.1. Also, an example implementation is provided illustrating possible compromises between the various contributing factors for each jitter parameter specified in the standard.

The example chosen is representative of a typical implementation using phase-lock loops (PLLs) as the clock recovery circuitry; however, it should be emphasized that this should in no way be interpreted as restricting implementations; the only requirements for compliance to the standard are stated in the main body of the standard.

C.1 Jitter contributions

This clause defines five types of jitter and eye closure sources; transferred jitter, PLL jitter, transmitter jitter, channel jitter, and equalizer error.

To ensure an acceptable bit error rate (BER), the system must be designed such that the incoming data stream is sampled during the open portion of the received signal eye. The effects considered in this annex can be considered as either narrowing the received signal eye opening (eye-closure), or moving the sampling edge away from the center of the eye (alignment error).

In figure C.1, the PMC that receives the eye under consideration is referred to as the current PMC, and the PMC that transmits to this current PMC is referred to as the upstream PMC. The eye at the input to the current PMC is clocked (transmitted) by the upstream PMC’s recovered clock.

C.1.1 Transferred jitter (1)

Transferred jitter is defined as the alignment error at the current PMC caused by the jitter in the incoming data stream at the upstream PMC, relative to the clock in the active monitor. This jitter is composed of a correlated component (limited by the FAPS parameter) and an uncorrelated component (limited in combination with the uncorrelated jitter introduced by the upstream PMC by the AUJA parameter).

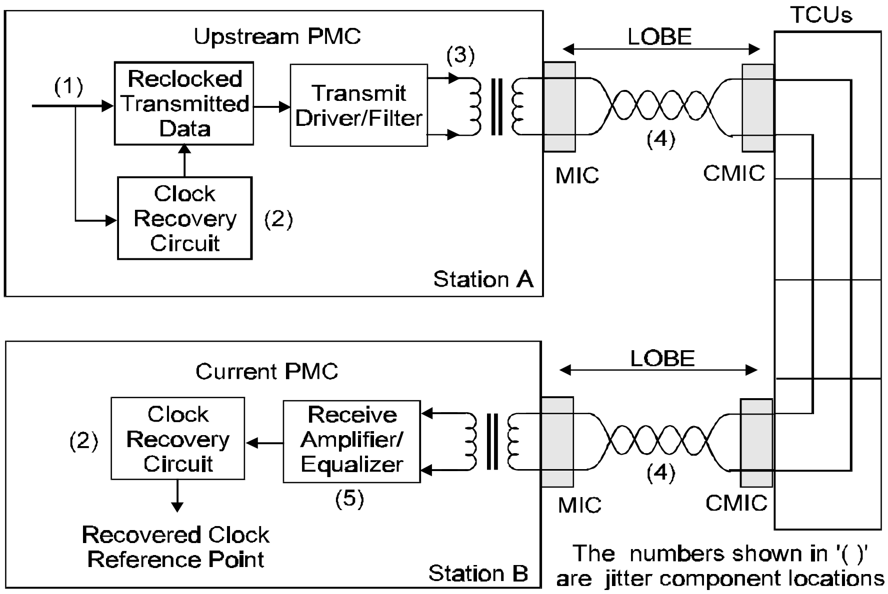

Figure C.1—Jitter components between stations

C.1.2 PLL jitter (2)

PLL jitter is defined as the alignment error produced by the jitter introduced by the upstream PMC's timing recovery circuitry. Typically, this jitter comprises a high-frequency component due to the presence of discrete charge pump pulses in the phase-lock loops often used for this application; and a random component due to noise in the timing recovery components. Both contributions are limited by the AUJA parameter.

The current PMC's timing recovery circuitry introduces the same type of jitter; in this case, the PLL jitter is limited by the JTOL parameter in addition to the AUJA parameter.

Due to imperfections in the clock recovery circuit, there may also exist an additional constant alignment error term at the current PMC, referred to as "static timing error." This is defined as the time between the center of the received eye and the clock edge used to sample the incoming data for a random data stream, averaged over a sufficient period so that all effects of uncorrelated and correlated jitter are removed.

C.1.3 Transmitter jitter (3)

Transmitter jitter is defined as the eye-closure produced by the transmit circuitry of the upstream PMC. This parameter, which is almost entirely correlated jitter, is limited by Tdiff01 and Tdiffmax parameters, and transmit duty cycle requirements.

C.1.4 Channel jitter (4)

Channel jitter is defined as the jitter introduced by the channel between the output of the MIC at the upstream PMC and the input to the MIC at the current PMC. Contributions to this jitter include any noise, crosstalk, or signal distortion introduced as a result of all connectors, cables, and passive TCUs in the channel.

Channel noise, crosstalk, and reflections that are not included in the channel specifications of 7.2.4 are treated as additional degradations to the test channel degradations for calculating the jitter budget. The dispersion characteristics of the cable in the test channel are assumed to be those of a perfect square-root loss cable and any deviations from this ideal for actual cables are similarly treated as additional degradations in determining the jitter budgets.

C.1.5 Equalizer error (5)

Equalizer error is defined as the eye-closure that results from the current PMC receiver circuitry's imperfect equalization of the phase and frequency response of the transmit filter, receive filter, and channel.

This jitter is almost entirely correlated with the data stream and is composed of potentially low-frequency components caused by the steady-state error between two worst-case data streams (often data streams consisting of all Data-ones and all Data-zeros), and of high-frequency intersymbol interference (ISI) components. The former lower-frequency jitter is limited by the FAPS and AJ parameters; the latter higher-frequency jitter by the jitter tolerance parameter JTOL.

C.2 Jitter budgets

C.2.1 Jitter budgets—elastic buffer elasticity and maximum PMC count

The maximum number of PMCs that can be supported on a ring, which includes repeaters, stations, and retiming concentrators, is limited by the size of the elastic buffer in the active monitor. The maximum accumulated jitter must be less than half the elasticity of the elastic buffer. On the other hand, the maximum number of MAC processes on a ring is limited by the timers used in the neighbor notification process, which is a maximum of 250. Therefore, the maximum number of stations that may be supported on a single ring is the lesser of 250 or the maximum permitted by the accumulated jitter.

The total accumulated jitter around the ring is limited by the requirements of 7.1.1:

	4 Mbit/s	16 Mbit/s
1) minimum elasticity	3 symbols	15 symbols
	750 ns	937 ns

This total accumulated jitter is composed of the accumulated jitter in the absence of crosstalk and reflections (AJ) and the accumulated jitter due to crosstalk and reflections (additional jitter). AJ is limited by the requirements of 7.1.1; the remainder of the elasticity is available for the other accumulated jitter sources.[31]

2) AJ	600 ns	750 ns
3) additional jitter	150 ns	187 ns

The maximum number of PMCs on a single ring is determined by the combination of the average contributions of AJ and the additional jitter at each PMC in the ring. Typically, AJ is dominated by the correlated jitter component due to equalizer error, and the additional jitter is dominated by the correlated jitter due to crosstalk. Under worst-case assumptions, the jitter accumulates linearly in the ring (note that this is only true for certain types of correlated data patterns) and thus the average contribution to AJ at a PMC, termed

---

[31] Note that the use of jitter-reducing devices in the ring may require allocation of some of the elasticity to accommodate re-centering of the elastic store in the jitter-reducing devices.

4) half-width of perfect eye	62.5 ns	15.6 ns
5) transmit asymmetry	6.0 ns	1.5 ns

the mean filtered correlated jitter (mFCJ)[32] and the average contribution to the additional jitter, termed the mean additional jitter (mAJ), can be added on a linear basis and the result divided into the total latency to determine worst-case maximum station count. Typical values of mFCJ and mAJ are as follows:

	4 Mbit/s	**16 Mbit/s**
6) mFCJ	2.0 ns	1.5 ns
7) mAJ	0.5 ns	1.0 ns
8) accumulated jitter/PMC	2.5 ns	2.5 ns

Therefore the

9) maximum number of PMCS is	300	375

Two examples are given of rings satisfying this requirement; first, for a ring built with passive concentrators, but with a small number of repeaters, where the maximum number of stations is limited by the MAC processes to 250; and second, for a ring using retiming concentrators where the maximum number of stations is limited by the requirements of the PMC layer.

Example 1: Passive concentrators

10) number of stations	250	250
11) possible number of other repeating elements	50	125

Example 2: Retiming concentrators

12) number of stations	144	180
13) number of retiming TCU ports	144	180
14) possible number of other repeating elements	12	15

C.2.2 Jitter budgets—filtered accumulated phase slope (FAPS)

The FAPS is limited by the requirements of 7.1.1; and is defined in the absence of any jitter contributions from crosstalk or reflections. The FAPS is determined by the transfer function of the timing recovery circuitry used, scaled by mFCJ.

In the case of a heavily over-damped phase-lock loop with minimal jitter peaking, the maximum phase slope builds rapidly over a small number of PHYs (typically 10–20) and shows only a small increase thereafter. This restriction also facilitates the calculation of the jitter budgets. For these reasons, it is recommended that the maximum jitter gain of the timing recovery circuit is limited to

15) jitter gain	<1.00069	<1.00058
	0.006 dB	0.005 dB

In this case; the FAPS is given to a good approximation by the following formula:

$$\mathrm{FAPS} = \frac{\mathrm{mFCJ}}{\frac{1}{K} - \mathrm{NDT}}$$

---

[32] mFCJ is filtered to remove uncorrelated jitter and high-frequency jitter that does not contribute to jitter accumulation.

where

K is the effective timing recovery bandwidth (in radians/s) of the PLL. The effective bandwidth, *K*, of the timing recovery circuitry, can be defined conveniently in terms of the maximum filtered alignment error produced at any PHY by the application of a constant phase slope to the input data stream. In the case where the data stream consists of a long series of Data-zero symbols (to remove all correlated jitter) and the alignment error is filtered and averaged to remove the effects of PLL jitter, the *K* value of the data recovery circuitry is given by the ratio of the gradient of the incoming phase slope to the maximum alignment error observed. For a first-order phase-lock loop, *K* is given by 2π times the 3 dB bandwidth of the loop (in Hz).

NDT is the net delay time of the PHY (the average bit-delay time minus the average phase delay through the PHYs when they are repeating data using the timing recovered from the received data).

Note that NDT is typically a positive number but may be negative in some implementations.

Typically, taking the *K* value as that of a perfect first-order phase-lock loop of 360 kHz (16 Mbit/s) or 90 kHz (4 Mbit/s) 3 dB bandwidth:

	4 Mbit/s	**16 Mbit/s**
16) NDT	500 ns	125 ns
17) K	0.565 Mrad/s	2.26 Mrad/s
18) FAPS	0.0035 ns/ns	0.0082 ns/ns

Therefore,

19) maximum allowable mFCJ	4.4 ns	2.6 ns

It should be noted that these maximum figures are greater than those assumed in 6), and that, as a result, any implementations that specify an mFCJ at the upper end of this range may not be able to support the maximum number of stations calculated in 10) or 12).

Also specified in 7.1.1 is the parameter DFAPS; in this example, nFAPS (as specified in C.2.5) is used:

20) DFAPS (ΔFAPS for 1 ns mFCJ)	0.00078 ns/ns	0.0032 ns/ns

which is well within the limits of 0.0012 ns/ns and 0.0048 ns/ns at 4 Mbit/s and 16 Mbit/s, respectively.

C.2.3 Jitter budgets—phase slope tolerance (JTOL)

The tolerance of a single PMC to an incoming phase slope is specified by 7.2.3.1 as

21) JTOL	0.007 ns/ns	0.0172 ns/ns

Assuming the *K* values from 17, and by the definition of *K* given above, this corresponds to a maximum alignment error in the current PMC of

22) tracking error	12.4 ns	7.6 ns

The half-width of the eye received at the current PMC is reduced from the ideal by the effects of transmit duty cycle distortion (specified in 7.2.2.1), and the ISI caused by imperfect equalization in the receiver for the effects of the transmit filter, channel, and receive filter. The full transmit asymmetry is considered to narrow the half-width of the received eye to accommodate those designs of clock-recovery circuit sensitive to only one polarity of incoming data edge:

	4 Mbit/s	**16 Mbit/s**
23) half-width of perfect eye	62.5 ns	15.6 ns
24) transmit asymmetry	6.0 ns	1.5 ns
25) resultant ISI	5.0 ns	2.0 ns

Therefore,

26) resultant half-eye width	51.5 ns	12.1 ns

The alignment error causing the offset of the sampling edge from the center of this eye is built up from the static timing error in the phase-lock loop, the filtered correlated jitter produced by the equalizer and channel characteristics, and the PLL jitter in the current PMC. With the definition of static timing error used, mFCJ appears as a peak-to-peak alignment error with an average value of zero. Typically,

27) static timing error	10 ns	2.0 ns
28) mFCJ/2 (peak value of filtered correlated jitter)	1.0 ns	0.75 ns
29) PLL jitter/2 (peak value of PLL jitter)	2.0 ns	0.8 ns

Therefore, the

30) margin available for tracking error [26) – 27) + 28) + 29)]	38.5 ns	8.5 ns

which provides a jitter margin [30) – 22)] of 26.1 ns at 4 Mbit/s and 0.95 ns at 16 Mbit/s.

The JTOLX parameter is designed to measure the phase slope tolerance in the presence of a degree of eye-closure caused by an interfering signal. For this example, the half-eye is closed 6 ns by the interfering signal.

31) eye-closure due to interfering signal	6.0 ns	6.0 ns
32) new resultant half-eye width [26) – 31)/2]	48.5 ns	9.1 ns

Therefore, the margin available for tracking during JTOLX testing is reduced to

33) margin available for tracking [32) – 27) + 28) + 29)]	35.5 ns	5.5 ns

The phase slope required to be tracked during this test is correspondingly reduced to JTOLX

34) JTOLX	0.007 ns/ns	0.0082 ns/ns

which, in the same timing recovery circuit as previously assumed, requires a tracking error of

35) tracking error	12.4 ns	3.6 ns

Therefore, there is a margin of 23.1 ns at 4 Mbit/s and 1.9 ns at 16 Mbit/s [33) – 35)].

C.2.4 Jitter budgets—real network implementations

The difference between the specified minimum single PMC phase slope tolerance (JTOL) and the specified maximum accumulated phase slope (FAPS) is required to accommodate the effects of accumulated uncorrelated jitter, crosstalk, and reflections.

The effects of the accumulation of uncorrelated jitter are limited by the requirements of 7.1.2 and the effects of crosstalk and reflections by analysis of the channels described in annex B. Use of channels other than those recommended in this annex may result in unacceptable error rates.

C.2.5 Outline of an implementation specific DFAPS measurement

DFAPS is a limitation on the effective bandwidth of the PMC timing recovery circuits. The AJ and FAPS measurements are made on channels that do not contain some sources of jitter such as crosstalk that may be present in real channels. A PMC with a good equalizer may produce a very small amount of correlated jitter and thus present excellent FAPS results even though its timing recovery circuit may have a very high bandwidth. This presents a potential problem when jitter is added that the equalizer cannot compensate for since the high bandwidth of the timing recovery circuit will then produce large amounts of phase slope. Additionally, some timing recovery circuits may have bandwidths that are a function of the amplitude of the input jitter and for larger jitter or phase slope may show behavior different from that in the FAPS measurement. Thus the DFAPS test ensures that the bandwidth of the timing recovery circuit is limited over the range of jitter that the PMC will deal with in real environments.

For implementations with timing recovery well modeled by a linear phase transfer function over the range of normal operation and with nonzero correlated[33] jitter for at least one choice of a specific Test_Chan_A, the following measurement could be used to verify conformance to DFAPS:

a) A ring with the specific test channel conformant with Test_Chan_A is constructed with the number of PMCs specified in table 32 or with a number sufficient to extrapolate to the result after the number specified in table 32.

b) A data pattern that causes time-isolated correlated jitter phase steps[34] at each PMC that accumulates through the PMCs.

c) The total correlated jitter and correlated jitter phase slope are measured for this pattern. To determine the correlated jitter and phase slope components of the accumulated jitter, it is necessary to average repeated measurements of the phase vs. time of the jitter triggered off a consistent point in the repeated data pattern. This removes the uncorrelated part of the jitter. The jitter waveform is then filtered by the filter pole F6 described in 7.1.1 to remove high-frequency components of the jitter typically generated by the PMC previous to the measurement and that typically do not accumulate in the ring. The change in the resultant jitter over the period of the phase slope is then the total correlated jitter and the maximum value of the magnitude of the derivative of the resultant jitter over the period of the phase slope is the correlated jitter phase slope.

d) The total correlated jitter is then divided by the number of PMCs to get the correlated jitter per station. The correlated jitter phase slope is then divided by the correlated jitter per station to yield the normalized phase slope.

e) A value of the normalized phase slope less than nFAPS below then meets the requirements of the DFAPS specification (nFAPS=DFAPS/1 ns).

	4 Mbit/s	**16 Mbit/s**
nFAPS	0.0012	0.0048 ns/ns

---

[33] Correlated means jitter that is correlated to the data pattern as opposed to jitter that is not such as jitter caused by external noise.

[34] Time-isolated correlated jitter phase steps are phase waveforms that have constant phase values on either side of the step in phase for a period of time greater than *n* times 1 over the effective bandwidth (in rad/s) of the timing recovery circuitry of the PHYs. This time isolation is necessary to ensure that the transient response of the chain of timing recovery circuits to the phase step is not corrupted by transient responses to other changes in the phase waveform.

f) This measurement can give misleading results if the correlated jitter per station is small such that measurement inaccuracies in determining the correlated jitter per station results in large uncertainties in the calculation of nFAPS. It can also be misleading where the actual correlated jitter at each station varies widely around the correlated jitter per station calculated above since this can result in apparent higher nFAPS numbers due to the bias in measuring the maximum phase slope at a point where several stations with higher correlated jitter are clustered together yet the correlated jitter per station normalization reflects a lower per station jitter.

C.2.6 Clarification of NDT specification

Figures C.2 and C.3 show the typical behavior of the magnitude of the FAPS and its derivative as a function of the net delay time (NDT). When the net delay time approaches one over the effective bandwidth of the timing recovery circuitry, the phase slope increases rapidly. It is desirable to operate in a region well away from the discontinuity (at 400 ns in the example) to ensure that FAPS are at reasonable levels in a homogeneous ring. It would be possible to operate on either side of the discontinuity but analysis has shown that in a heterogeneous ring with some PHYs operating with NDTs lower than the critical value and some operating with NDTs higher than the critical value, it is possible to end up with the ring operating at the discontinuity. Thus it is necessary to ensure that all stations operate on one side of the discontinuity. For historical reason, the lower NDT side (left) has been chosen. Ensuring that the derivative of the magnitude of the FAPS is a positive as shown will ensure that the PHY is operating in the correct region. This can be tested in principle by adding data latches timed off the recovered clock to the transmit path and measuring the effect on phase slope; however, implementation details may make it impossible to test directly. Conformance can be shown by presenting design data or measurement of NDT and BW that show that under normal operating conditions NDT <1/BW.

$$\mathrm{MFAPS(mFCJ,\ BW,\ NDT)} = \left| \mathrm{mFCJ} \frac{1}{\frac{1}{\mathrm{BW}} - \mathrm{NDT}} \right| \qquad \left(\begin{array}{l} \text{Magnitude of FAPS is a function of} \\ \text{mFCJ, NDT, and retiming bandwidth.} \end{array} \right)$$

$$\mathrm{dMFAPS(mFCJ,\ BW,\ NDT)} = \frac{|\mathrm{mFCJ}|}{\left[\left| \frac{1}{\mathrm{BW}} - \mathrm{NDT} \right| \cdot \left(\frac{1}{\mathrm{BW}} - \mathrm{NDT} \right) \right]} \qquad (\text{Derivative of magnitude of FAPS})$$

$$i = 0...100 \qquad \mathrm{NDT}_i = \frac{1}{100} \cdot 1000 \cdot 10^{-9} \qquad (\text{Setup range of NDTs from 0 to 1 ms})$$

$$\mathrm{BW} = 2 \cdot \mathrm{p} \cdot 400 \cdot 10^3 \qquad (400\ \text{kHz bandwidth PLL})$$

$$\mathrm{mFCJ} = 1 \qquad (1\ \text{unit of jitter})$$

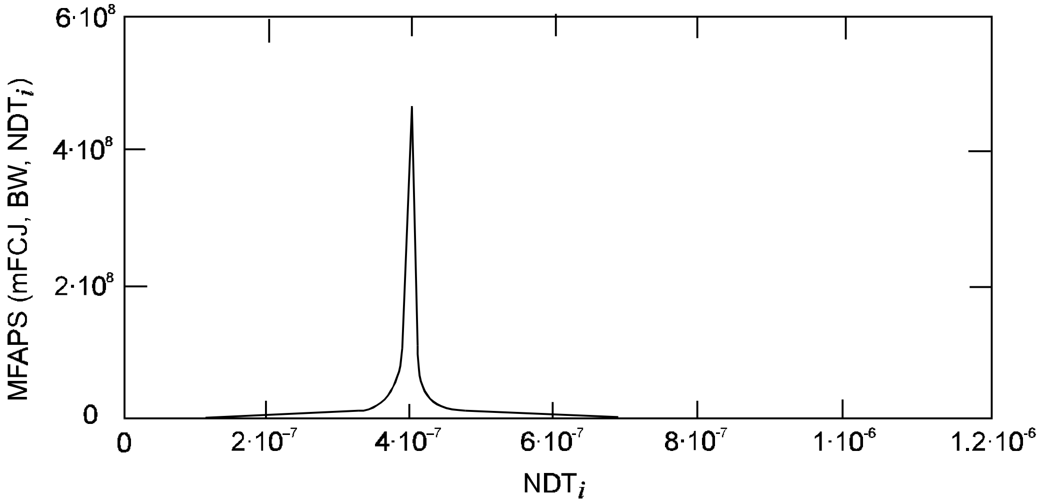

Figure C.2—MFAPS (mFCJ, BW, NDT$_i$)

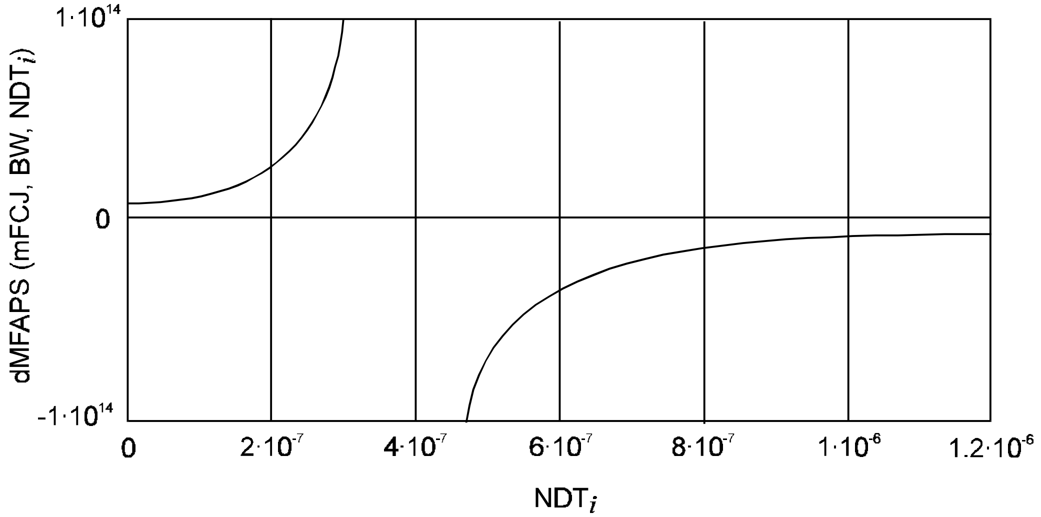

Figure C.3—dMFAPS (mFCJ, BW, NDT$_i$)

C.2.7 Model for calculation of accumulated uncorrelated alignment error

The following model can be used to determine the accumulated uncorrelated jitter alignment error (AUJA) from a chain of retiming elements. The model assumes an input sequence of time stamps (every signal element) from the last clock recovery circuit in the chain during the data patterns specified in 7.1.2.

$$\mathrm{OUT}_n = \mathrm{OUT}_{n-1} * X + \mathrm{IN}_n - \mathrm{IN}_{n-1}$$

where

$X = e^{-(U\,2\mathrm{p}F6)}$

OUT_n = AUJA at n

OUT_{n-1} = previous AUJA at $n-1$

IN_n = input time stamp data at n

IN_{n-1} = input time stamp data at $n-1$

n = index to time stamp

$F6$ = 360 kHz for 16 Mbit / s; 90 kHz for 4 Mbit / s

$U = 31.25 \cdot 10^{-9}$ s for 16 Mbit / s; $125 \cdot 10^{-9}$ s for 4 Mbit / s

Annex D

(informative)

Transmitter example

An example of a filter that may meet the transmitter masks and Tdiff requirements when applied to a square wave output is described by the following magnitude and differential time delay specifications. This filter is provided only as an example and is not to be considered a specification of the transmitter. The full specification of the transmitter is in 7.2.2 of the main body of the specification.

Table D.1—Transmitter insertion loss specification

Frequency (MHz)	Maximum insertion loss (dB)	Minimum insertion loss (dB)
0.05	3.0	—
4	0.5	—
8	0.5	—
16	0.8	—
24	1.7	—
30	—	1.0
32	—	7.0
36	—	15.0
44	—	30.0
48	—	30.0

Table D.2—Transmitter delay specification

F1 (MHz)	F2 (MHz)	Maximum differential delay (DD) (ns)	Minimum DD (ns)
2	4	0.4	0.0
8	4	0.0	–0.3
8	16	1.25	0.5
8	24	4.8	2.5
8	28	9.5	4.3

Differential delay (DD) is the delay between frequencies *F*1 and *F*2 defined as follows:

$$\mathrm{DD}\left[F1,\ F2\right] = \frac{\mathrm{PHASE\_at\_}F1}{360 \cdot F1} - \frac{\mathrm{PHASE\_at\_}F2}{360 \cdot F2}$$

where *F*1 and *F*2 are in Hz, PHASE is in degrees, and DD is in seconds.

Annex E

(informative)

Safety and environmental considerations

E.1 Network safety

This annex sets forth a number of recommendations and guidelines related to safety concerns; the list is neither complete nor does it address all possible safety issues. The designer is urged to consult the relevant local, national, and international safety regulations to ensure compliance with the appropriate requirements. For further information on safety issues, refer to IEC 60060 [B2] and IEC 60950 : 1991.

LAN cable systems described in this annex are subject to at least four direct electrical safety hazards during installation and use. These hazards are as follows:

a) Direct contact between LAN components and power, lighting, or communications circuits.
b) Static charge buildup on LAN cables and components.
c) High-energy transients coupled onto the LAN cable system.
d) Voltage potential differences between safety grounds to which various LAN components are connected.

Such electrical safety hazards must be avoided or appropriately protected against for proper network installation and performance. In addition to provisions for proper handling of these conditions in an operational system, special measures must be taken to ensure that the intended safety features are not negated during installation of a new network or during modification or maintenance of an existing network. Fault tolerance requirements are defined further in this annex.

It is assumed that the equipment that is attached to the token ring network is properly grounded and not left floating or serviced by “a doubly insulated ac power distribution system.” The use of floating or insulated equipment and the consequent implications for safety are beyond the scope of this standard.

E.1.1 Installation

Care should be taken to follow sound installation practice, as defined by any applicable local codes and regulations. In particular, care should be taken to ensure that uninsulated network cable conductors do not make electrical contact with unintended conductors or frame ground.

E.1.2 Telephony voltages

The use of building wiring brings with it the possibility of wiring errors that may connect telephony voltages to token ring equipment. Other than voice signals (which are low voltage), the primary voltages that may be encountered are the “battery” and ringing voltages. Although there is no universal standard, the following maximums generally apply. Battery voltage to a telephone line is generally 56 Vdc applied to the line through a balanced 400 Ω source impedance. Ringing voltage is a composite signal consisting of an ac component and a dc component. The ac component is up to 175 V peak at 20–60 Hz with a 100 Ω source resistance. The dc component is 56 Vdc with a 300–600 Ω source resistance. Large reactive transients can occur at the start and end of each ring interval. Although token ring equipment is not required to survive such wiring hazards without damage, application of any of the above voltages should not result in any safety hazard.

NOTE—Wiring errors may impose telephony voltages differentially across token ring transmitters or receivers. Because the termination resistance likely to be present across a receiver's input is of substantially lower impedance than an off-hook telephone instrument, receivers will generally appear to the telephone system as off-hook telephones. Therefore, full ring voltages will be applied only for short periods. Transmitters that are coupled using transformers will similarly appear like off-hook telephones (though perhaps a bit more slowly) due to the low resistance of the transformer coil.

E.2 Fault tolerance

The receiver and transmitter should be designed to tolerate the application of short circuits between all leads at the MIC for an indefinite period of time and should resume normal operation after such faults are removed. The magnitude of the current through any such fault should not exceed 300 mA.

E.3 Environment

E.3.1 Electromagnetic compatibility

The twisted-pair link should comply with applicable local and national codes for the limitation of electromagnetic interference.

E.3.2 Temperature and humidity

The station and lobe media are expected to operate over a reasonable range of environmental conditions related to temperature, humidity, and physical handling (such as shock and vibration). Specific requirements and values for these parameters are considered to be beyond the scope of this part of ISO/IEC 8802. It is recommended that manufacturers indicate, in the literature associated with the station attachment, the operating environmental conditions over which the specifications of this standard will be met to facilitate selection, installation, and maintenance.

Annex F

(informative)

Low-level finite state machines (FSMs)

This annex presents the FSMs contained in figures 23 through 25 of clause 4 in the same detail as those contained in ISO/IEC 8802-5 : 1992, figures 4-3, 4-5, and 4-7. An overview of the changes made by this part of ISO/IEC 8802 (refer to annex G for a summary of the changes) to the ISO/IEC 8802-5 : 1992 FSMs is as follows:

a) Transitions used in the Standby Monitor FSM (figure 4-5) to join the ring are now in the Join Ring FSM, figure F.2.
b) The Operational FSM (figure 4-3) is now the Transmit FSM, figure F.3.
c) Transitions used in the Standby Monitor FSM (figure 4-5) and in the Active Monitor FSM (figure F.4) for active and standby monitor functions are now in the Monitor FSM, figure F.4.

F.1 Interpreting the FSMs

The Join, Transmit and, Monitor FSMs can operate simultaneously. Each FSM can have only *one* State active at a time.

Figure F.1 illustrates states and transitions as follows:

a) STATES as vertical lines with the state name at the top of the line.
b) TRANSITIONS as horizontal lines containing a number to indicate state/transition number (e.g., 12) where the first digit is the state being exited and the second digit is the state being entered.
c) TRANSITIONS occur when the required condition(s) is(are) met and the ARROWHEAD indicates the state entered. For example, TRANSITION 13A exits STATE 1 and enters STATE 3.
d) Some transitions are awkward to illustrate because of location. TRANSITION 31 exits STATE 3 and enters STATE 1.
e) OPTIONAL conditions are enclosed in brackets (e.g., [condition 6]).

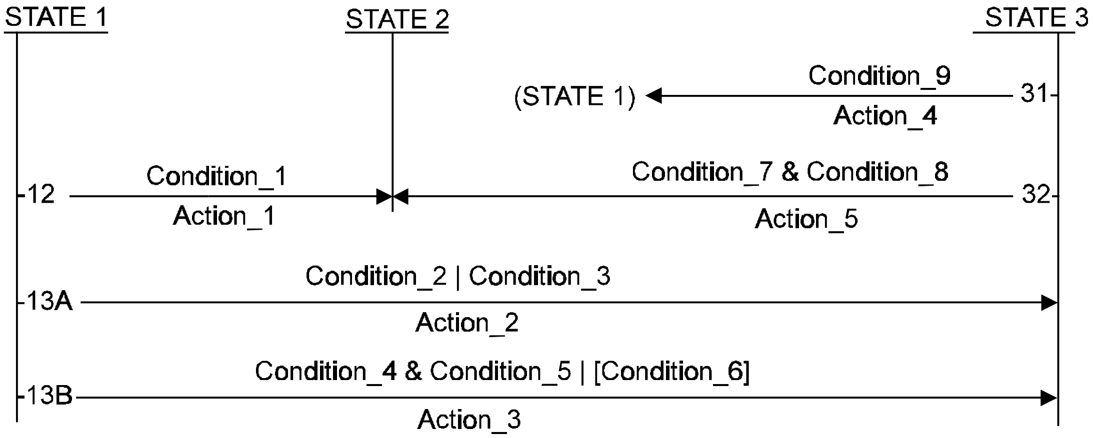

Figure F.1—State/transitions (S/T) illustration

Figure F.1 illustrates the following state transitions:

a) STATE 1 to STATE 2 (transition 12) due to Condition_1, executes Action_1.
b) STATE 1 to STATE 3 (transition 13A) due to Condition_2 OR Condition_3, executes Action_2.
c) STATE 1 to STATE 3 (transition 13B) due to Condition_4 AND Condition_5 OR optional Condition_6, executes Action_3.
d) STATE 3 to STATE 1 (transition 31) due to Condition_9, executes Action_4.
e) STATE 3 to STATE 2 (transition 32) due to Condition_7 AND Condition_8, executes Action_5.

Clause 4, table 6 EVENT/CONDITION column contains the current state, but it is not shown in figures F.2 through F.4 since the current state is illustrated by the vertical line.

Table 6 also contains transitions without a state/transition (S/T) entry and, although these transitions are not shown in the figures F.2 through F.4 FSMs, they are required for proper operation.

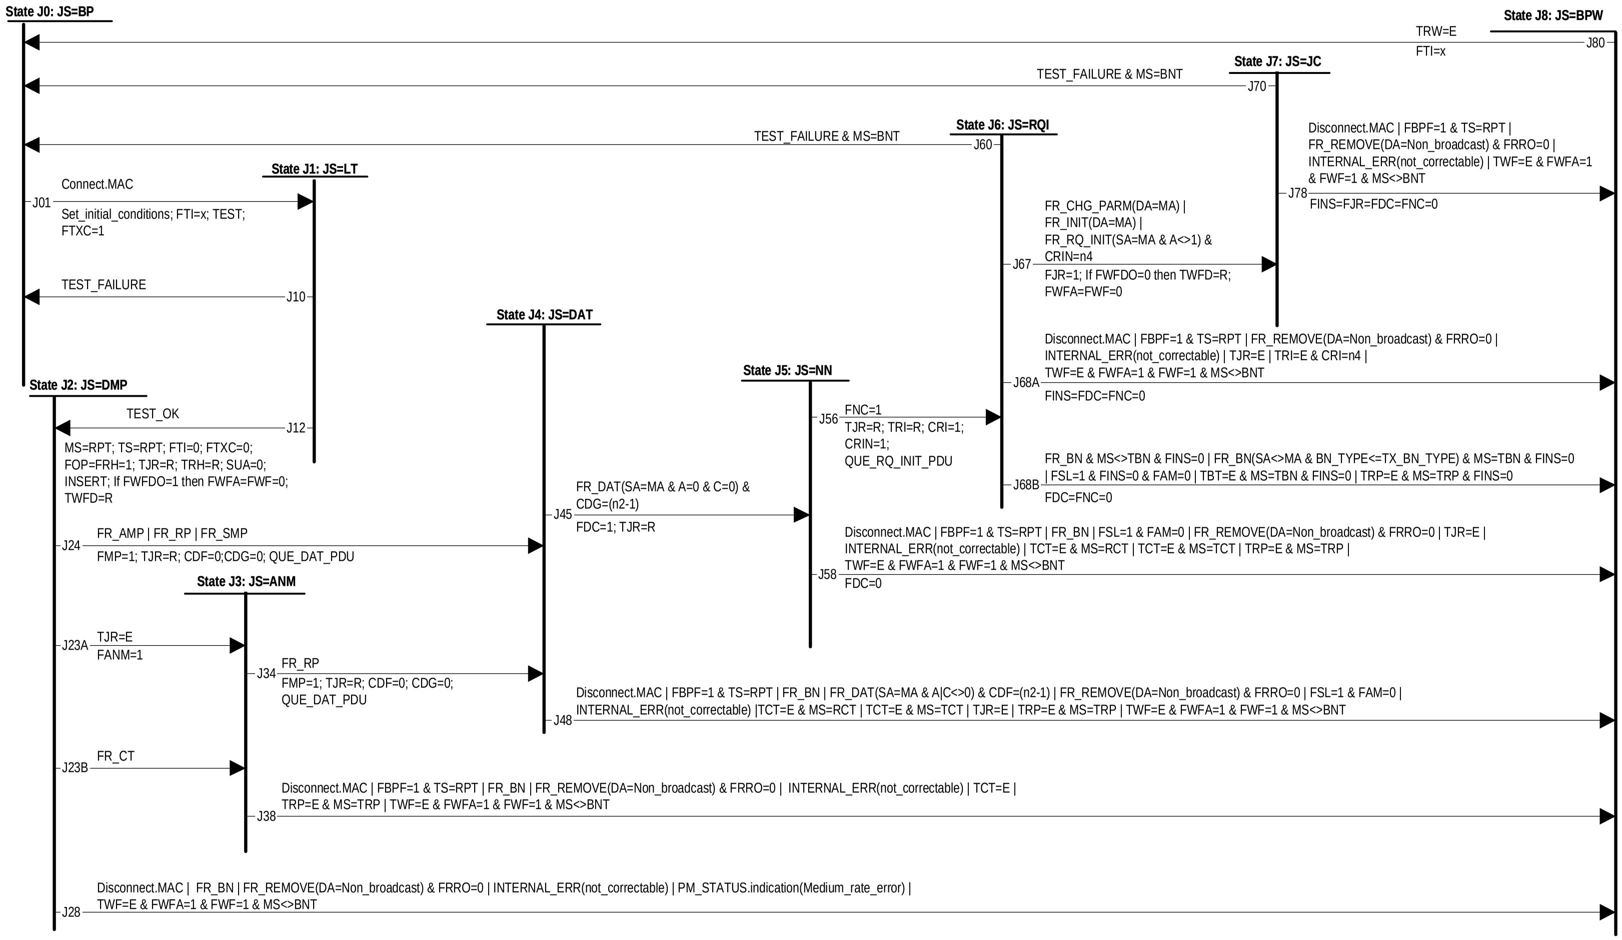

Figure F.2—Join (JS=state) finite state machine

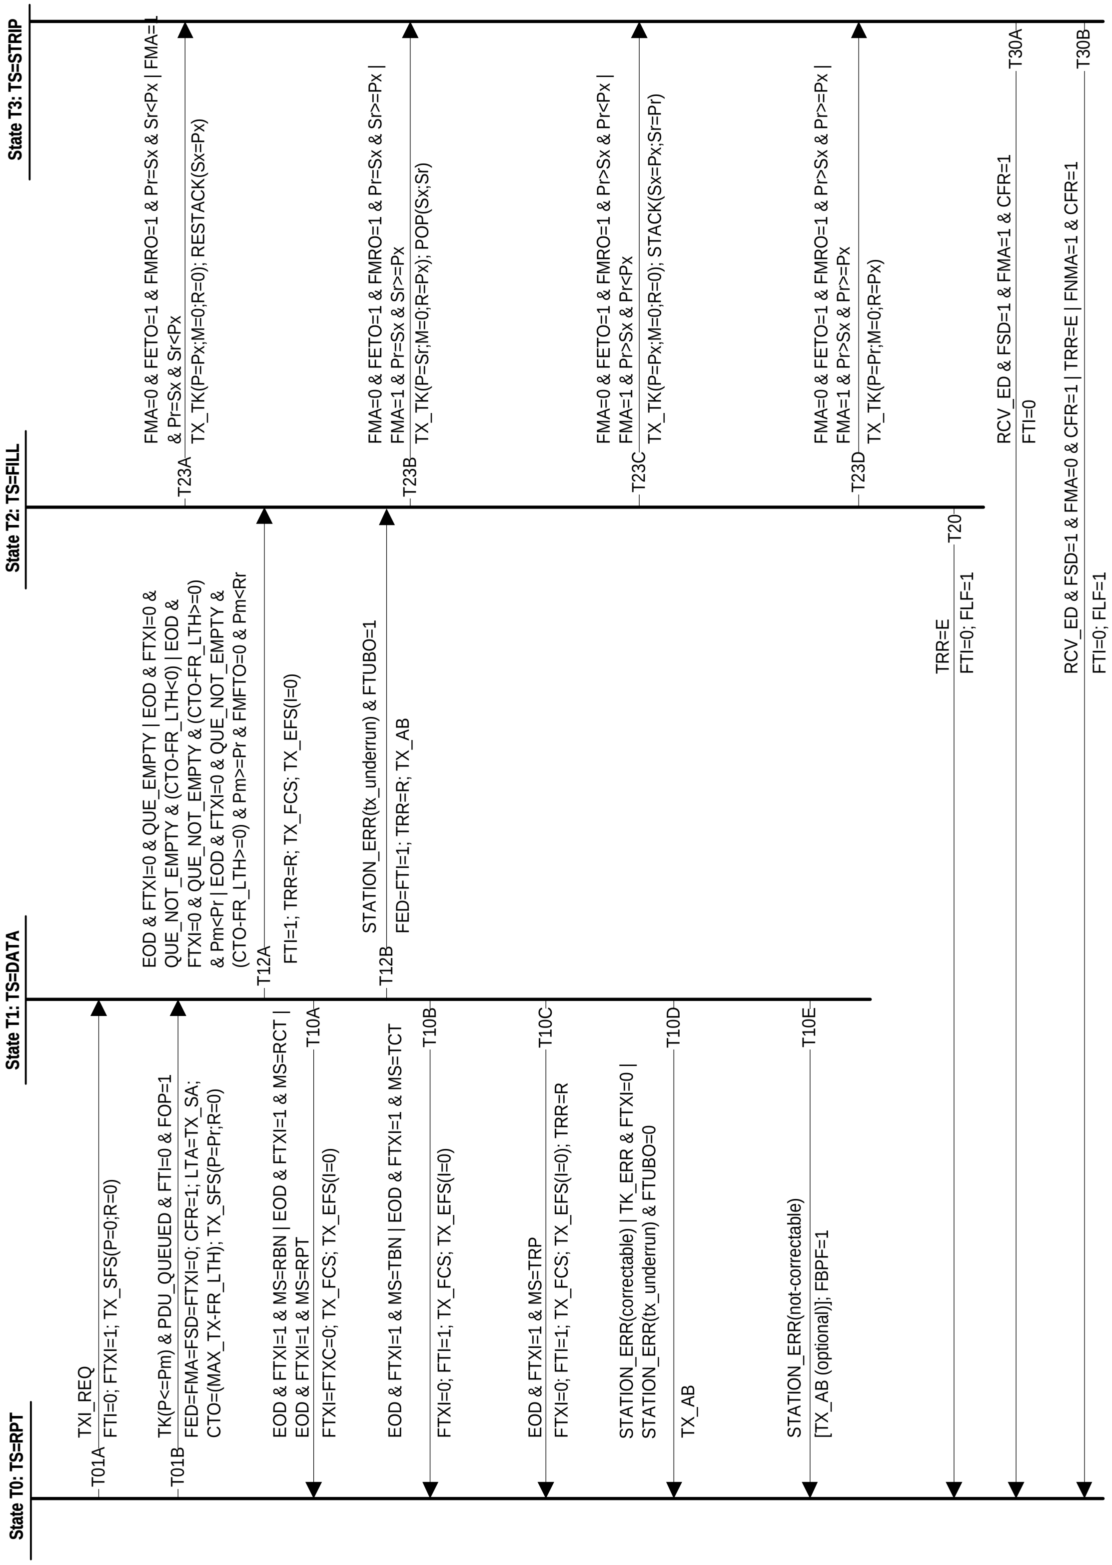

Figure F.3—Transmit (TS=state) finite state machine

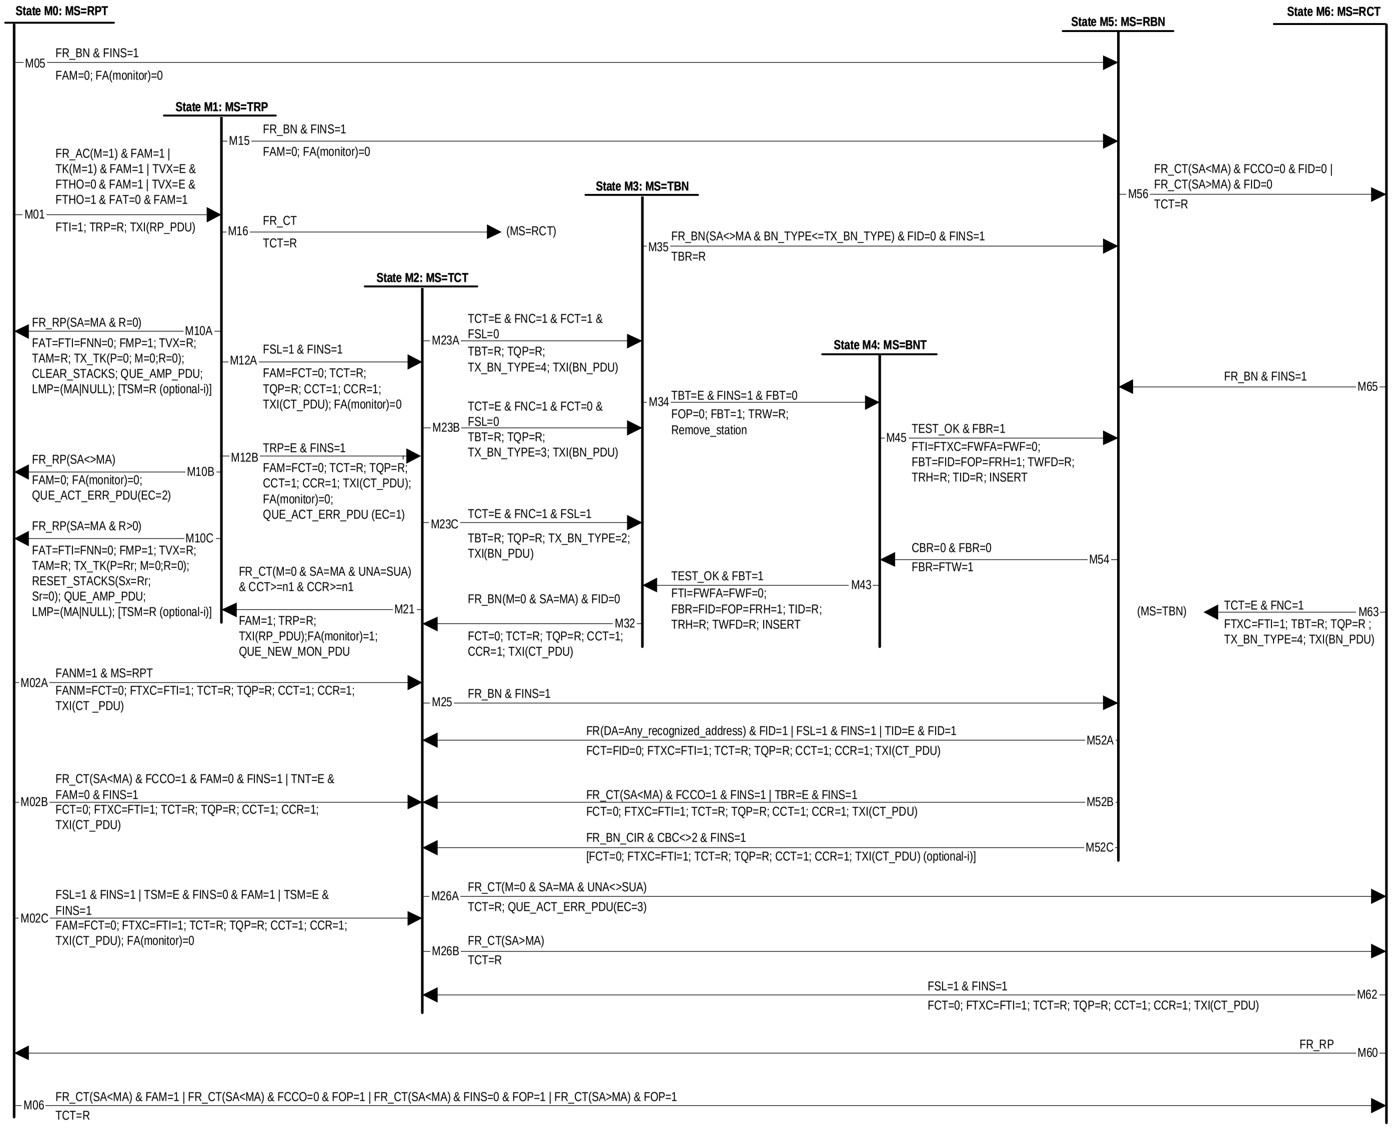

Figure F.4—Monitor (MS=state) finite state machine

Annex G

(informative)

Differences from 1992 version

G.1 Overview

The purpose of this annex is to identify major differences between ISO/IEC 8802-5 : 1992 and this version of the International Standard, providing a brief description and rationale for the change. These changes were made initially in the 1995 edition of this International Standard, and have been carried over into this edition.

G.2 MAC layer differences

Growth of specification: In addition to adding specification for clarity and enhancing operation from the previous version, there were many functions that were not obvious (because they were implied or stated briefly in the text). The processes were made more explicit by adding detail to the FSMs.

G.2.1 Partitioning of state machines

In the previous version of the document there was a separate Standby Monitor FSM and Active Monitor FSM with the Join sequence part of the Standby Monitor FSM. This version separates the Join functions into a separate FSM and combines the standby and active monitor functions into a single FSM in order to properly represent the processes and because the functions of the monitor process are independent of the Join state.

G.2.2 Enhancements for Join FSM

The Join sequence now explicitly requires a station to test its lobe prior to insertion into the ring. It no longer allows a station to join the ring without completing the duplicate address tests (DAT). It enhances the DAT to be tolerant of single bit errors by allowing more than one test frame. It provides for delays in TCU switching (TRW, TID, and TRH). It prevents stations from participating in neighbor notification until DAT is completed and requires neighbor notification before participation in the beacon process is allowed. Flag FBHO is used to select where in the join process a station is allowed to participate in beaconing. The requesting of parameters is now explicitly stated in the FSM. Timer TJR has been defined (replacing TSM in the join action) to monitor progress of the join process. Timer TRI is used to pace the request initialization frames. Action on wire fault has been added.

G.2.3 Enhancements for the Monitor FSM

The beacon process has been enhanced. Quicker resolution or purge has been provided for. Added specification for implied functions and events (e.g., error counting). Provided for assured delivery of error reports. Flag FTEO is used to differentiate between stations that do and do not isolate errors in tokens. Flag FTHO is used to differentiate between different methods of monitoring for token validity (timer TVX).

G.2.4 Modification to the claiming process

Allowing stations to always contend (FCCO) for active monitor (unless the station was the active monitor). Preventing a bad monitor from monopolizing the active monitor role by inhibiting of the active monitor from claiming unless it is the one that detects failure of AMP frames to propagate around the ring (TSM). Detection of a station stuck in claim. The use of TCT and TRP (instead of TNT) to expedite ring recovery.

G.2.5 Enhancement to beacon process

The requirement for stations in the fault domain to remove from the ring and perform self-test is explicitly stated. Stations now monitor the ring for correct beacon process and detect when the beacon transmitter is deceased. Protection against no-owner beacon frames continually circulating the ring was added.

G.2.6 Enhancements to the transmit FSM

Enhanced operation of frame counter (CFR) to count whole frames instead of just counting ending delimiters and to be able to compensate for burst errors (FBEO). Requirements for interframe gap have been added as well as the limitations on altering the interframe gap. Transmit immediate function was added. Enhancements to compensate for overstripping and to prevent understripping were added.

G.2.7 Filtering

Filtering of signal loss (TSL) and wire fault (TWF) are specified to prevent false actions.

G.2.8 Additional differences are

a) Increased restriction on undefined frame formats
b) Removed implied action on ZZZZZZ bits (FC field)
c) Removal of 2-octet addressing
d) The use of counter CTO to enforce the token holding time (4.2.4.1)
e) Pacing of beacon and claim frames
f) The use of CBR to enforce TBR
g) Removal of hierarchical addressing
h) Use of timer TSM by the active monitor
i) Priorities of MAC frames changed to accommodate high performance priority operation.
j) Flag, operational (FOP) was added to specify when the station is required to perform the MAC protocol.

G.3 PHY layer differences

The changes to the PHY clauses have been made to include operation on UTP at 4 and 16 Mbits (replacing the 4 Mbit U recommended practice, ISO/IEC 10738 : 1993), to better specify the limits on jitter in large rings, to clarify ring access control specifications, to add passive and active concentrators, and to include more definition of the channel. In addition, four informative annexes were added to help interpret the PHY layer requirements.

The overall organization of clauses 5 and 7 of the previous version has been maintained in the present version with these exceptions:

a) Clause 5 now includes the MAC-PHY interface definitions formerly found in clause 6, with changes to reflect the MAC differences.
b) Clause 5 gives an example of a station's data path to illustrate MAC-PHY interaction.
c) Clause 7 distinguishes media-independent and media-dependent specifications, as UTP systems differ from STP systems.

d) Clause 7 has an expanded channel clause to address UTP and concentrators.

e) A new clause 8 is added to specify passive and active concentrators.

G.3.1 UTP operation

Most subclauses of clause 7 include new information to specify UTP operation. The specification of the receiver and transmitter has been updated as well as that for the channel.

A different minimum supported PHYs number is given for UTP at 4 Mbit/s (7.2).

The transmitter output voltage has an added specification for UTP, with a smaller range than STP to increase the available crosstalk margin (7.2.2.3).

The transmitter output waveform is now specified directly with a time-domain template, to allow more straightforward conformance test, to limit variations between transmitters, and to help limit high-frequency components (7.2.2.2).

The return loss of the transmitter (7.2.2.4) and the receiver (7.2.3.2) are specified to limit the jitter due to the impedance discontinuities with the channel.

The receiver jitter tolerance specification now includes an alternate channel with added noise to simulate the effects of channel crosstalk noise (7.2.3.1).

The channel clause includes specifications for UTP (7.2.4.2) and adds a noise environment specification (7.2.4.3) that includes crosstalk effects. The channel now explicitly includes passive and active concentrators (7.2.4.1).

A new MIC for UTP is added and its pinout and mechanical definition are referenced (7.2, 7.2.5.2, 8.1, 8.2).

G.3.2 Jitter specification

To better specify the receiver and transmitter and limit their contribution to jitter in large rings, two overall changes in clause 7 have been made:

a) Some jitter specifications are in “phase slope” units, measured in the time domain, to correlate better with measurement methods and to use a parameter that relates more to interoperability.

b) Accumulated jitter and accumulated uncorrelated jitter are specified for a group of stations rather than one station.

The Transmit Asymmetry (TA) specification is replaced by the TDCD (7.2.2.1), Tdiff01 and Tdiffmax specifications. Tdiff01 and Tdiffmax (7.2.2.2.1) also help interoperability with active concentrators and bound the characteristics of receive equalizers.

The Jitter Bandwidth specification is replaced by accumulated jitter (7.1.1), limiting the maximum phase slope produced by a group of stations responding to cable distortion (attenuation and phase), and by DFAPS, limiting the increase in phase slope of these stations in the presence of crosstalk and other disturbers. A station delay specification is added (7.1.3), as it can impact accumulated jitter produced by a group of stations that includes more than one implementation.

Accumulated uncorrelated jitter is specified in terms of a limit on the alignment error it would produce in a clock recovery circuit rather than simple peak-to-peak jitter amplitude (7.1.2).

Receiver jitter tolerance is now specified with a square wave modulated frequency error, rather than a sine wave modulated phase error. Static frequency error of the jitter source is specified explicitly. The measurement units are ns/ns phase slope, rather than UI jitter amplitude (7.2.3.1).

G.3.3 Ring access control (RAC)

The RAC specification has been expanded to clarify the current and voltage requirements for both the phantom drive and compliant load (7.2.1.1). Fault reporting conditions have been made more explicit, and a specification of detectable fault resistances is added (7.2.1.2).

G.3.4 Concentrators

Explicit specifications for concentrators are added in clauses 7 and the new clause 8. Both passive concentrators and active retiming concentrators are described.

In clause 7, the channel includes either passive or active concentrators (7.2.4.1), with the components of the transmission path described explicitly. Flat loss is specified only for the passive channel and is limited to 2 dB rather than 15 dB to allow improved receiver implementations (7.2.4.2.1).

Clause 8 specifies passive and active retiming concentrators, and includes specifications for lobe ports, for trunk ports (concentrator to concentrator connections), and for RAC.

G.3.5 Channel

UTP is added to the channel specifications in clause 7, and both 100 Ω and 150 Ω are given for channel impedance. Explicit specification of cable types replaces the eye pattern specification previously used (7.2.4.1.3). Insertion loss and flat loss are specified, but with less flat loss allowed and lower insertion loss at 4 Mbit (7.2.4.2.1).

A new subclause is added to specify noise, primarily due to crosstalk in UTP cables (7.2.4.3).

Test channels (for phase slope and jitter tolerance testing) are specified explicitly (7.2.4.4).

G.3.6 Annexes

Annex B gives examples for calculating allowable cable lengths based on typical cable characteristics, concentrators, and transmission paths. It also includes notes on use of 120 Ω cabling.

Annex C is a tutorial on jitter, including a description of sources in a typical system and an example jitter budget calculation.

Annex D illustrates a possible frequency domain description (amplitude and delay) of a transmit filter.

Annex E gives guidelines for safety and environment, with references to other standards.

Annex F presents a detailed FSM diagram for each of the MAC state machines. The method used to designate (number) the state transitions has been changed to indicate both the old state and the new state for correlation with table 6.

Annex H provides a frame parsing algorithm as an example of how to verify the validity of MAC frames.

Annex I provides recommended use of priorities for token ring.

Annex P (new in this Edition) provides bit error rate (BER) criteria for lobe media testing.

Annex H

(informative)

Example of a MAC frame parsing algorithm

H.1 Overview

The purpose of this annex is to present an algorithm that may be used for parsing MAC frames addressed to the station.

H.2 MAC frame syntax checking

The Syntax Checking Algorithm has the following MAC frame requirements:

a) The station must be able to receive a MAC frame (see 3.1.2 for definition).
b) Only one vector is allowed in the received MAC frame.
c) The vector length must agree with the length of the MAC frame information field (see 3.3.2).
d) No required subvectors can be duplicated in the MAC frame.
e) Only the first occurrence of a duplicate optional subvector is processed by the MAC.
f) Unknown optional subvectors are ignored by the MAC.

H.3 Syntax checking algorithm

The following is an ordered list of syntax checks done on a MAC frame:

a) The MAC frame containing a RI field (RII=1) is checked and is discarded if source routing is not supported or when any of the following RI field conditions are detected.
    1) The RI length is equal to 0.
    2) The RI length is odd.
    3) The RI length is longer than that supported by the station with 18 bytes being the minimum that must be supported.

b) The MAC frame is checked to determine if it is long enough to contain the vector length, vector class, and vector ID fields (4 or 6 bytes: VL=2 bytes, VC=1 byte, and VI=1 or 3 bytes, see 3.3.2). When the MAC frame length is not long enough to contain the entire vector, it is handled as follows:
    1) The MAC frame with a source class of X'0' is discarded.
    2) Otherwise, a RSP MAC frame is sent with a response code of “Short_MAC (X'8001')” and the frame is discarded.

    Note that the RSP MAC frame source and destination class field may be invalid since the frame could have been too short to contain them.

c) If the received MAC frame has a destination class of other than X'0', the frame is indicated to the management interface (MGT_UNITDATA.indication) without further syntax checking.

d) If both the destination and source classes are X'0' and the frame contains a RI field, the frame is discarded.

e) The vector length is checked to determine if it is equal to the length of the MAC frame information field. If the length is not equal, then the frame is discarded and, if the source class is not X'0', then a RSP MAC frame is sent with response code of “Vector Length Error (X'8002').”

f) The vector identifier (VI) value is checked and, if greater than X'14', the received MAC frame is handled as follows.

    1) If the station is configured not to indicate destination class X'0' MAC frames to the management interface the frame is discarded and, if the source class is not X'0', a RSP MAC frame is sent with a response code of “Unrecognized Vector ID ('8003').”

    2) If the station is configured to indicate destination class X'0' MAC frames to the management interface, then the MAC frame is indicated to the management interface without further syntax checking.

g) The MAC frame source class field is checked to determine if it is valid for the VI. If the source class is not appropriate for the VI then the frame is discarded and, if the source class is not equal to 0, a RSP MAC frame is sent with a response code of “Inappropriate Source Class (X'8004').”

    Frames with source class X'4' or X'8' are recognized as Network Management MAC frames under the following rules:

    1) If the station receives a MAC frame with source class X'8', the frame is accepted and syntax checking continues.

    2) If the source class is X'04' and the station has function class X'8' enabled, then a RSP MAC frame is sent with a response code of : “Inappropriate Source Class (X'8004').”

h) If the received MAC frame is longer than that supported by the station then the frame is discarded and, if its source class is not X'0', the RSP MAC frame is sent with a response code of “MAC Frame Too Long (X'8009').”

i) Each subvector within the MAC frame information field is checked as follows.

    1) If the subvector length is zero then the frame is discarded and if the source class is not X'0' then a RSP MAC frame is sent with a response code of “Subvector Length Error (X'8005').”

    2) If the subvector identifier is equal to X'09' (Correlator) and its length is not 4, then the frame is discarded and, if the source class is not X'0', then a RSP MAC frame is sent with a response code of “Subvector Length Error (X'8005').”

    3) If an optional subvector is duplicated and it passes the subvector syntax checking, the subvector is ignored.

    4) If the subvector type is optional and not known to the MAC, then it is accepted by the MAC and syntax checking continues.

    5) If the subvector identifier is a duplicate occurrence of a required subvector, then the frame is discarded and, if the source class is not X'0', then RSP MAC frame is sent with a response code of “Subvector Unknown (X'8008').”

    6) If the subvector is valid for this VI, the subvector length is checked to be equal to the prescribed length for this subvector. If not equal, the frame is discarded and, if the source class is not X'0', then a RSP MAC frame is sent with a response code of “Subvector Length Error (X'8005').”

j) The station determines whether any required subvectors are missing. If a required subvector is missing the frame is discarded and if the source class is not X'0' then a RSP MAC frame is sent with a response code of “Required Subvector Missing (X'8007').”

k) If the received MAC frame passes the above vector and subvector syntax checking and the function is not supported or is disabled, the frame is discarded and, if the source class is not X'0', then a RSP MAC frame is sent with a response code of “Function disabled (X'800A').”

After the MAC frame has passed the syntax checks stated above, it is processed as a verified MAC frame by the FSMs specified in clause 4, tables 7.1–7.6.

Annex I

(informative)

Recommended use of priorities for token ring

This annex provides a set of recommendations for the use of token ring access priorities on or across token ring LANs to support delay sensitive (e.g., multimedia) traffic.

I.1 Definition of terms

The following definitions apply to this annex:

transmit latency: The delay time between the application's request for transmit and the station's MAC request for transmit on the ring. This time is variable.

access latency: The delay time between the station's MAC request for transmit and when the station starts to transmit the frame (token wait time). This time is variable.

station latency: The delay time between the application's request for transmit and when the station starts to transmit the frame. Station latency is the sum of transmit latency and access latency. This time is variable.

receive latency: The delay time between the station's reception of a frame's last octet of data and when the receiving application is notified of the frame's presence. This time is variable.

end-to-end (E2E) frame latency: The delay time between the source application's request for transmit and when the destination application is notified that the source application's frame has been received.

bridge latency: The delay time between the bridge's reception of a frame's first octet and when the bridge's MAC starts to transmit the first octet of the forwarded frame on the ring. This time is variable.

end-to-end (E2E) frame jitter: The time delta between different E2E frame latency intervals. E2E frame jitter is caused by variations in the E2E frame latency.

user priority: User priority is specified by the application when requesting transmission of data. User priority is encoded into each frame as the YYY bits in the frame control field and also establishes the request priority Pm (see 4.1.13.1).

requested access priority: Each frame to be transmitted has a requested access priority (Pm). The MAC compares Pm to the priority of the received token to determine if a token may be captured to transmit data. The priority of the transmitted frame (P) can be equal to or less than the requested access priority (Pm).

time sensitive service (TSS): Time-critical traffic such as that associated with playback of audio or video from a server. Information integrity (human perception) is dependent upon avoidance of network delays that impact the continuous playback by the client, or the ability to synchronize different media streams (e.g., audio and video).

real time critical service (RTCS): Time-critical traffic such as that associated with multiparty collaborative connections involving real-time sharing or transfer (two-way traffic). One example is desktop videoconferencing with the exchange of images, video, and/or voice traffic.

I.2 Introduction

Transmission of delay-sensitive traffic over the token ring network introduces both E2E frame jitter and E2E frame latency. E2E frame jitter is primarily caused by contention for a common physical medium. E2E frame latency is principally composed of delays due to shared access and also the variable delay associated with queuing frames.

The token ring priority mechanisms provide control over both the E2E frame jitter and E2E frame latency as follows:

a) User priority allows higher priority frames to be queued such that they are transmitted before any previously queued lower priority frames. Upon reception, higher priority frames may be delivered to the receiving applications before any lower priority frames. As a result, transmit and receive latency can be effectively reduced for higher priority frames.

b) Access priority allows higher priority frames to access the ring before lower priority frames. As a result, access latency can be effectively reduced for higher priority frames.

The use of multiple priorities allows for the management of the latency and jitter of delay-sensitive data. The use of higher user priority and access priority helps reduce station latency which minimizes E2E frame latency and E2E frame jitter.

Mechanisms to control the use of ring bandwidth to prevent exceeding a practical level of utilization are beyond the scope of this part of ISO/IEC 8802, but should be considered an essential extension to ISO/IEC 8802-2 : 1994.

Clause I.3 categorizes data traffic into 3 types. Clause I.4 provides recommendations on the use of token ring priorities to provide the needed mechanism for protecting each category of delay-sensitive traffic.

I.3 Quality of service requirements

Applications have certain throughput and latency requirements that necessitate different quality of service. The categories of latency-related quality of service needs are as follows:

a) Non-time critical service
b) Time sensitive service (TSS)
c) Real time critical service (RTCS)

I.3.1 Non-time critical service

Non-time critical service is defined for traffic whose proper operation is not sensitive to E2E frame latency or E2E frame jitter. Applications used for evaluating non-time critical service traffic are file servers and print servers.

I.3.2 Time sensitive service (TSS)

TSS is defined for synchronized traffic whose proper operation is sensitive to E2E frame jitter. E2E latency is less critical for TSS traffic because buffering at the destination provides the needed compensation for E2E frame jitter. E2E frame latency must be managed because it directly relates to E2E frame jitter. Applications used for evaluating TSS traffic are multimedia servers and video playback servers. These applications tolerate a maximum E2E frame latency of 300 ms.

I.3.3 Real time critical service (RTCS)

RTCS is defined for synchronized traffic whose proper operation is sensitive to both E2E latency and E2E frame jitter. Because of the total latency requirement, less buffering is available for jitter compensation, and therefore E2E frame latency and jitter must be kept to a minimum. Applications used for evaluating RTCS traffic are audio/video conferencing and interactive video. These applications tolerate a maximum E2E frame latency of 100 ms.

I.3.4 Priority and latency relationship

Since non-time critical traffic does not have definite latency requirements, it should have the lowest access priority of the three categories.

For TSS traffic to meet the 300 ms E2E frame latency, stations must have an access latency of 100 ms or less. This requires a higher access priority than non-time critical service traffic.

For RTCS traffic to meet the 100 ms E2E frame latency, the station must have an access latency of 10 ms or less. This requires a higher priority than any other traffic category.

I.4 Priority discussion

This clause includes an overview of token ring priority fields, a station's priority structure, and concludes with station and bridging priority recommendations.

I.4.1 Token-ring priority fields

The token-ring LLC frame has the frame header format shown in figure I.1. The priority operations on the ring are obtained through use of fields within the access control (AC) field and the frame control (FC) field.

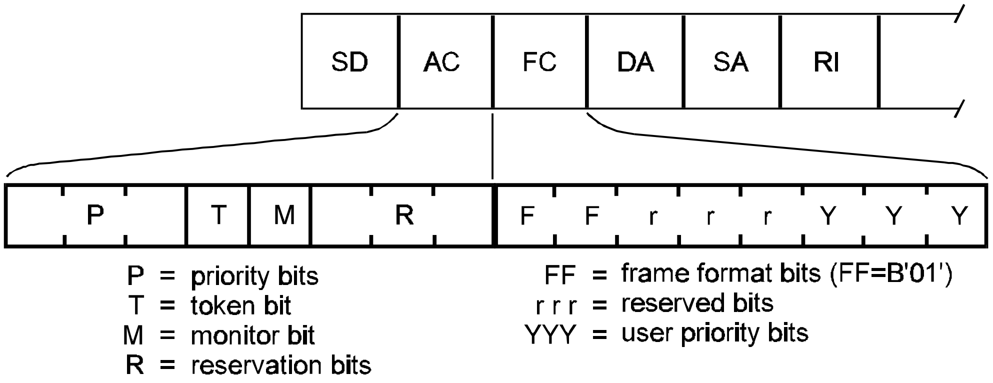

Figure I.1—LLC frame header format

Token priority is specified by the P bits of the ac field within a token or a frame. In combination with the R bits, these bits regulate which station gets access to the token ring.

User priority is specified by the LLC MA_UNITDATA request. It establishes the request access priority and it is conveyed end-to-end by the YYY bits within the FC field. Both the token priority and user priority are encoded such that B'000' is the lowest priority and B'111' is the highest priority.

It is recommended that a user priority of seven (B'111') not be used for LLC traffic.

I.4.2 Station priority structure

Figure I.2 represents an architectural structure that results in a lower latency for higher priority traffic.

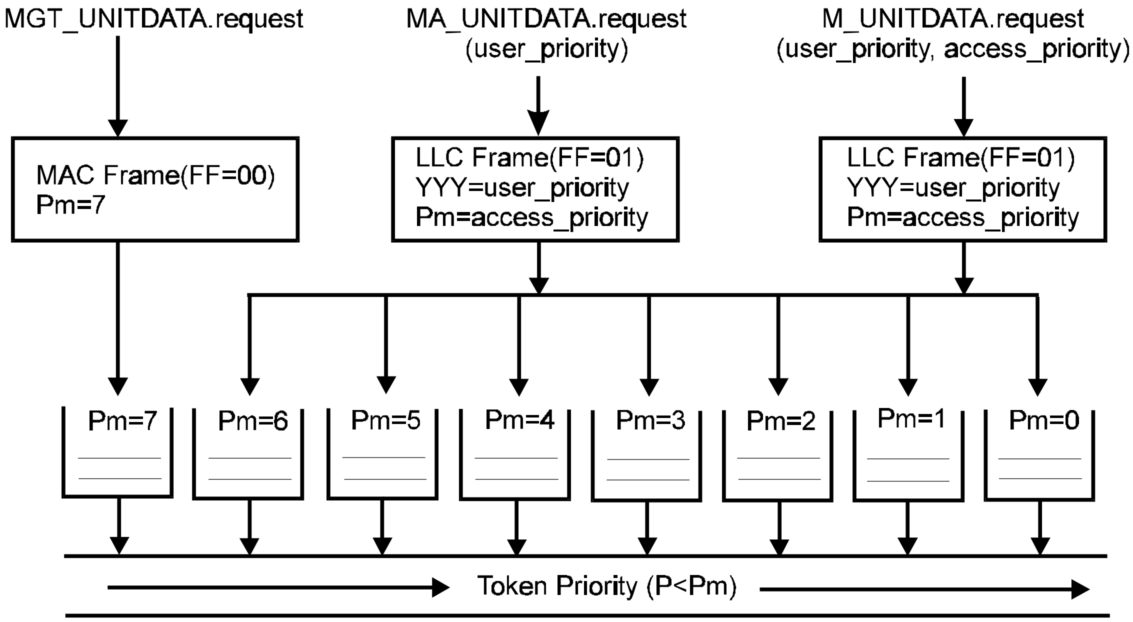

Figure I.2—Priority structure within a station

Each of the eight frame priorities is represented as a separate queue. The assumption is that higher priority frames will be transmitted before lower priority frames, but the frames will not be reordered within a given priority. Furthermore, it is assumed that all non-recovery MAC frames originated by the station will be queued for transmission with a requested access priority (Pm) equal to seven.

The user priority is specified as part of the MA_UNITDATA.request and MA_UNITDATA.indication primitives over the LLC/MAC service interface.

a) User priority from the MA_UNITDATA.request primitive is encoded into the YYY bits of the frame allowing the originator's user priority to be carried by the LLC frame end-to-end.
b) Upon receipt, the frame's YYY bits are used to set the value of the user priority in the MA_UNITDATA.indication primitive.

A station transmits a queued frame when the requested access priority (Pm) associated with the frame is equal to or greater than the priority (P) of the token received.

I.4.3 Station priority recommendation

The following recommendations provide control over E2E frame latency to support delay-sensitive sessions. These recommendations are summarized in table I.1.

a) To minimize E2E frame latency the time sensitive service TSS traffic should be queued with a user priority of five.

b) To minimize E2E frame latency and E2E frame jitter, the RTCS traffic should be queued with a user priority of six.

c) Due to the importance of management traffic, it is further suggested that LAN management LLC traffic be queued with a user priority of four.

d) Clause 3.3.5.1 recommends that all MAC frames transmitted on a token ring be queued for transmission with priority 7. Following this recommendation is advised.

e) User data traffic that is not time-critical should be transmitted with a user priority of zero.

Services for other user traffic employing user priorities 1, 2, and 3 are not discussed in this annex.

Table I.1—Recommended use of priorities by stations

Services offered	User priority (YYY bits)	Requested access priority (Pm)
MAC frame traffic	N/A	7
RTCS traffic	B'110'	6
TSS traffic	B'101'	5
LAN management traffic	B'100'	4
Other user traffic	B'011'	3
Other user traffic	B'010'	2
Other user traffic	B'001'	1
Other user traffic	B'000'	0

The following recommendations help reduce the frame jitter associated with high-priority traffic:

a) Stations that employ multiple frame transmission per token should honor reservations for a higher priority token as provided by a FMFTO value of 0 (see 3.5.6), thus expediting the release of a token at the frame’s reservation priority (from the frame’s AC field R bits).

b) The maximum information field size for all traffic on the token ring should be constrained to 4399 octets when delay-sensitive traffic is traversing the ring.

I.4.4 Bridging priority recommendations

Bridge latency for delay-sensitive traffic should be less than non-delay sensitive traffic to minimize the delay-sensitive traffic’s E2E frame latency. ISO/IEC 10038 : 1993 allows token ring MAC bridges to forward all frames with an outbound access priority equal to the default Pm value of 4 or with a Pm based on the user priority encoded in the FC field YYY bits.

This annex further recommends that LLC frames with user priorities greater than 4 should be forwarded with an access priority equal to the user priority as shown in table I.2. No specific recommendations are made for LLC frames with user priorities less than 4.

Table I.2—Recommended use of priorities by token ring bridges

Received user priority (YYY Bits)	Transmitted requested access priority (Pm)	Assumed service
N/A	7	MAC frame traffic
B'110'	6	RTCS traffic
B'101'	5	TSS traffic
B'100'	4	LAN management/bridged traffic
B'011'	3 or 4	Other user traffic
B'010'	2 or 4	Other user traffic
B'001'	1 or 4	Other user traffic
B'000'	0 or 4	Other user traffic

Annex J

(informative)

Bibliography

[B1] ASTM D4566-94, Test methods for electrical performance properties of insulations and jackets for telecommunications wire and cable.[35]

[B2] IEC 60060, High-voltage test techniques.[36]

[B3] IEEE Std 802.1F-1993, IEEE Standards for Local and Metropolitan Area Networks: Common Definitions and Procedures for IEEE 802 Management Information (ANSI).[37]

[B4] IEEE Std 802.5c-1991 (Reaff 1997), IEEE Recommended Practice for Dual Ring Operation with Wrapback Reconfiguration (ANSI).

---

[35] ASTM publications are available from the Customer Service Department, American Society for Testing and Materials, 1916 Race Street, Philadelphia, PA 19103, USA.

[36] IEC publications are available from IEC Sales Department, 3 rue de Varembé, Case Postale 131, CH-1211, Genève 20, Switzerland/Suisse. These publications are also available in the United States from the Sales Department, American National Standards Institute, 11 West 42nd Street, 13th Floor, New York, NY 10036 USA.

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

Annex P

(normative)

Bit error rate (BER) criteria for lobe media testing

This annex presents analysis that can be used when designing lobe media testing for the Token Passing Access Protocol (TKP) or Transmit Access Protocol (TXI) operations. It is based on the assumption that frame error rate affects token ring operation as if the BER were a Gaussian variable with probabilities in the range of the calculations. Although line noise will often produce a burst error rather than a single error, each burst generally corrupts only a single frame. Therefore, the noise-induced frame error rate will be approximately the same as the frame error rate, assuming that each error-producing noise event produces only a single bit error. The analysis below derives the error rate based on this single bit error per event assumption, and can be applied to line noise error rate observed in token ring networks.

The BER requirement for a station operating on a conformant link and attached to a conformant concentrator is specified in clause 7, and is not specifically validated by the lobe media test (LMT). The purpose of the LMT is twofold:

a) The LMT must prevent the network attachment of stations whose BER would adversely affect the network performance.
b) The LMT must prevent rejection of a station that would meet the BER criteria if attached.

A further requirement of the lobe media test is that it must not unduly delay attachment of the station to the network.

These criteria are met by choosing an LMT pass/fail criteria of BER = 10^{-6}, where links with a BER of 10^{-7} or better would have a ≥99% chance of passing the test, and links with a BER of 10^{-5} or worse would have a ≥99% chance of failing the test. During the LMT, the station shall pass, with a probability of >99%, a lobe that has a sustained bit error rate of 10^{-7}, or better, and shall fail, with a probability of >99%, a lobe that has a sustained bit error rate of 10^{-5}, or worse. The following analysis provides the mathematical basis for designing an appropriate LMT.

Bit Error Rate Testing: Analysis

Given an LMT that employs “m” frames, each with an equal length of “n” bits, and that the test passes if no more than a single frame contains errors.

What is the probability of the LMT failing that test given

- All bit errors are independent events, and
- The probability of a single bit error is q.

Analysis

The probability of completing the test without detecting any bit errors = $(1-q)^{nm}$.

The probability of a *single* frame being in error = 1 – (probability that it arrives without error): Probability_of_error = $1-(1-q)^{n}$; i.e., the first frame has an error.

The probability that the rest of the frames are error free is the probability that there are no errors in the next n(m–1) bits:

The probability of having m–1 frames without error = $(1-q)^{n(m-1)}$.

Therefore, the probability that the first frame has at least one error, and the rest of the frames are error free =

$$(1-(1-q)^n)(1-q)^{n(m-1)}$$

This expression is also equal to the probability that all frames but the second are error free, = probability that all frames but the third are error free.

Therefore, the probability that exactly one frame has errors is the sum of these m terms, 1 (one) for each frame, and the probability of successfully completing the test with a single frame having an error is equal to

$$m(1-(1-q)^n)(1-q)^{n(m-1)}$$

Therefore, the total probability of completing the tests successfully is

$$(1-q)^{nm} + m(1-(1-q)^n)(1-q)^{n(m-1)}$$

Let p = 1–q and the result is

$$\begin{aligned} p^{nm} + m(1-p^n)p^{n(m-1)} &= p^{n(m-1)}(p^n + m(1-p^n)) \\ &= p^{n(m-1)}(m-(m-1)p^n) \end{aligned}$$

This equation can be plugged into a spreadsheet for various values of n, m, and p to show what is needed to have the probability of passing the test of ≥99% if BER is $\leq 10^{-7}$, and the probability of failing the test of ≥99% if BER $\geq 10^{-5}$.

There is a broad range of solutions that satisfies these criteria. This range includes LMTs of one million bits in 256-bit to 1024-bit equal length frames with a pass/fail criteria of

— The lobe media test passes if 0 or 1 frames are with errors.
— The lobe media test fails if 2 or more frames are with errors.

NOTE—For this analysis, a frame is not considered to be with error if the error is due to a token with error. This condition is required to satisfy the initial assumption of this analysis that each frame has an equal and known length. Circulating tokens would contribute an unknown number of bits to each frame transmitted, and might cause LMT to fail a good lobe.